

Essays on Moral Realism

By

Alexander Hyun

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy
(Philosophy)

at the

UNIVERSITY OF WISCONSIN-MADISON

2017

Date of final oral examination: 7/14/2017

This dissertation is approved by the following members of the Final Oral Committee:

Russ Shafer-Landau, Professor, Philosophy

Michael Titelbaum, Associate Professor, Philosophy

John Bengson, Associate Professor, Philosophy

Sarah Paul, Associate Professor, Philosophy

Paul Kelleher, Associate Professor, Medical History and Bioethics

© Copyright by Alexander B. Hyun 2017
All Rights Reserved

Contents

Acknowledgments	iii
Abstract	iv
Introduction.....	1
0.1 The Nature of Moral Realism.....	1
0.11 Descriptive Cognitivism	1
0.12 Success Thesis	2
0.13 Objectivity Thesis	3
0.2 The Plan.....	5
0.21 Chapter 1.....	5
0.22 Chapter 2.....	6
0.23 Chapter 3.....	7
Chapter 1: Moral Realism, Practical Reasons, and the Argument from Epistemic Duties... 8	
1.1 Introduction	8
1.2 The Argument from Epistemic Duties	11
1.2.1 Defense of Premise (4)	14
1.2.2 Defense of Premise (1)	21
1.2.3 Defense of Premise (2)	24
1.3 Critique of Arguments for Reasons Internalism	33
1.3.1 Manne’s Reasons Internalism.....	33
1.3.2 Why Manne’s Argument for MI Probably Goes Wrong Somewhere	34
1.3.3 Identifying the Flaw in Manne’s Argument for MI.....	38
1.3.4 Markovits’ Reasons Internalism.....	41
1.3.5 Markovits’ Analogical Argument.....	42
1.3.6 Markovits’ Argument from Rationality.....	46
1.3.7 Markovits’ Argument from Motivational Disparity.....	53
1.4 Conclusion.....	60
Chapter 2: A Defense of the Companions in Guilt Argument for Moral Facts..... 61	
2.1 Introduction	61
2.2 The Companions in Guilt Argument.....	61
2.2.1 The Internalist Argument for Moral Error Theory	63
2.2.2 The Argument from Irreducible Normativity	65

2.2.3 The Supervenience Argument	70
2.2.4 The Epistemological Queerness Argument	80
2.3 Heathwood’s Objection from Epistemic Naturalism	82
2.4 Ingram’s Objection from Epistemic Non-Cognitivism.....	87
2.5 Streumer’s Objection from Global Normative Error Theory.....	97
2.6 Conclusion.....	102
Chapter 3: Pure Moral Deference, Social Goods, and Fair Play.....	103
3.1 Two Puzzles of Pure Moral Deference	103
3.2 Against Anti-Realist Accounts.....	111
3.2.1 Simple Non-Cognitivism.....	111
3.2.2 Sophisticated Non-Cognitivism.....	112
3.2.3 Subjectivism	116
3.2.4 Constructivism.....	119
3.3 A Dilemma for the Rebutting Puzzle of Pure Moral Deference	122
3.4 A Realist-Friendly Account	127
3.4.1 The Social Goods Account.....	129
3.4.2 The Fair Play Account.....	138
3.5 Conclusion.....	142
Bibliography	143

Acknowledgments

I am extremely grateful to the professors in the philosophical community at UW-Madison for creating a terrific environment in which to grow as a philosopher. Special thanks are due to Russ Shafer-Landau, Mike Titelbaum, and John Bengson. Thanks to Russ for providing stimulating seminars that convinced me to specialize in meta-ethics, for being an unsurpassable model of philosophical rigor and clarity, for being unfailingly friendly and professional, and for providing helpful and prompt feedback on all of my dissertation work. And I thank Mike and John for pushing me to be a deeper thinker and to be unsatisfied with producing superficial argumentation.

Thanks also to my philosophical peers over the last decade for teaching me, challenging me, and inspiring me. For playing one or more of these roles, I especially want to thank Carlos Medina, Jason Zarri, Hayley Clatterbuck, Nic Owen, Kevin Dewan, Casey Hart, Justin Morton, Eric Sampson, Justin Horn, Jeff Behrends, Daniel Crow, Leigh Vicens, David O'Brien, and Ben Schwan.

I am grateful to my Bible Study group for helping me to turn my mind habitually towards matters that are more important than normative reasons and moral realism, thereby helping me to keep some much-needed perspective in graduate school. And I thank my parents, John and Charlotte Hyun, for their unwavering support of my endeavor to become an excellent philosopher.

Finally, I thank my wife, Megan, for her grace, love, and encouragement throughout the latter half of my graduate school years. Words feel entirely insufficient to express the depths of my gratitude for her.

Abstract

This dissertation is a defense of moral realism. By *moral realism*, I mean the conjunction of three claims: (i) *Descriptive cognitivism*, according to which moral judgments are descriptive beliefs that aim to represent the world accurately; (ii) *The success thesis*, according to which some moral judgments are true; and (iii) *The objectivity thesis*, according to which the true moral judgments are objectively true, in the sense that their truth does not constitutively depend on the attitudes of some actual or idealized agent. The purpose of my dissertation is to argue in favor of the success and objectivity theses.

In Chapter 1, I argue in favor of externalism about normative reasons, thereby defending both the success and objectivity theses from influential objections. Roughly, externalism about normative reasons states that there are some external reasons for action, i.e., reasons to do some act that do not depend on the desires of the agent whose reasons they are. I argue for externalism by appealing to epistemic normativity. While others have appealed to epistemic normativity to defend externalism, such appeals are normally aimed at undermining arguments against externalism. In contrast, I develop a more ambitious use of epistemic normativity that aims to provide a direct argument for the truth of externalism. Specifically, I argue that there exist *practical epistemic facts* – facts to the effect that we epistemically ought to perform certain actions – and that these facts entail the existence of external reasons for action. I also bolster this argument for externalism by seeking to refute the formidable challenges to externalism that have recently been offered by Kate Manne and Julia Markovits. Much of my rebuttal of Manne’s argument is now reproduced in my ‘Internalism, Ideal Advisors and the Conditional Fallacy,’ which appears in *The Journal of Ethics and Social Philosophy*.

In Chapter 2, I defend a version of the increasingly influential ‘companions in guilt’ argument for moral facts, thereby establishing the success thesis. My favored version of this argument goes as follows: (1) If there are no moral facts, then there are no practical epistemic facts; (2) there are practical epistemic facts; (3) so, there are moral facts. The second premise, which is known as the ‘Ontological Premise,’ is defended at length in Chapter 1. I offer a presumptive case for the first premise, which is known as the ‘Parity Premise,’ by arguing that the four most formidable arguments against moral facts suggest equally-plausible arguments against practical epistemic facts. I then argue that my argument’s atypical appeal to practical epistemic facts allows it to address recent objections to the companions in guilt argument that have been offered by Christopher Heathwood and Stephen Ingram.

In the third and final chapter of my dissertation, I respond to the ‘puzzle of pure moral deference,’ a challenge to the objectivity thesis that has been most forcefully pressed by Sarah McGrath. According to this challenge, moral anti-realism can explain why moral deference seems intuitively problematic to many of us, whereas moral realism cannot explain why this is so; and we therefore have reason to accept anti-realism instead of realism.

I develop three independent rebuttals to this challenge. First, I object to the four main anti-realist accounts of our discomfort with moral deference, thereby undermining the claim that moral anti-realism provides an explanation of this discomfort. Second, I develop a dilemma for the proponent of the puzzle of pure moral deference, arguing that either the anti-realist cannot provide the needed explanation, or else the realist can do so. Finally, I offer a novel, realist-friendly account of our discomfort with moral deference that builds on extant realist accounts. In brief, I argue that a lot of people’s discomfort can plausibly be explained by appealing to the fact that moral deference is both unfair and bad for society.

Introduction

This dissertation is a defense of moral realism – roughly, the view that there are objective moral truths. In this introduction, I first explain the nature of the version of moral realism that I will be defending. I then explain how each of the main sections of this dissertation contributes to the case for moral realism.

0.1 The Nature of Moral Realism

0.11 Descriptive Cognitivism

A *moral sentence* is a sentence that, in virtue of its surface logical and grammatical form, appears to predicate a moral property to something. For example, the following are moral sentences: ‘Telling the truth was the morally right thing to do,’ ‘Eating meat is morally wrong,’ and ‘Hitler is evil.’ *Moral judgments* are the states of mind – whatever they are – that are expressed by utterances of moral sentences. The first commitment of moral realism is a thesis about the nature of moral judgments:

Descriptive Cognitivism: Moral judgments are descriptive beliefs.

By a *descriptive belief*, I mean a belief with representational content; and by *representational content*, I mean content that represents the world as being a certain way, or in other words, as being such that some fact obtains within it. A belief’s representational content is true if and only if it represents the world as being a certain way and it really is that way; and a descriptive belief is true if and only if it has true representational content. Descriptive cognitivism therefore entails that

moral judgments are capable of being true or false, and that the truth value of a given moral judgment is determined by whether it accurately represents the moral facts.¹

To illustrate this position, suppose that I say, ‘Hitler is evil.’ The descriptive cognitivist holds that I have expressed a descriptive belief with representational content that represents the world as being such that it is a moral fact that Hitler is evil. This descriptive belief is true if and only if it is a moral fact that Hitler is evil.

In adopting descriptive cognitivism, moral realists distinguish themselves from *non-cognitivists*, who hold that moral judgments are not descriptive beliefs, but rather are desire-like mental states. Desires do not have content that represents the world as being a certain way, and consequently they are not truth apt in the same way that descriptive beliefs are truth apt. According to the non-cognitivist, moral judgments are mental states that resemble desires in these respects. For instance, various defenders of non-cognitivism have proposed that moral judgments are approvals and disapprovals, emotions, acceptances of norms, acceptances of plans, and non-descriptive beliefs.²

0.12 Success Thesis

Those who accept descriptive cognitivism fall into one of two camps. *Moral error theorists* are descriptive cognitivists who hold that there are no moral facts for moral judgments to represent. Moral error theorists therefore hold that all moral judgments are untrue.³ In contrast, realists adopt:

Success Thesis: Some moral judgments are true.

¹ Defenses of this popular view of moral judgments include Enoch (2011), Huemer (2005), Shafer-Landau (2003), and Smith (1994).

² Ayer (1952) argues that moral judgments are attitudes such as approvals, disapprovals, and emotions; Gibbard (1990) and Gibbard (2003) argue that moral judgments are acceptances of norms and acceptances of plans, respectively; and Horgan and Timmons (2000) argues that moral judgments are *non-descriptive beliefs*, or beliefs without representational content. Some non-cognitivists, such as Blackburn (1993), hold that moral judgments are truth apt, but only in a ‘minimal’ way.

³ This sort of moral error theory is defended by Joyce (2001, 2006), Mackie (1977), and Olson (2014), among others.

In adopting the success thesis, moral realists take on a commitment to moral facts that our moral judgments sometimes succeed in representing.

0.13 Objectivity Thesis

Those who accept both descriptive cognitivism and the success thesis can be divided into two camps: moral subjectivists and moral objectivists. To explain the nature of these positions, it will be helpful to introduce a distinction between subjective properties and objective properties. A property P is a *subjective property* if and only if part of what it is for a thing to have P is for some agent either (i) to have some attitude or reaction to that thing, or (ii) to be such that they would have some attitude or reaction to that thing in certain counterfactual conditions.⁴ As Michael Huemer observes, funniness is an example of a subjective property because what it is for something to be funny is roughly for it to be the case that others are disposed to laugh at it and be amused by it. An *objective property* is a property that is not a subjective property. Squareness is an example of an objective property, for it is not the case that part of what it is for something to be square is for some agent to have some (actual or counterfactual) attitude or reaction to it.

This distinction between subjective and objective properties suggests an analogous distinction between subjective and objective facts. A *subjective fact* is a fact to the effect that something has some subjective property. An *objective fact* is a fact to the effect that something has some objective property.

As I mentioned above, descriptive beliefs represent the world as being such that some fact obtains within it. We can say that a descriptive belief is *subjectively true* if and only if it is true and the fact that it accurately represents as obtaining is a subjective fact. A descriptive belief is *objectively true* if and only if it is true and the fact that it successfully represents as obtaining is an

⁴ I am drawing on Huemer (2005: 2-3)'s helpful elucidation of the notion of subjective properties.

objective fact. With the above terminology in place, we can understand the moral realist's final commitment as follows:

Objectivity Thesis: The true moral judgments are objectively true.

Assuming the truth of the moral judgment that Hitler is evil, the proponent of the objectivity thesis adopts both of the following: (i) this moral judgment is a descriptive belief that represents the world as being such that the moral fact of Hitler's being evil obtains; and (ii) it is not the case that part of what it is for Hitler to be evil is for some agent to have some (actual or counterfactual) attitude or reaction to Hitler. Claim (ii) is sometimes expressed by saying that Hitler's being evil does not 'constitutively depend' on the attitudes or reactions of agents towards Hitler.⁵

In adopting the objectivity thesis, the moral realist separates herself from meta-ethical versions of subjectivism, cultural relativism, constructivism, and divine command theory. Contrary to these views, proponents of the objectivity thesis hold that moral truths are not constituted by the individual's subjective feelings, culture's attitudes of approval and disapproval, the attitudes that some actual or hypothetical agent would have upon completing some type of constructive procedure, or God's attitudes and commands.

For the purposes of this dissertation, *moral realism* is the conjunction of descriptive cognitivism, the success thesis, and the objectivity thesis. The position I seek to defend is therefore opposed to classical non-cognitivism, expressivism, error theory, and all subjectivist theories.

⁵ Proponents of the objectivity thesis include moral realists from both naturalistic and non-naturalistic camps, including Boyd (1988), Cuneo (2007), Enoch (2011), Huemer (2005), Railton (1986), and Shafer-Landau (2003, 2006).

0.2 The Plan

Each of the three chapters that together comprise this dissertation aims to defend one or more of the components of moral realism. As I explain below, Chapter 1 defends the success and objectivity theses from influential objections; Chapter 2 provides a positive argument in support of the success thesis; and Chapter 3 defends the objectivity thesis from a challenge that arises from the recent debate over moral deference. In this dissertation, I take for granted the truth of descriptive cognitivism in order to focus on defending moral realism on other fronts.⁶

0.21 Chapter 1

Internalism about normative reasons – roughly, the view that all reasons for action are dependent on the desires of the agent whose reasons they are – is the basis for two influential challenges to moral realism. First, John Mackie and Richard Joyce appeal to internalism in the course of defending moral error theory, which is incompatible with moral realism’s success thesis. Second, Gilbert Harman appeals to internalism in the course of defending his brand of moral relativism, which conflicts with moral realism’s objectivity thesis.

In Chapter 1, I defend moral realism from these two objections by defending externalism about reasons for action, which is the view that there do exist reasons of the sort ruled out by internalism. I defend such desire-independent reasons by appealing to epistemic normativity. While others have appealed to epistemic normativity to defend externalism, such appeals are normally aimed at undermining arguments against externalism. In contrast, I develop a more ambitious use of epistemic normativity that aims to provide a direct argument for the truth of

⁶ Although I will not be defending descriptive cognitivism directly, my second chapter provides indirect support for this view. In Chapter 2, I defend an argument for the existence of moral facts. Suppose that this argument succeeds. Since moral facts exist, embracing descriptive cognitivism does not commit one to believing that ordinary moral discourse is radically flawed. And this means that an important consideration in favor of non-cognitivism – namely, that it is necessary to avoid this level of skepticism about ordinary moral discourse – is refuted. So, the argumentation I provide in Chapter 2 indirectly undermines non-cognitivism and supports descriptive cognitivism. I borrow this point from Cuneo (2007: 42-44).

externalism. Specifically, I argue that there exist some *practical epistemic facts* – facts to the effect that we epistemically ought to perform certain actions – and that these facts entail the existence of external reasons for action. I also bolster this argument for externalism by seeking to refute the formidable challenges to externalism that have recently been offered by Kate Manne and Julia Markovits.

0.22 Chapter 2

In Chapter 2, I establish moral realism’s success thesis by arguing in favor of the existence of moral facts. My strategy will be to defend a companions in guilt argument. Proponents of this type of argument seek to show that moral facts are importantly similar to epistemic facts whose existence are not in question, and that we should therefore believe in moral facts.⁷ While others have appealed to epistemic facts to defend moral facts, they almost always appeal to *doxastic epistemic facts* – facts to the effect that some doxastic attitude possesses some epistemic property, such as epistemic justification. In contrast, the version of this argument that I will defend goes as follows:

- (1) **Parity Premise:** If moral facts do not exist, then practical epistemic facts do not exist.
- (2) **Ontological Premise:** Practical epistemic facts do exist.
- (3) Therefore, moral facts exist.

By the end of Chapter 1, I have already defended the Ontological Premise at length. My primary task in Chapter 2 is to defend a presumptive case for the Parity Premise by arguing that the four most formidable arguments against moral facts suggest equally-plausible arguments against practical epistemic facts. After I do this, I consider recent objections to companions in guilt

⁷ It would be more apt to refer to this type of argument as a companions in *innocence* argument since it tries to identify companions for moral facts that have impeccable existential credentials; but I will use the customary designation in what follows.

arguments that have been developed by Christopher Heathwood and Stephen Ingram, arguing that my argument's atypical appeal to practical epistemic facts allows it to withstand these objections.

0.23 Chapter 3

In the third and final chapter, I respond to the 'puzzle of pure moral deference,' a challenge to the objectivity thesis that has been most forcefully pressed by Sarah McGrath. Proponents of this challenge begin by observing that moral deference strikes many of us as off-putting—a surprising fact in light of our relative comfort with deference about most other matters. They then argue that moral anti-realism alone can explain this observation, and that we therefore have reason to accept anti-realism instead of realism.

I develop three independent lines of response against this challenge. First, I object to the four main anti-realist accounts of our discomfort with moral deference, thereby undermining the claim that moral anti-realism provides an explanation of this discomfort. Second, I develop a dilemma for the proponent of the puzzle of pure moral deference, arguing that either the anti-realist cannot provide the needed explanation, or else the realist can do so. Finally, I offer a novel, realist-friendly account of our discomfort with moral deference that builds on the extant realist accounts. In brief, I argue that a lot of people's discomfort can plausibly be explained by appealing to the fact that moral deference is both unfair and bad for society.

Chapter 1: Moral Realism, Practical Reasons, and the Argument from Epistemic Duties

1.1 Introduction

Normative reasons for action are facts that favor some agent's performing some act or refraining from performing some act.⁸ For example, the fact that one is hungry is normally a normative reason to eat something. And the fact that putting one's hand in the fire would cause one terrible agony is a normative reason not to do so. One popular family of theories about normative reasons for action is *reasons internalism*.⁹ According to reasons internalism, it is a necessary condition on there being a reason for an agent to do some action that she have a set of desires of the appropriate sort.¹⁰ Different reasons internalists hold different views about what the 'appropriate sort' of desires amounts to. For example, a very simple version of reasons internalism holds that there is a normative reason for an agent to ϕ only if her set of desires includes a desire to ϕ .¹¹ A more sophisticated version states that there is a normative reason for an agent to ϕ only if her set of desires includes a desire the satisfaction of which is promoted by her ϕ -ing.¹² And Bernard Williams' influential version states that there is a normative reason for an agent to ϕ only if her set of actual desires is such that there is a 'sound deliberative route' from her set of desires to the conclusion to ϕ .¹³ What unites these views—what makes them all versions of reasons internalism—is that they all hold it to be a necessary condition on there being a normative reason

⁸ While I will assume that reasons are *facts*, nothing important in the arguments that follow depends on this assumption.

⁹ Other labels for this popular theory include 'the desire-based theory,' 'the Humean theory,' and 'subjectivism.'

¹⁰ Sometimes, reasons internalists are characterized as holding that there is a necessary link between having a reason to do an action and *being motivated* to do that action. For the purposes of this paper, I wish to focus on versions of reasons internalism that posit a necessary connection between reasons and *desires*.

¹¹ Enoch (2011: 260) calls this view 'Really Strong Existence Internalism.'

¹² Williams (1979: 102) calls this view 'the sub-Humean model.'

¹³ See Williams (1979, 1989). For still other versions of reasons internalism, see Hubin (2003), Manne (2014), Markovits (2014), Schroeder (2007), Smith (1994, 1995), and Tubert (2016).

for an agent to do some act that she have a set of desires that is appropriately related to her doing that act.

Aside from its intrinsic interest, reasons internalism is important because it forms the basis of two influential objections to moral realism. The first is *the internalist argument for moral relativism*:

- (1) There is a normative reason for an agent to ϕ only if she has desires of the right sort, such as a desire that is promoted by ϕ -ing.
- (2) An agent has a moral duty to ϕ only if there is a normative reason for her to ϕ .
- (3) So, an agent has a moral duty to ϕ only if she has desires of the right sort, such as a desire that is promoted by ϕ -ing.¹⁴

Since different people have different desires, the conclusion of this argument implies that different agents have different moral duties depending on which desires they happen to have. This kind of moral relativism stands in conflict with moral realism's objectivity thesis. The first premise of the internalist argument for moral relativism is reasons internalism. The second premise is *moral rationalism*, according to which morality is robustly normative in the sense that moral requirements entail corresponding reasons for obedience.

The second important internalist challenge to moral realism is *the internalist argument for moral error theory*, which Richard Joyce has defended in its most sophisticated form. The argument goes as follows. For any person S and action ϕ :

- (1) If S morally ought to ϕ , then S morally ought to ϕ regardless of whether ϕ -ing serves any of her desires.
- (2) If S morally ought to ϕ regardless of whether ϕ -ing serves any of her desires, then there is a reason for S to ϕ regardless of whether ϕ -ing serves any of her desires.
- (3) So, if S morally ought to ϕ , then there is a reason for S to ϕ regardless of whether ϕ -ing serves any of her desires.
- (4) But it's never the case that there is a reason for a person to do some action regardless of whether doing that action serves any of her desires.

¹⁴ This relativistic argument is defended by Harman (1975, 2000).

(5) So, it's never the case that a person morally ought to do some action.¹⁵

The truth of the conclusion of this argument would refute moral realism's success thesis.¹⁶ Premise (4) is an expression of reasons internalism. Premise (2) is meant to be plausible in virtue of its following from moral rationalism. And Joyce defends (1) by observing that we ordinarily do not take back our judgment that someone ought to act in a certain way once we discover that acting in that way fails to promote her desires.¹⁷ For instance, we would not take back our judgment that the Nazis ought to have refrained from committing genocide if we learned that refraining from genocide would have only frustrated their desires. Premise (1) is therefore quite credible.

Much more could be said in elucidation of these internalist challenges to moral realism, both of which have generated a large literature. But the very brief explication that I have offered above is enough for our purposes. In this chapter, I aim to defend moral realism from both of these internalist challenges. My strategy is to provide and vindicate a new argument for *reasons externalism*, the view that there are some desire-independent reasons for action of the sort ruled out by reasons internalism. If this argument succeeds, then reasons internalism is false and moral realism is safe from the above two internalist arguments. After defending my argument for reasons externalism in section 1.2, I go on in section 1.3 to respond to four arguments for reasons internalism that have recently been offered by Kate Manne and Julia Markovits.

¹⁵ This type of argument for moral error theory is defended by Mackie (1977), Joyce (2001), and Olson (2011).

¹⁶ It is worth noting that there is a gap between the conclusion of the internalist argument for moral error theory and the falsity of the moral realist's success thesis. This is because: (i) the conclusion of this argument directly entails the falsity (or at least, the non-truth) only of those moral judgments according to which someone morally ought to do some action; and (ii) there are other types of moral judgment, such as those that ascribe moral virtues to people and those that ascribe moral goodness to states of affairs. But in my view, Joyce (2001: 175-177) persuasively fills this argumentative gap by arguing that all moral judgments entail the truth of some moral judgment of the sort that's directly targeted by the internalist argument for moral error theory.

¹⁷ *Ibid.*, 42-43.

1.2 The Argument from Epistemic Duties

Other philosophers have appealed to epistemic normativity to defend reasons externalism (henceforth, just ‘externalism’). But most of these philosophers use epistemic normativity to show only that some arguments against externalism fail. They do not draw on this kind of normativity to establish directly the more interesting conclusion that externalism is true. Consider an example. John Mackie famously argued against the existence of desire-independent reasons for action.¹⁸ On one interpretation of Mackie’s argument, he argues against such reasons on the grounds that their existence would not cohere with *metaphysical naturalism*, the view that the only sort of things that exist are those that can be investigated by science.¹⁹ Since (according to Mackie) metaphysical naturalism is true, and since the existence of desire-independent reasons for action would conflict with naturalism, Mackie concludes that there are no desire-independent reasons for action: that is, externalism is false.

Proponents of externalism have appealed to epistemic reasons for belief to rebut this argument from naturalism. Derek Parfit, for example, argues that epistemic reasons for belief are irreducibly normative. Since this is so, metaphysical naturalism implies that there are no epistemic reasons for belief. But such reasons do exist. So metaphysical naturalism is false and the argument from metaphysical naturalism against externalism fails.²⁰

As I said, the use of epistemic normativity to undermine arguments against externalism has been used by other philosophers, as well.²¹ But I believe that externalism’s proponents can get

¹⁸ Mackie (1977).

¹⁹ Timmons (1999: 12, 50) offers this interpretation.

²⁰ Parfit (2011: 110).

²¹ For two more examples of externalism’s proponents appealing to epistemic normativity to undermine arguments against externalism, see Millgram (1996) and Suikkanen (2011). Millgram (1996) draws on some considerations about epistemic reasons for belief in order to identify where Bernard Williams’ famous action-explanation argument against externalism goes wrong. Suikkanen (2011) appeals to epistemic normativity in order to identify where Harry Frankfurt (2004, 2006)’s argument for his love-based theory of reasons for action goes wrong. Since Frankfurt understands love as a type of desire, his love-based theory conflicts with externalism.

more mileage out of epistemic normativity than their work suggests. I will argue that epistemic considerations generate a powerful argument for the conclusion that there exist some external, desire-independent reasons for action. These considerations therefore demonstrate that externalism is true, rather than showing only that some arguments against externalism fail.

My argument for externalism relies on a certain view about the content of our epistemic duties. Some philosophers hold that the only epistemic duties that we have are those which I will call *epistemic duties of belief*. These are duties of a distinctively epistemic sort (distinct from moral and prudential duties, for example) to hold certain doxastic attitudes. Richard Feldman defends a version of the view that all epistemic duties are epistemic duties of belief by arguing that our only epistemic duty is to have doxastic attitudes that fit the evidence that we presently possess.²² This narrow view of the content of our epistemic duties is rejected by other epistemologists who hold that epistemic duties can favor actions in addition to doxastic attitudes. These actions include the acts of gathering evidence in certain ways rather than others, marshalling additional evidence for certain propositions, asking an expert for her opinion, maintaining a system of belief, and even breathing.²³ Call these sorts of epistemic duties *epistemic duties of action*.

My argument for externalism depends on the broader view about the content of our epistemic duties that allows for epistemic duties of action. Before I defend that view, however, let me lay out my overall argument, so that we can more clearly see the role that this expansive view of epistemic duties plays. For any person S and action ϕ :

²² Feldman (2002).

²³ Kornblith (1983) proposes that we can have epistemic duties to go about gathering evidence in certain ways rather than others. Hall and Johnson (1998) argue that we have an epistemic duty to seek more evidence regarding all propositions about which we are uncertain. Booth (2006) argues that we can have reasons ‘from the epistemic point of view’ to perform certain kinds of actions, including asking an expert for their opinion. Chrisman (2016) argues that we are sometimes under epistemic duties to engage in various activities constitutive of maintaining our system of belief. And Talbot (2014) claims that we oftentimes have an epistemic duty to breathe since breathing is sometimes necessary to acquire true beliefs.

- (1) If S epistemically ought to ϕ , then S epistemically ought to ϕ regardless of whether ϕ -ing serves any of her desires.
- (2) If S epistemically ought to ϕ regardless of whether ϕ -ing serves any of her desires, then there is a reason for S to ϕ regardless of whether ϕ -ing serves any of her desires.
- (3) So, if S epistemically ought to ϕ , then there is a reason for S to ϕ regardless of whether ϕ -ing serves any of her desires.
- (4) Sometimes, we epistemically ought to perform some action (for example, gather more evidence, go about the activity of information-gathering in certain ways, or take a critical thinking class).
- (5) So, sometimes there is a reason for us to perform some action regardless of whether performing that action serves any of our desires.

Call this the *argument from epistemic duties*. Premise (4) expresses commitment to the broader view of the content of our epistemic duties that I mentioned above.²⁴ Premise (1) is the thesis that epistemic duties are such that an agent's having an epistemic duty to do some action does not depend on her having a desire of the right sort. For example, (1) implies that if I have an epistemic duty to reflect critically on my beliefs, then my having this epistemic duty does not depend on my possession of a desire promoted by reflecting critically on my beliefs. Premise (1) therefore asserts that epistemic duties are 'categorical' in the sense elucidated by Philippa Foot when she argued that there's a sense in which both moral and etiquette requirements are categorical.²⁵

Underlying premise (2) is a thesis that we can call *epistemic rationalism*: If a person S epistemically ought to ϕ , then there is a reason for S to ϕ . Thus, premise (2) depends on the view that epistemic duties are robustly normative in the same sense in which proponents of moral rationalism hold moral duties to be robustly normative.

The conclusion of the argument from epistemic duties states that there are reasons for action that are independent of the agent's desires in the way ruled out by internalism. So,

²⁴ I am assuming that duties and oughts are the same in all relevant respects. I am not aware of any difference between them that makes a difference for my argument.

²⁵ Foot (1972).

externalism is true if (1), (2), and (4) can be established. I shall argue that we should accept all of these premises.

1.2.1 Defense of Premise (4)

Before defending the fourth premise of my argument, it may be helpful to offer some remarks about how I am conceiving of ‘the epistemic.’ I will seek to investigate what we ought to do from the epistemic point of view. This task would be easier if we had before us an analysis of the concepts *the epistemic point of view*, *the epistemic domain*, or *epistemic duty*. But I will not be offering analyses of these concepts, for I am not sure what to think about this difficult matter.²⁶ This does not mean that we cannot make progress towards understanding the content of our epistemic duties, however. Compare this to the case of the moral domain. It is quite difficult to analyze the concept of *morality*. ‘What makes a duty a *moral* duty?’ is a difficult question that moral philosophers have not yet resolved.²⁷ But this has not prevented ethicists from making progress in understanding the content of our moral duties. Such progress is possible because we have moral intuitions, or intuitions that certain behaviors are favored *from the moral point of view*. For instance, it is intuitive that giving to famine relief is something that, from the moral point of view, we sometimes ought to do. Analogously, I assume that there is another point of view—one that it feels natural to call ‘the epistemic point of view’ or ‘the epistemological point of view’—that is such that we sometimes intuit that some bit of behavior is favored from this point of view. A paradigm example of behavior that is oftentimes favored from this point of view is believing that *p* when one has strong evidence that *p*.

²⁶ Others have tried to analyze *the epistemic*. Two proposals that I find plausible and that are congenial for my arguments are those of Cuneo (2007: 56-59, 75) and Hazlett (2013: 267-69).

²⁷ See Gert and Gert (2016) for a helpful overview of attempts to fix the domain of moral propositions and judgments.

I find premise (4) to be quite plausible partly on intuitive grounds. Consider the following actions:

- seeking out evidence about some matter before forming a judgment about it
- enrolling in a critical thinking class
- consulting an expert's opinion
- seeking out conversations with people who hold different moral and political views

Surely we sometimes ought to behave in these ways. Sometimes, we *morally* ought to do these things. For example, perhaps I have promised you that I would consult that expert's opinion. And sometimes, we *prudentially* ought to perform these actions. For example, perhaps it would be in my self-interest to graduate and I can graduate only if I take that critical thinking course. But each of these actions are also such that we sometimes ought to do them from the epistemic point of view. Intuitively, we fail to act as we *epistemically* ought to act when we (say) strive to surround ourselves only with people who hold the same views that we hold.

To further support this appeal to intuition, consider a case from Terence Cuneo:

Sam receives some critical feedback on a project in progress from his colleague Margaret. Margaret is astute and perceptive; however, she does tend to voice her views with verve, and the force of her comments is in keeping with this trait. Upon reading Margaret's comments, Sam finds himself believing that she doesn't really appreciate his work. He thereupon finds himself annoyed and ultimately angry with her...He forms the intention to ignore the generous criticism that Margaret has offered him and subsequently acts upon it.²⁸

Cuneo suggests that it is natural to think that 'from an epistemic point of view' Sam ought not to ignore Margaret's criticism, and I agree.²⁹ Premise (4) is intuitively true. This is evidence that it is true.

²⁸ Cuneo (2007: 73).

²⁹ *Ibid.*

Some people may lack these intuitions favoring (4), however. For such people, I offer the following bolstering argument. Sometimes, performing some action (such as enrolling in a critical thinking class) is a necessary means to believing in accordance with one's evidence. One always ought to believe in accordance with one's evidence, from the epistemic point of view. And the following *normative transmission principle* is true: If person S ought to X from normative point of view N, and Y-ing is a necessary means for S to X, then according to N, S ought to Y. And so, sometimes one ought to do some action, from the epistemic point of view. That is, premise (4) is true. Call this the *argument from transmission*.

I believe that the only promising strategy for rebutting this argument is rejecting the normative transmission principle.³⁰ But this principle is quite attractive for two reasons. First, it is intuitive. It would be mystifying were a normative system to require that one do something, yet be indifferent about whether one took the necessary means to doing that very thing. Second, the normative transmission principle seems to render correct judgments when it comes to the moral and prudential domains. Whenever a person morally ought to do X (for instance, meet someone where one promised to meet her) and is such that she must do some act Y (for instance, wake up early) in order to do X, it is surely true that she morally ought to do Y. And whenever a person prudentially ought to do X (for instance, pursue a fulfilling career) and is such that she must do some act Y (for instance, go to college) in order to do X, it is surely true that she prudentially ought to do Y. The truth of the normative transmission principle would therefore account for certain facts about the moral ought and the prudential ought, and this counts as a reason to accept it.

³⁰ One might object to the premise that one always epistemically ought to believe in accordance with one's evidence. Such an evidentialist view is admittedly controversial, but rejecting this view is not a promising strategy for rejecting the core idea of the argument from transmission. So long as there is *something* one epistemically ought to believe, the argument can get off the ground. For no matter what one thinks one epistemically ought to believe, we can always think of a situation in which one can believe that thing only by way of doing some action.

John Broome objects to the normative transmission principle with the following purported counter-example.³¹ Suppose that you prudentially ought to see your doctor, and that staying home from work is a necessary means for you to do this. Further, suppose that as a matter of fact, it is false that you would see your doctor were you to stay home from work: Were you to stay home from work, you would ‘simply sit around feeling anxious.’³² In this case, it seems plausible that the following are true: (i) you ought to go to the doctor from the prudential point of view; (ii) staying home from work is a necessary means for you to do this; yet (iii) it is not the case that you ought to stay home from work from the prudential point of view. So, the argument goes, the normative transmission principle is false.

I am not sure whether I share Broome’s intuitions. But if the reader is convinced by this type of purported counter-example, then she should just replace the normative transmission principle with the *modified normative transmission principle*: If person S ought to X from normative point of view N, and Y-ing is a necessary means for S to X, and it’s true that were S to do Y then she would do X, then according to N, S ought to Y. This modified normative transmission principle is sufficient for my purposes. For sometimes, all of the following are true: (i) one epistemically ought to believe in accordance with one’s evidence, (ii) one can believe in accordance with one’s evidence only by way of performing some action A, and (iii) were one to perform action A, one would end up believing in accordance with one’s evidence. I conclude that we sometimes epistemically ought to do some action.

The argument from transmission, along with (4)’s intuitive appeal, provides a strong presumptive case for (4). We should therefore accept (4) in the absence of compelling reasons to reject it. I do not believe that there are any compelling reasons to reject (4), and to bolster this

³¹ Broome (2013: 126).

³² *Ibid.*, 126.

contention, I will rebut the best argument against (4) of which I am aware. I have in mind the following argument from Feldman:

(F1) ‘Epistemological duties are duties that one must carry out in order to be successful from an intellectual (or epistemological) perspective.’³³

(F2) ‘Epistemological success consists in having reasonable or justified cognitive attitudes.’³⁴

(F3) So, epistemological duties are duties that one must carry out in order to have reasonable or justified cognitive attitudes.

(F4) No actions are such that one must perform them in order to have reasonable or justified cognitive attitudes.³⁵

(F5) So, epistemological duties are never duties to perform some action.³⁶

I grant (F1) and (F2) for the sake of the argument. We should reject (F4). Sometimes, we must perform certain actions in order to have justified doxastic attitudes. I will assume, along with Feldman, that justified doxastic attitudes are those that fit one’s evidence. So, my claim is that (F4) is false because we sometimes must perform certain actions in order to achieve having doxastic attitudes that fit our evidence. One situation of this sort is when someone is quite bad at discerning which propositions are supported by her evidence, and is such that she can get better at this only by enrolling in a critical thinking class. In such a scenario, there is an action—enrolling in a class—that one must perform in order for one to have justified doxastic attitudes. So, (F4) is false.

Here is another problem case for (F4). We cannot think carefully when we are terribly sleep-deprived or malnourished. Thinking carefully is often a prerequisite for the acquisition of justified doxastic attitudes. So, we often are in the situation of needing to get rest and to maintain an adequate diet in order to have justified doxastic attitudes. Getting rest and maintaining a certain diet are actions, and so, (F4) is false.

³³ Feldman (2002: 376).

³⁴ *Ibid.*, 379.

³⁵ *Ibid.*, 380.

³⁶ I will assume that by ‘epistemological duty,’ Feldman means what I mean when I say ‘epistemic duty.’

Feldman seems to anticipate this sort of objection to (F4). He writes:

There may be some things that will help one do a better job of having the attitudes that are justified. For example, enrolling in critical thinking courses may help. But this is a sort of *indirect* way of getting the proper attitudes. Perhaps one could plausibly describe that sort of thing as what's desirable without being required.³⁷

Feldman here concedes that enrolling in a class may help one to acquire justified doxastic attitudes, and I have observed that it is further true that such action is sometimes *required* to have justified attitudes. But Feldman denies that this is a problem for his argument because such action is only an 'indirect way' of getting justified attitudes. This makes it seem as if Feldman's final argument relies on the following premise, rather than on (F4):

(F4)* No actions are such that (i) one must do them in order to have justified cognitive attitudes, and (ii) in doing them, one thereby *directly* acquires justified cognitive attitudes.

Feldman can argue that (F4)* is not threatened by the fact that enrolling in a class, getting rest, and maintaining an adequate diet are sometimes required for having justified cognitive attitudes. Such actions are not counter-examples to (F4)* because such actions are not *direct* ways of acquiring justified attitudes.

But if Feldman does ultimately endorse (F4)* as his fourth premise, he must actually be offering the following argument:

(F1)* Epistemological duties are duties that (i) one must carry out in order to be successful from an epistemological perspective, and (ii) are such that in carrying them out, one thereby *directly* realizes epistemological success.

(F2) Epistemological success consists in having justified cognitive attitudes.

(F3)* So, epistemological duties are duties that (i) one must carry out in order to have justified cognitive attitudes, and (ii) are such that in carrying them out, one thereby *directly* acquires justified cognitive attitudes.

³⁷ *Ibid.*, 380, my emphasis.

(F4)* No actions are such that (i) one must do them in order to have justified cognitive attitudes, and (ii) in doing them, one thereby *directly* acquires justified cognitive attitudes.
 (F5) So, epistemological duties are never duties to perform some action.

This alternative reading of Feldman's argument fails for a different reason: (F1)* is dubious. First, the only reason on offer to accept that epistemic duties must fulfill condition (ii) in (F1)* seems to be that this allows Feldman to avoid the objection I have offered against (F4). But this makes the inclusion of (ii) in (F1)* objectionably ad hoc. Second, (F1)* seems suspect when we notice that analogous claims in other normative domains are plainly false. For example, consider the moral analogue of (F1)*:

Moral duties are duties that (i) one must carry out in order to be successful from a moral perspective, and (ii) are such that in carrying them out, one thereby *directly* realizes moral success.

We sometimes have a moral duty to do acts that only *indirectly* realize moral success. One may have a moral duty to put a check in the mailbox for some charitable organization, despite the fact that this action only *indirectly* realizes the moral success of relieving others' suffering. So, the moral analogue of (F1)* is false.

Similarly, the prudential analogue of (F1)* seems false. Here is the prudential analogue:

Prudential duties are duties that (i) one must carry out in order to be successful from a moral perspective, and (ii) are such that in carrying them out, one thereby *directly* realizes prudential success.

We sometimes prudentially ought to do things that only *indirectly* realize prudential success. One may have a prudential obligation to go to college, despite the fact that this only indirectly realizes the prudential success of acquiring a fulfilling career. So, the prudential analogue of (F1)* is false.

I submit that (F1)* is suspect because its analogues in other normative domains are false. It may be that the epistemic domain is unlike others in such a way that its characteristic duties are tied to the *direct* realization of its characteristic success. But this would be reasonable to believe only if we can find some relevant difference between the epistemic domain and these other domains that would explain why this might be so, and Feldman does not identify any such relevant difference.

To recap: Feldman's argument either relies on (F4) or on (F1)*, and both of these claims are dubious. Feldman's argument should therefore be rejected. I conclude that we have most reason to accept premise (4) of the argument from epistemic duties.³⁸

1.2.2 Defense of Premise (1)

Recall that premise (1) states that if S epistemically ought to do some act, then she epistemically ought to do that act regardless of which desires she has. To illustrate, (1) entails that if you epistemically ought to pay attention to the expert's constructive criticism, then you epistemically ought to do this even if you don't have any desires that are promoted by paying attention to her criticism. In other words, (1) expresses the view that epistemic duties of action are *categorically applicable*.

To see why we should accept premise (1), it is helpful to begin by reflecting on the sorts of cases that have led many philosophers to believe that epistemic duties *of belief* are categorically applicable. Consider the case of Ella. Ella desires only one thing, namely psychological contentment. She has acquired decisive evidence that her father has cancer. So intuitively, she

³⁸ It is worth noting that premise (4) is immune to the most well-known objection to epistemic duties, namely Alston (1988)'s argument from doxastic involuntarism. The targets of Alston's argument are epistemic duties of belief. Assuming that a very strong sort of determinism is false, we do sometimes have control over which actions we perform; so nothing like Alston's argument against epistemic duties of belief will work as an argument against epistemic duties of action.

ought to believe that her father has cancer, from the epistemic point of view. But believing that her father has cancer will in no way serve her desire for psychological contentment. So, Ella's epistemic duty of belief is one that she has independently of her desires: it applies to her categorically.³⁹

Or consider a well-known case that Thomas Kelly offers in the context of his influential critique of instrumentalist theories of epistemic rationality:

If, despite my utter lack of interest in the question of whether Bertrand Russell was left-handed, I stumble upon strong evidence that he was, then I have strong epistemic reasons to believe that Bertrand Russell was left-handed. Indeed, my epistemic reasons will be no different than they would be if I had acquired the same evidence deliberately, because I *did* have the goal of finding out whether Russell was left-handed. Once I come into possession of evidence which strongly supports that claim that p, then I have epistemic reasons to believe that p, regardless of whether I presently have or previously had the goal of believing the truth about p, or any wider goal which would be better achieved in virtue of my believing the truth about p.⁴⁰

Kelly puts his point in terms of epistemic reasons, but we can just as easily draw a conclusion about epistemic duties. You ought to believe that Bertrand Russell was left-handed when you have completely decisive evidence that he was, even if you have no goals or desires promoted by believing that Bertrand Russell was left-handed. Reflection on cases like these make it reasonable to believe that epistemic duties of belief are categorically applicable.

These cases can be modified in such a way that they provide us with good reason to believe that epistemic duties of action are categorically applicable. First, consider Ella's case again, but let us now add two further stipulations. First, Ella's psychological fragility makes it the case that she cannot believe that her father has cancer unless she seeks a special kind of cognitive psychotherapy. Undergoing such therapy is a necessary means for her to believe that her father has

³⁹ Cuneo (2007: 207-208) introduces the case of Ella.

⁴⁰ Kelly (2003: 625).

cancer. Second, Ella has no desires that would be promoted by seeking this special kind of cognitive psychotherapy. Just as in the original case, Ella's having the decisive evidence that she does implies that she ought to believe that her father has cancer, from the epistemic point of view. Because of the normative transmission principle that I defended above, it follows that Ella ought to seek the cognitive psychotherapy, from the epistemic point of view. And because Ella has no desires that would be promoted by acting in this way, this epistemic duty of action must apply to Ella categorically. That is, her possession of an epistemic duty to seek cognitive psychotherapy does not depend on her having some appropriate set of desires. This line of reasoning shows that at least some epistemic duties of action are categorically applicable.

Next, consider the Bertrand Russell case again, but let us add two further stipulations. First, let us suppose that the person who stumbles across decisive evidence that Russell was left-handed has a strange condition that makes her unable to believe that Russell was left-handed unless she swallows a certain pill. And second, let us suppose that this person has no desires that would be promoted by swallowing this belief-enabling pill. Just as in the original case, this person's having decisive evidence that Bertrand Russell was left-handed seems to make it the case that she epistemically ought to believe that this is so. Because of the normative transmission principle, it follows that this person epistemically ought to swallow the pill. And because this person has no desires that would be promoted by swallowing the pill, this epistemic duty of action must apply to her categorically. That is, her possession of an epistemic duty to swallow the pill does not depend on her having some appropriate set of desires. Again, we see that some epistemic duties of action are categorically applicable.

Of course, premise (1) makes the stronger claim that *all* epistemic duties of action are categorically applicable. And it is of course not logically incoherent to hold that while some

epistemic duties of action are categorically applicable, other epistemic duties of action are not. But I submit that the foregoing considerations generate considerable pressure to accept premise (1).

Consider a general thesis about the epistemic domain:

Epistemic Categoricity: All epistemic duties are categorically applicable.

This thesis is theoretically attractive in light of the following facts, which we can appreciate in light of the above discussion: (i) epistemic duties of belief are categorically applicable, as evidenced by the original Ella and Bertrand Russell cases; (ii) the epistemic duties of action that are depicted in the modified Ella and Bertrand Russell cases are categorically applicable; and (iii) using reasoning similar to the sort employed above, it would be easy to demonstrate that many other epistemic duties of action are also categorically applicable. Reflection on these facts directs our attention to a number of cases in which Epistemic Categoricity holds true. The truth of Epistemic Categoricity would therefore unify the epistemic domain, as well account for the categoricity of many particular epistemic duties. On these grounds, I suggest that we have reason to accept Epistemic Categoricity, and consequently, that we have reason to accept premise (1) of the argument from epistemic duties.⁴¹

1.2.3 Defense of Premise (2)

Recall premise (2) of the argument from epistemic duties:

(2) If S epistemically ought to ϕ regardless of whether ϕ -ing serves any of her desires, then there is a reason for S to ϕ regardless of whether ϕ -ing serves any of her desires.

⁴¹ In case the argument in this paragraph seems too quick for some readers, I will note that the core idea of my argument from epistemic duties survives even if not all epistemic duties of action are categorically applicable. To see that this is so, suppose that just a single epistemic duty of action is categorically applicable: namely, Ella's epistemic duty to seek cognitive psychotherapy. This fact, combined with premise (2) of the argument from epistemic duties, entails that Ella has a reason to seek cognitive psychotherapy regardless of her desires. This conclusion establishes reasons externalism.

And recall that underlying this premise is epistemic rationalism, the thesis that if a person S epistemically ought to do act ϕ , then there is a reason for S to ϕ . Epistemic rationalism entails (2), so I will defend (2) by defending epistemic rationalism.

My argument for epistemic rationalism appeals to the connections between epistemic duties, negative attitudes such as blame and criticism, and reasons. It goes as follows (let ' ϕ ' stand for some action):

(E1) If a person S epistemically ought to ϕ , then S would be criticizable if she were to refrain from ϕ -ing.

(E2) If S would be criticizable if she were to refrain from ϕ -ing, then there is a reason for S to ϕ .

(E3) Therefore, if a person S epistemically ought to ϕ , then there is a reason for S to ϕ .

Call this the *epistemic argument from criticizability*. When I say that a person is *criticizable if she refrains from ϕ -ing*, I mean that it is legitimate or justified for others to be critical of her if she fails to ϕ . And to be *critical* of someone is to have some attitude towards that person that casts her in a negative light, such as disapproval, disrespect, reproach, and blame. We should distinguish between being critical of someone and expressing that criticism. Being critical of someone need not involve any such overt expression of negative attitudes. We should also distinguish between a person's being criticizable in the sense I have elucidated and its being accurate to think that a person has violated a norm. A person who violates a rule of etiquette is such that it is accurate to say that she has violated a norm, but this doesn't necessarily mean that others are justified in having some critical attitude such as reproach towards her.

Premise (E1) asserts that it is appropriate to be critical of a person when she violates her epistemic duties of action. One reason to accept this premise is that the analogous principle about epistemic duties *of belief* is true. When someone believes a proposition that she epistemically ought not to believe, it seems that she is justly criticized, blamed, or reproached. To illustrate, consider

a friend of mine who once told me that he holds his religious beliefs because they make him feel happy. Believing something just because doing so makes one feel happy is epistemically wrong: we ought not to do this, from the epistemic point of view. Just as the analogue of (E1) for duties of belief predicts, my friend seems criticizable for believing as he does. We are inclined to be critical of my friend, and having some critical attitude seems appropriate.⁴² So, we have reason to accept (E1), unless some relevant difference between epistemic duties of action and epistemic duties of belief is identified that makes it plausible that only the latter sort of duty is such that violations of them would justify criticism.⁴³

A second reason to believe (E1) is that it is directly intuitive that violations of epistemic duties of actions make one worthy of criticism. Recall the case of Sam, and suppose that he chooses to ignore his friend's insightful criticism because it makes him uncomfortable to listen to criticism of his ideas. Sam fails to act in the way that he ought to, from the epistemic point of view. And sure enough, we are critical of Sam: we disapprove of him for ignoring the criticism. And this critical attitude seems appropriate, called for, legitimate. (E1) correctly predicts that Sam's behavior merits criticism, and this is further reason to accept (E1).

Premise (E2) is quite plausible because it accounts for two facts. First, (E2) accounts for the fact that a common strategy for defending oneself from criticism for failing to ϕ is to argue that there was no reason for one to ϕ . For example, suppose that I pass out muffins to everybody except Sam, and that you reproach me for failing to offer Sam a muffin, as well. It would be natural

⁴² Of course, it will sometimes be quite inappropriate for us to express this critical attitude, say by making fun of him. But that is a different matter.

⁴³ Many epistemologists have observed that violations of epistemic duties of belief seem to render a person subject to legitimate criticism. For example, Nicholas Wolterstorff writes: 'We say to each other such things as, "You should have known better than to think that Borges was an English writer," "You should be more trusting of what our State Department says," and "You should never have believed him when he told you that the auditors had approved that way of keeping books." Not only do we *regret* the knowledge and ignorance of our fellow human beings, their beliefs, disbeliefs, and non-beliefs; we *reproach* them, *blame* them, *chastise* them" (Wolterstorff, 2005: 326). See also Grimm (2009: 253-56), Hazlett (2013: 134), and Heil (1983: 362).

and sensible for me to defend myself by pointing out that Sam despises muffins. In pointing this out, I am demonstrating that the consideration that you took to be a reason to offer Sam a muffin—namely, that Sam would enjoy a muffin, too—is no reason at all to offer Sam a muffin. Given (E2), this natural strategy for defending myself makes sense. For if (E2) is true, then this is a way to show that your critical attitude towards me for failing to give Sam a muffin is not justified.

Second, (E2) accounts for the fact that the most obvious cases in which a person does not merit criticism for failing to perform an action are those in which a person refrains from performing an action that there clearly was no reason for her to do. For example, one is normally not criticizable for putting one's left sock on before one's right sock, thereby failing to put one's right sock on before one's left sock. (E2) implies that this is so since there is normally no reason to put one's right sock on before one's left sock. Since (E2) accounts for the two facts that I have just described, we have reason to accept (E2).⁴⁴

I have argued that we have good reason to accept (E1) and (E2), and thus that we have good reason to accept epistemic rationalism. But now I want to argue that the critics of moral realism to whom I am primarily responding in this paper face special pressure to accept (E1) and (E2). Recall that my primary goal is to defend moral realism from the internalist argument for moral relativism and the internalist argument for moral error theory. And recall from section 1.1 that moral rationalism is a crucial premise of both of these anti-realist arguments. My anti-realist interlocutors are therefore committed to there being a good reason to accept moral rationalism. Their challenges to moral realism fail if there is no such reason. This is why it is of great importance that the epistemic argument from criticizability is just the epistemic analogue of the most popular and formidable argument for moral rationalism. This argument for moral rationalism,

⁴⁴ See Portmore (2011) and Vogelstein (2013: 1091-1094) for further defense of this connection between criticizability and reasons.

which we can call the *moral argument from criticizability*, appeals to the connections between reasons, negative attitudes like criticism or blame, and moral duties. Here is a version of this kind of argument (let ‘ φ ’ stand for some action):

(M1) If a person S morally ought to φ , then S would be criticizable if she were to refrain from φ -ing.

(M2) If S would be criticizable if she were to refrain from φ -ing, then there is a reason for S to φ .

(M3) Therefore, if a person S morally ought to φ , then there is a reason for S to φ .⁴⁵

Premise (M2) is identical to (E2). Thus, my anti-realist interlocutors cannot reject (E2) without thereby rejecting a key premise in the most important argument for moral rationalism. They therefore face pressure to accept (E2).

It also seems to me that my anti-realist interlocutors face pressure to accept (E1). Suppose that they reject (E1) by denying that we are subject to legitimate criticism whenever we violate our epistemic duties. They would then be committed to thinking that our intuitions about when a person is criticizable suffer from systematic error. For as I argued above, it is intuitive that a person who violates her epistemic duties is thereby criticizable. If we learn that our intuitions about when a person is criticizable suffer from this sort of systematic error, then we have acquired reasons for doubting the trustworthiness of these intuitions. And this undermines our grounds for accepting (M1); for the only grounds for accepting (M1) is that it is intuitive that those who violate their moral duties are thereby subject to warranted criticism. So, my anti-realist interlocutors cannot reject (E1) without thereby being committed to doubting a key premise in the most important argument for moral rationalism. They therefore face pressure to accept (E1). I conclude that the

⁴⁵ Philosophers who offer arguments for moral rationalism along these lines include Darwall (2003: 15), Joyce (2001), Portmore (2011), Shafer-Landau (2003: 192-193), Skorupski (1999: 42-43), and Vogelstein (2013).

anti-realist critics to whom I am primarily responding face pressure to accept both of the premises of my argument for epistemic rationalism.

I concede that the anti-realist can avoid this special pressure to accept (E1) and (E2) by defending moral rationalism without appealing to the moral argument from criticizability. But the significance of this concession is minimized by the following two considerations. First, the fact of the matter is that there are few arguments for moral rationalism in the literature. The anti-realist therefore deprives herself of a crucial source of support for moral rationalism if she decides to avoid reliance on the moral argument from criticizability. And second, it is quite difficult to think of an alternative argument for moral rationalism that does not have an equally-plausible epistemic analogue. It is therefore difficult for the anti-realist to avoid being committed to the success of some argument for epistemic rationalism even if she seeks to defend moral rationalism without appealing to the moral argument from criticizability.⁴⁶

Opponents of epistemic rationalism may try to undermine my argument for it by providing putative counter-examples to (E1). Suppose that Stephanie morally ought to ignore her friend's constructive criticism. We can imagine that by ignoring this criticism, Stephanie could thereby save one hundred people from certain death. On this basis, she ignores her friend's criticism. It seems that we should not be critical of Stephanie for failing to do what she epistemically ought to do. That is, it seems inappropriate to have negative attitudes such as disapproval or reproach towards Stephanie just because she failed to pay attention to her friend's constructive criticism.

⁴⁶ For example, Brink (1992) briefly suggests the following line of reasoning for moral rationalism: "Agents typically engage in moral deliberation in order to decide what to do; people give moral advice with the aim of guiding others' conduct; and most of us are quite sensitive to moral criticism. One explanation of these attitudes and expectations is that we think moral requirements give agents reasons for action" (1-2). If this argument for moral rationalism succeeds, then an analogous argument for epistemic rationalism probably succeeds. This is because (i) agents typically engage in deliberation about what they epistemically ought to do in order to decide what to do; (ii) people give epistemic advice with the aim of guiding others' conduct; and (iii) most of us are quite sensitive to epistemic criticism.

This case is therefore a counter-example to (E1). More generally, (E1) seems false in light of cases in which someone epistemically ought to ϕ , yet is morally required to refrain from ϕ -ing.⁴⁷

I am unsure whether I share the intuitions that are driving this objection. Perhaps in situations in which one epistemically ought to do one action and morally ought to do some incompatible action, one is subject to legitimate criticism no matter what one does. If so, then this type of case is not a counter-example to (E1). But if the reader shares the intuitions that are driving this objection, then I can offer her a slightly-modified version of my argument that avoids it entirely. Consider the following version, letting ' ϕ ' be some action:

(E1)* If a person S epistemically ought to ϕ , then S would be *pro tanto* criticizable if she were to refrain from ϕ -ing.

(E2)* If S would be *pro tanto* criticizable if she were to refrain from ϕ -ing, then there is a reason for S to ϕ .

(E3) So, if a person S epistemically ought to ϕ , then there is a reason for S to ϕ .

Call this the *modified epistemic argument from criticizability*. The only difference between this argument for epistemic rationalism and my original one is that this argument refers to *pro tanto* criticizability instead of referring to criticizability. When I say that a person is *pro tanto criticizable if she refrains from ϕ -ing*, I mean that there is some reason for others to be critical of her if she fails to ϕ . A person can be *pro tanto* criticizable for refraining from ϕ -ing without being criticizable for so refraining. To illustrate, suppose that I promised to meet you at the park, but keeping this promise would require me to refrain from picking up my child from school. Arguably, I am not criticizable for refraining from keeping my promise to you in order to pick up my child from school: it seems inappropriate for others to be critical of me for prioritizing my child's wellbeing in this way. Yet it does seem that I am *pro tanto* criticizable if I refrain from keeping my promise

⁴⁷ I thank Rebecca Chan and Ron Aboudi for pressing this objection. A parallel objection could be constructed by directing attention to a case in which one's epistemic duties conflicts with one's *prudential* duties.

to you in order to pick up my child from school. The fact that I would be breaking a promise is surely a (outweighed) reason for others to be critical of me.

The support for (E1) that I offered above can be slightly adjusted so that it serves as support for (E1)*. One reason to accept (E1)* is that the analogous principle about epistemic duties of *belief* is true. That is, a person is *pro tanto* criticizable for failing to believe *p* if she epistemically ought to believe *p*. To illustrate, suppose that I have decisive evidence for the claim that it is raining (I see that it is raining outside), and consequently, that I epistemically ought to believe that it is raining. And suppose that I fail to believe that it is raining, and in fact, that I manage to delude myself into believing that it is not raining. The analogue of (E1)* for epistemic duties of belief implies that I am *pro tanto* criticizable for violating my epistemic duty in this case. And this seems exactly right. So, we have reason to accept (E1)*, unless some relevant difference between epistemic duties of action and epistemic duties of belief is identified that makes it plausible that only the latter sort of duty is such that a person's violations of them would contribute to others' justification for being critical of her.

A second reason to believe (E1)* is that it is directly intuitive that there is at least something to be said in favor of being critical of someone when she violates an epistemic duty of action. Recall once against the case of Sam, and suppose that he ignores his friend's insightful criticism because listening to constructive criticism of his ideas makes him uncomfortable. Sam fails to act in the way that he ought to, from the epistemic point of view. We are inclined to disapprove of Sam for ignoring the criticism, and intuitively, there is something to be said in favor of our having this disapproving attitude. (E1)* correctly predicts that Sam's behavior is *pro tanto* criticizable, and this is reason to accept (E1)*.

Further, (E1)* is not vulnerable to the purported counter-example to (E1) that was described above. For (E1)* does not have the implausible implication that Stephanie is criticizable if she ignores her friend's constructive criticism in order to save 100 people. Of course, (E1)* does imply that Stephanie is *pro tanto* criticizable if she ignores her friend's criticism in order to save 100 people. But upon reflection, this seems exactly right. Others have some reason to be critical of Stephanie for violating her epistemic duty, but this reason is outweighed by much stronger reasons *not* to be critical of her for violating it—reasons deriving from the fact that violating her epistemic duty was necessary to save 100 people. So, we have good reasons to accept (E1)* and no clear reasons to reject it.

The support for (E2) that I gave above can be slightly adjusted so that it serves as support for (E2)*. (E2)* is quite plausible because it accounts for two facts. First, as I argued above, a common strategy for defending oneself from criticism for failing to ϕ is to argue that there was no reason for one to ϕ . Given (E2)*, this natural strategy for defending oneself makes sense. For if (E2)* is true, then this is a way to show that others' critical attitudes towards one is not favored by any reason. Second, (E2)* accounts for the fact that the most obvious cases in which a person is not *pro tanto* criticizable for failing to perform an action are those in which a person refrains from performing an action that there is no reason for her to do. For example, one is normally not *pro tanto* criticizable for putting one's left sock on before one's right sock, thereby failing to put one's right sock on before one's left sock. (E2)* implies that this is so since there is normally no reason to put one's right sock on before one's left sock. Since (E2)* accounts for the two facts that I have just described, we have reason to accept (E2)*.

I conclude that (E1)* and (E2)* constitute good grounds for accepting epistemic rationalism. So, we have good grounds for accepting premise (2) of the argument from epistemic duties even if the purported counter-example to (E1) described above succeeds.

Because we have good grounds for accepting all of the premises of the argument from epistemic duties, this argument gives us good reason to believe that there are at least some desire-independent reasons for action. We therefore have good grounds for accepting reasons externalism, and this secures moral realism from the internalist arguments for moral relativism and moral error theory.

1.3 Critique of Arguments for Reasons Internalism

In the last few years, internalists have developed new arguments for internalism and against externalism, and it may be objected that these arguments provide us with decisive reason to reject externalism, the argument from epistemic duties notwithstanding. In the remainder of this chapter, I aim to close off this line of response by arguing that the recent arguments for internalism by Kate Manne and Julia Markovits are unsuccessful.

1.3.1 Manne's Reasons Internalism

Kate Manne develops an argument for the following version of internalism:

Manne's Internalism (MI): There is a reason for person S to ϕ only if S would end up in a state of being (somewhat) motivated to ϕ , following an idealized process of being reasoned with about the matter of what to do.⁴⁸

A difference between MI and the type of internalism that I have focused on above is that MI is put in terms of *motivation* whereas the type that I have focused on is put in terms of *desires*. But on the assumption that a person is motivated to perform an action only by having a desire that is

⁴⁸ Manne (2014: 109).

served by performing that action, MI entails that a necessary condition on there being a reason for a person to perform an action is that she have the right sort of desires: namely, desires the possession of which ensures that she will end up motivated to perform that action after ‘an idealized process of being reasoned with about the matter of what to do.’ So, MI is the sort of internalist thesis that can play the required role in the internalist argument for moral relativism and the internalist argument for moral error theory. It is therefore important for my purposes to rebut Manne’s argument in favor of MI.

Manne understands the ‘idealized process’ mentioned in this principle to be a completed conversational process in which S’s ‘ideal advisor’ reasons with S about what S ought to do.⁴⁹ We are supposed to imagine S’s ideal advisor as being ‘a flesh and blood human being’ who is ‘possessed of all the relevant information’ and who is ‘fully procedurally rational (or at least as fully procedurally rational as any actual human being could be).’⁵⁰ Further, S’s ideal advisor is virtuous, wise, well-disposed toward her advisee, and is ideally suited for ‘getting through’ to S, morally.⁵¹ According to Manne, then, an idealized process of an agent’s being reasoned with about the matter of what to do is a completed process in which an advisor with all of these features reasons with the agent about what she ought to do.

1.3.2 Why Manne’s Argument for MI Probably Goes Wrong Somewhere

Before engaging with Manne’s argument for MI directly, I will show that there are intuitive counter-examples to MI, thereby showing that her argument for MI probably goes wrong somewhere. A counterexample to MI is a case in which there is a reason for an agent S to ϕ , yet it is false that S would end up in a state of being (somewhat) motivated to ϕ , following an idealized

⁴⁹ *Ibid.*, 97.

⁵⁰ *Ibid.*

⁵¹ *Ibid.*

process of being reasoned with about the matter of what to do. Consider two cases that are intuitively of this sort.

First, there are cases of extreme introverts. Suppose that Ms. Introvert is just like any ordinary human being in this respect: if she goes too long without having a conversation with another person, she starts to feel terribly lonely and sad. But further, suppose that Ms. Introvert is unusually introverted, in this way: all social interaction thoroughly exhausts her, such that it takes weeks for her to psychologically recuperate after conversing with someone. In fact, when she has more than one conversation in a month, she becomes completely miserable.

Now, suppose that Ms. Introvert has not interacted with anybody for a very long time, and she starts to feel lonely. Fortunately, her best friend from high school is in town, and she knows that a visit with her old friend would buoy her spirits. Intuitively, Ms. Introvert has excellent reason to call her friend and arrange to meet up with her. Yet it is surely false that Ms. Introvert would end up in a state of being motivated to call her friend, following an idealized process of being reasoned with over the matter of what to do. For following such a process, Ms. Introvert would have had her fill of socializing, and calling her old friend would be the last thing she would want to do. So, Ms. Introvert's case seems to be a counterexample to MI.

Second, there are cases of extremely proud agents. Suppose that Ms. Prideful prides herself on her rationality, and she cannot stand the thought that there might be people who are more rational than she is. In fact, she has the following extreme disposition: whenever it becomes clear to her that she is interacting with someone who is far more procedurally rational than she is, she becomes overwhelmingly depressed, losing all of her ordinary cares and concerns. Fortunately, there are no actual human beings who are far more rational than Ms. Prideful, so she is able to lead

a completely ordinary life. Ms. Prideful is bizarre, but it certainly seems possible for there to be an agent with her psychological profile.

Now, suppose that it is an ordinary day, and Ms. Prideful starts to feel famished. Intuitively, she has a reason to eat a snack. Yet it is surely false that Ms. Prideful would end up in a state of being motivated to eat a snack, following an idealized process of being reasoned with over the matter of what to do. For following such a process, Ms. Prideful would realize that she is interacting with someone who is far more procedurally rational than she is, and she would therefore have none of her ordinary motivations, such as the motivation to eat when she is hungry. So, Ms. Prideful's case is another intuitive counterexample to MI.

Another way to put the objections to MI that I have offered is that it commits the conditional fallacy, a well-known problem for analyses of reasons according to which the possession of a reason depends on the truth of some counter-factual. As Robert Johnson helpfully explains:

The fallacy consists of overlooking, in various ways, dependencies between the *analysandum* and the antecedent and consequent of the conditional in the *analysans*. For instance, one might overlook a set of counter-examples for one's analysis in which the *analysandum* is obviously true, yet because the antecedent of the conditional in the *analysans* is in some way incompatible with the consequent in the *analysans*, the *analysans* is false.⁵²

The above two counterexamples show that MI commits the conditional fallacy in this way. Consider again the case of Ms. Introvert. It is obviously true that Ms. Introvert has a reason to call her friend; yet because Ms. Introvert's conversing with an ideal advisor is incompatible with her subsequently having a motivation to call her friend, the following counter-factual is false: if Ms. Introvert were reasoned with by an ideal advisor about what to do, then she would end up motivated

⁵² Johnson (1999: 54).

to call her friend. Similarly, it is obviously true that Ms. Prideful has a reason to eat a snack; yet because her conversing with an ideal advisor is incompatible with her subsequently having a motivation to eat a snack, the following counter-factual is false: if Ms. Prideful were reasoned with by an ideal advisor about what to do, then she would end up motivated to eat a snack. By overlooking the ways in which MI's consequent's antecedent can conflict with its consequent, Manne overlooks these counterexamples.

And now that we see that I have basically been relying on the conditional fallacy to advance my criticisms, it may be thought that Manne has replies. In a lengthy footnote, Manne offers some thoughts on her general strategy for addressing conditional fallacy worries.⁵³ She suggests that we may 'prevent some of the well-known conditional fallacy worries which afflict 'ideal agent' models from afflicting [her] account too' by being 'very careful about how we individuate conversations' that transpire between an agent and her ideal advisor. To illustrate, suppose that one raises the following conditional fallacy worry for Manne's view. Mr. Thirsty wants to drink gin, and he believes that the glass in his hand is filled with gin. But in fact the glass is full of petrol, which Mr. Thirsty would hate to drink. Intuitively, Mr. Thirsty has a reason to inquire into the contents of the glass. But the following counter-factual seems false: if Mr. Thirsty were reasoned with by an ideal advisor, then he would be motivated to inquire into the contents of the glass. For at the end of a conversation with his fully informed advisor, Mr. Thirsty will have been told about the true contents of the glass, and thus will not be the least bit motivated to inquire into the matter.

Manne replies to this conditional fallacy worry by suggesting that 'conversations should be individuated in a more fine-grained way.'⁵⁴ There is a (short) conversation between Mr. Thirsty and his ideal advisor that goes something like this:

⁵³ Manne (2014: 105).

⁵⁴ *Ibid.*, 105.

Ideal Advisor: The stuff in that glass is not what you think it is. Would you like to know what's really in it?

Mr. Thirsty: Yes.

Ideal Advisor: Then I recommend that you inquire into the contents of the glass.

Since Mr. Thirsty is motivated to inquire into the contents of the glass by the end of this short conversation between him and his ideal advisor, Manne suggests that her view does not entail the counter-intuitive claim that Mr. Thirsty does not have a reason to inquire into the contents of the glass. It therefore seems that Manne's considered judgment is that an agent has a reason to ϕ only if there is *some* (perhaps short) conversational process of being reasoned with by an ideal advisor about what to do, such that if the agent were to go through this process, then she would end up in a state of being motivated to ϕ . And on this refined statement of MI, the above conditional fallacy worry is not a problem.

I agree that this general strategy for responding to conditional fallacy worries provides a plausible line of response to the specific conditional fallacy worry just discussed. But it seems hopeless as a response to the conditional fallacy worries that I have offered above. For Ms. Introvert (we may now stipulate) becomes exhausted after even the shortest of social interactions; and Ms. Proudful becomes depressed immediately upon conversing with somebody who is superior to her with respect to procedural rationality. So, Manne cannot avoid the specific conditional fallacy worries that I have raised by being careful about how she individuates conversations between agents and their ideal advisors. I conclude that we have excellent reason to believe that Manne's argument for MI must go wrong somewhere.

1.3.3 Identifying the Flaw in Manne's Argument for MI

Manne's argument for MI is as follows:

- (1) A reason for an agent S to ϕ is a consideration which would be apt to be cited in favor of S's ϕ -ing by her ideal advisor, who is reasoning with her in an ideal way about what she ought to do.
- (2) There is a consideration which would be apt to be cited in favor of S's ϕ -ing by her ideal advisor, who is reasoning with her in an ideal way about what she ought to do, only if S would end up in a state of being (somewhat) motivated to ϕ following an idealized process of being reasoned with in this way.
- (3) So, an agent S has a reason to ϕ only if S would end up in a state such that she would be (somewhat) motivated to ϕ , following an idealized process of being reasoned with in this way.⁵⁵

My position is that premise (1) is false. Underlying premise (1) is Manne's 'practice-based approach' to practical normativity. On this approach, abstract normative notions are analyzed in terms of their role within some type of *normative behavior*—that is, 'behavior by means of which we give voice to ideas about what to do, and also what should happen.'⁵⁶ Premise (1) expresses the thought that the concept of a *reason* is correctly analyzed in terms of its role within the normative behavior of reasoning with someone about what to do. But this premise should be rejected.

First, consider Manne's support for premise (1). So far as I can see, she gives only the following argument in favor of it:

The above proposal [that reasons just are the considerations that would be apt to be cited by one's ideal advisor] seems to me attractive partly insofar as it secures a close connection between reasons for action and the activity of reasoning with a person about what she ought to do. It is hard to believe that the entities and the activity could come too far apart. Surely the connection goes deeper than the common etymological root of the corresponding English words.⁵⁷

I fully agree that there must be some non-etymological explanation of the fact that there is a 'close connection' between reasons and the activity of reasoning with a person about what she ought to

⁵⁵ *Ibid.*, 109. Manne presents her argument in more detail, but the details do not matter for my purposes.

⁵⁶ *Ibid.*, 94.

⁵⁷ *Ibid.*, 98.

do. I also agree that premise (1) provides such an explanation. But this counts as a strong reason to accept (1) only if it is the best explanation of the *explanandum*, and it is far from clear that this is so. While Manne proposes to explain this connection by analyzing reasons in terms of the activity of reasoning, we could alternatively explain this connection by analyzing the activity of reasoning in terms of reasons. For instance, the following account of the nature of reasoning with someone seems plausible:

Reasoning: Agent S_1 's reasoning with agent S_2 about what S_2 ought to do just is the activity in which S_1 helps S_2 come to recognize some consideration as a reason for action.

This analysis accounts for the fact that reasons and reasoning with someone are closely connected, and it obviously does not require the truth of premise (1). The availability of this account therefore greatly reduces the strength of Manne's argument for premise (1).

In addition to the failure of Manne's support for premise (1), there are decisive reasons to reject it. Premise (1) entails the following:

(A) If R is a reason for an agent S to ϕ , then R is a consideration which would be apt to be cited in favor of S 's ϕ -ing, by her ideal advisor, who is reasoning with her in an ideal way about what she ought to do.

And (A) is subject to the same sorts of counterexamples that we saw to afflict MI. To illustrate, consider for a final time the case of Ms. Introvert. The following is surely true: *That Ms. Introvert is lonely and craving social interaction* is a reason for her to call her friend. But the following is false: *That Ms. Introvert is lonely and craving social interaction* is a consideration which would be apt to be cited in favor of Ms. Introvert calling her friend, by her ideal advisor, who is reasoning with Ms. Introvert in an ideal way about what she ought to do. For once the ideal advisor is reasoning with Ms. Introvert, Ms. Introvert is neither lonely nor craving social interaction, and it

would not be apt for an ideal advisor to appeal to a false claim while reasoning with Ms. Introvert about what to do. (A) is therefore false. And so premise (1) should be rejected. I conclude that Manne has failed to provide us with good grounds for accepting her version of reasons internalism.

1.3.4 Markovits' Reasons Internalism

Julia Markovits argues for the following version of internalism:

Julia Markovits' Internalism (JMI): What it is for there to be a reason for agent S to perform some action ϕ is for it to be the case that S's desires are such that there is some fact that shows ϕ -ing to be relevantly related to these desires.⁵⁸

Markovits understands an agent's *desires* in a very broad way, such that they include the various motivational states that make up what Bernard Williams calls the 'motivational set.'⁵⁹ This 'may include, in addition to ordinary present desires, evaluations, attitudes, projects, commitments, etc.'⁶⁰ She does not offer an exhaustive account of what it means for ϕ -ing to be *relevantly related* to S's ends, but she does offer some examples of relevant relations. A fact shows that ϕ -ing is relevantly related to S's ends if any of the following is true: (i) the fact shows that ϕ -ing is a means to the fulfillment of one of S's desires, (ii) the fact shows that ϕ -ing is constitutive of the fulfillment of one of S's desires, and (iii) the fact shows that ϕ -ing is valuable in consequence of the value of the fulfillment of one of S's desires.⁶¹

⁵⁸ Markovits (2011, 2014, 2016a, 2016b).

⁵⁹ Williams (1979: 105).

⁶⁰ Markovits (2011: 258)

⁶¹ Markovits (2014) sometimes writes as if her internalist thesis is a thesis about our reasons to perform certain *actions*, and that's how I will understand it in what follows (52). But she sometimes writes as if her internalist thesis is a thesis about a broader class of reasons that includes reasons to have certain intentions, to adopt certain ends, to protect certain things, and to respect certain things (56, 59).

1.3.5 Markovits' Analogical Argument

Markovits sets up her first argument for JMI by inviting us to consider what internal and external reasons *for belief* would have to be like.⁶² Since an *internal reason for action* is a reason for action that depends on the agent's desires in the way specified by JMI, an *internal reason for belief* would have to be a reason for belief that favors believing *p* in virtue of showing that believing *p* stands in consistency and coherence relations with the rest of the agent's beliefs. For example, 'the fact that shining a coherent beam of light through two parallel slits produces a certain pattern on the screen behind them may give the physicist performing the experiment an internal reason to believe that light is a wave.'⁶³ This would be so if the fact that shining the beam through the slits produces this pattern shows that the belief that light is a wave coheres well with the rest of the physicist's beliefs. Similarly, since an *external reason for action* is a reason for action that is independent of the agent's desires in the way ruled out by JMI, an *external reason for belief* would have to be a reason for belief that favors believing *p* independently of whether it shows belief in *p* to stand in consistency and coherence relations with the rest of the agent's antecedent beliefs.

Markovits argues that there are some external reasons for belief. She writes:

It seems like, for *some P*, there are *external* reasons to believe *P* – considerations that count in favor of believing *P* (by providing evidence for *P*) but *not* in virtue of the relation they show *P* to stand in to our antecedent beliefs. Here are some plausible examples: The fact that I feel pain seems to give me a reason to believe *I'm in pain*, regardless of what else I already believe. And the fact that I have an experience of redness gives me a reason to believe *I'm having an experience of redness* regardless of what else I believe.⁶⁴

Markovits then observes that these plausible examples of external reasons for belief have something in common: they favor beliefs that 'are not merely uncontroversial, but tend to be self-

⁶² Markovits (2011, 2014) defends this argument.

⁶³ *Ibid.*, 60.

⁶⁴ *Ibid.*

evident, incorrigible, indubitable, or in some other way plausibly immune to error.’⁶⁵ For example, the belief that *I’m in pain* seems to be a belief of this sort since it seems impossible for me to believe that I’m in pain while not actually being in pain. On the basis of reflection on these cases of beliefs that are supported by external reasons for belief, Markovits infers that *all* external reasons for belief are such that the beliefs that they favor are beliefs that are in some way immune to error.⁶⁶ For ease of exposition, I will say that a reason is an *infallible reason* if and only if that reason favors something (such as a belief or an action) in such a way as to make it in some way immune to error, such as by making it self-evident, indubitable, or incorrigible. Markovits’ claim can then be put as follows: All reasons for belief that qualify as external reasons are infallible reasons.

Having made these observations, Markovits presents the following *analogical argument for JMI*. Because the only reasons for belief that qualify as external reasons are infallible reasons, we have some grounds for believing that the parallel claim about reasons for action is true. That is, we have some grounds for believing that the only reasons for action that qualify as external reasons are infallible reasons. Such reasons would be reasons that render some action immune to error in some way. But there are no actions that are immune to error in any relevant way.⁶⁷ So, we have some grounds for believing that there are no reasons for action that qualify as external reasons. And so, we have some grounds for accepting JMI.

I will grant the following premises for the sake of the argument: (i) that there are some external reasons for belief, (ii) that such reasons would have to be reasons for belief that are

⁶⁵ *Ibid.*, 62.

⁶⁶ *Ibid.*, 62-63.

⁶⁷ As Markovits writes: ‘Are there any ends that are uncontroversial, largely immune to erroneous adoption, and therefore not the kinds of things we feel people must offer further justification for caring about, beyond telling us they care about them? I don’t think there are any such ends’ (2014: 63).

independent of the agent's antecedent beliefs, (iii) that the only external reasons for belief that exist are infallible reasons, and (iv) that there are no reasons for action that are infallible reasons.⁶⁸

I will focus my criticisms on the following step of the analogical argument for JMI:

Analogy Premise: Because the only reasons for belief that qualify as external reasons are infallible reasons, we have some grounds for believing that the parallel claim about reasons for action is true.

Markovits does not explicitly argue for this crucial premise, but there is a natural line of reasoning in its support that I think she has in mind, and which goes as follows. The Analogy Premise follows from the following principle:

Practical-Theoretical Analogy Principle: If something is true of reasons for belief, then this counts as some grounds for believing that the same is true of reasons for action.⁶⁹

And the Practical-Theoretical Analogy Principle is at least *prima facie* plausible.⁷⁰ In its support, one can appeal to a general presumption in favor of unified accounts of a given phenomenon over those accounts that are less unified. If this general methodological approach is sound, then it would follow that there is some presumption in favor of holding that reasons for action are a certain way

⁶⁸ But it is worth making two comments about these premises. First, I find (ii) unmotivated, and Markovits never defends it. In the context of most discussions of normative reasons, calling a reason an *external* reason amounts to saying that it's in some way independent of the desires (broadly construed) of the agent whose reason it is. It therefore seems quite natural to understand external reasons for belief as reasons for belief that are in some way independent of the desires (as opposed to beliefs) of the agents whose reasons they are.

Second, premise (iii) would be rejected by many epistemologists today. For example, foundationalists who hold that the fact that a belief was produced by properly-functioning cognitive faculties counts as an external reason for that belief would deny that all external reasons for belief are infallible reasons. Markovits acknowledges this point, and in response briefly presents an alternative version of the analogical argument for JMI that avoids commitment to (iii) – see Markovits (2011: 271-272) and Markovits (2014: 63-64). My criticism of Markovits' main version of the analogical argument for JMI also applies to this alternative version. Thanks to William Melanson for helpful discussion on this point.

⁶⁹ Brunero (forthcoming: 3-11) also interprets Markovits to be employing something like this Practical-Theoretical Analogy Principle.

⁷⁰ There is surprisingly little reflection on the theoretical-practical analogy principle in the literature, but see Brunero (forthcoming: 3) and Finlay and Schroeder (2012) for brief discussions. The brief support for the theoretical-practical analogy principle that I give in the main text is quite similar to the support that these authors offer in its favor.

if reasons for belief are that way; for an account of normativity is more unified the more similar that reasons for action and reasons for belief are to each other on that account. Since there's a presumption in favor of the Practical-Theoretical Analogy Principle, there's a presumption in favor of the Analogy Premise.

Assume for the moment that the Practical-Theoretical Analogy Principle is true. It would follow that there are grounds for believing that there are some reasons for action that are independent of the desires of the agent whose reasons they are. To see this, consider for a second time Thomas Kelly's case about Bertrand Russell. If you stumble across strong evidence that Russell was left-handed, then it seems clear that you have strong reasons to believe that he was left-handed – *regardless of what desires you have*. Intuitively, there is reason for you to believe that Russell was left-handed even if you have no desires that are served by so believing. Since this is so, it is a truth about reasons for belief that some of them obtain independently of the desires of the agent whose reasons they are. And so, on the assumption that the Practical-Theoretical Analogy Principle is true, we should think that there are some grounds for believing that it is a truth about reasons *for action* that some of them obtain independently of the desires of the agent whose reasons they are. Thus, on the assumption that the Practical-Theoretical Analogy Principle is true, there are grounds to *reject* JMI. Let this objection to JMI be called *the analogical argument against JMI*.

I have argued that if the Practical-Theoretical Analogy Principle is true, then the analogical argument against JMI is sound. And again, on my understanding of Markovits' analogical argument for JMI, it is sound only if the Practical-Theoretical Analogy Principle is true. Therefore, Markovits' analogical argument for JMI is sound only if the analogical argument against JMI is sound. This makes the analogical argument for JMI unhelpful as a source of support for internalism about reasons for action.

1.3.6 Markovits' Argument from Rationality

Another one of Markovits' arguments for JMI begins with the following observation: 'only rational creatures have reasons.'⁷¹ That is, the only beings who are such that there can be reasons for them to perform certain actions are beings capable of advanced rational choice and thought. Adult human beings have reasons; but tables, chairs, acorns, ice cream cones, bugs, dogs, and babies do not. These observations form the basis of Markovits' *argument from rationality*:

- (1) Only rational creatures have reasons.
- (2) JMI provides an explanation of (1).
- (3) Externalism about reasons for action does not provide an explanation of (1).
- (4) If there are two competing theories, T_1 and T_2 , of domain D , and T_1 explains a significant fact F within D whereas T_2 does not, then F counts significantly in favor of accepting T_1 over T_2 .
- (5) So, the fact that only rational creatures have reasons counts significantly in favor of accepting JMI over externalism about reasons for action.⁷²

I grant premises (1), (3), and (4). Premise (2) is mistaken. Far from explaining why (1) is true, JMI implies that (1) is false. To see that this is so, consider a typical dog that I will call 'Fido.' Fido has desires, such as the desire to eat. Fortunately for him, he hears that his owner has just filled his bowl with food. Consequently, there is some fact that shows that the action of running to his bowl is a means to the fulfillment of one of Fido's desires. In particular, the fact that there is food in his bowl is such a fact. So, Fido's desires are such that there is some fact that shows that the performance of some action is relevantly related to these desires. According to JMI, this means that there is a reason for Fido to perform that action. So, JMI implies that there is a reason for Fido to perform a certain action. As Markovits acknowledges, Fido is not a rational creature.⁷³

⁷¹ Markovits (2011: 273).

⁷² *Ibid.*, 273-276. I am filling in some details to make the argument from rationality more precise. In particular, Markovits does not explicitly state premise (4), but she is clearly taking it for granted.

⁷³ Markovits makes it clear that she does not consider non-human animals to be rational creatures (*Ibid.*, 273).

Therefore, JMI implies that premise (1) of the argument from rationality is false. Since a theory that implies $\sim p$ cannot be such as to provide an explanation of why p , premise (2) must be false.⁷⁴

It might be objected that JMI cannot imply that there is a reason for Fido to run to his bowl, for Fido is not a true agent. If Fido is not an agent, then JMI couldn't imply that there are reasons for Fido to act in certain ways, for JMI is a theory about what it is for there to be a reason *for an agent* to do something. Markovits is not clear about what she means by an 'agent.' But if she is using the term 'agent' in such a way that non-human animals cannot count as agents, then JMI would not imply that non-human animals lack reasons. And if JMI does not imply that non-human animals lack reasons, then JMI could not provide an explanation of why non-human animals lack reasons, in which case premise (2) of the argument from rationality would be false. So, the internalist cannot defend the argument from rationality by arguing that Fido and other non-rational, desire-bearing creatures are not true agents.

One might also challenge my objection's premise that typical dogs are not rational creatures. Typical dogs can engage in some impressive reasoning: for example, they can infer that the sound of the door opening means that somebody is at the door, and they can decide to wait until their owner is out of the room before trying to open the refrigerator. It may be argued that such capacities are sufficient to make Fido a rational creature. From this it would follow that the fact that JMI implies that Fido has a reason for action does not show that JMI implies that premise (1) is false. Markovits herself would not press this objection since she affirms that dogs are not

⁷⁴ In light of my critique of premise (2) of the argument from rationality, it seems that the sorts of considerations that Markovits marshals in developing this argument actually furnish us with the following argument *against* JMI: (1) Only rational creatures have reasons; (2) JMI implies that (1) is false; (3) So, JMI is false. By my lights, this argument against JMI is not compelling since I am inclined to think that some beings that are not rational creatures (such as dogs and human babies) do have reasons. But Markovits should find this argument against JMI compelling since she accepts that only rational creatures have reasons, and since I have just demonstrated that JMI implies otherwise.

rational creatures, but a proponent of her argument from rationality might insist that this is something that she ought not to affirm.⁷⁵

In response, I will note that the basic thought behind my objection to the argument from rationality is that there is something that fulfills the following conditions: (i) it is not a rational creature, (ii) it has some desire *D*, and (iii) there is some action that's such that *D* would be served were it to perform that action. My basic objection works so long as there is some possible being that meets conditions (i)-(iii). If my critic is not convinced that Fido satisfies all of these conditions, then she should consider whether there is some other possible being that does satisfy them. I suggest that these conditions are met by many actual and conceivable beings, such as some non-human animals who are less cognitively sophisticated than dogs, as well as some human infants.

The objection that I have raised against the argument from rationality is fairly obvious, and it may seem odd that it would have been overlooked by Markovits; but I can explain this oversight. Markovits presents the argument from rationality in her 2011 article, 'Why Be an Internalist about Reasons?' In this article, she is actually discussing two distinct versions of internalism. JMI is her favored version.⁷⁶ The other internalist thesis she discusses is as follows:

Julia Markovits' Internalism 2 (JMI2): What it is for there to be a reason for agent *S* to perform some action ϕ is for it to be the case that *S* would be motivated to ϕ were she perfectly employing her faculties of procedural rationality.

For an agent to be such that *she would be motivated to ϕ were she perfectly employing her faculties of procedural rationality* is for it to be the case that (i) she has rational faculties, and (ii) she would be motivated to ϕ were she fully procedurally rational. And for an agent to be *fully procedurally rational* is (in part) for her set of desires to perfectly conform to the 'requirements of internal

⁷⁵ Thanks to William Melanson for pressing this objection.

⁷⁶ *Ibid.*, 260-261.

consistency and coherence.⁷⁷ For example, a fully procedurally rational agent would be inclined to form a desire to take the means necessary to satisfy other desires that she has; for the formation of such a desire would make her set of desires exhibit a higher degree of coherence.

Markovits claims that the argument from rationality (along with all of the arguments that she offers for internalism) can serve as support for both JMI and JMI2.⁷⁸ But when she is developing and defending the argument from rationality, her discussion is clearly offered in terms of JMI2.⁷⁹ It seems that Markovits' assertion that the argument from rationality serves as support for JMI as well as for JMI2 is due to a failure to think through how this argument would have to go in order to truly support JMI. I have argued that once we think this through carefully, it is clear that this argument is not successful.

Of course, it may be suggested that what I have shown is that internalists such as Markovits should just endorse JMI2 rather than JMI, for I have not yet offered any reason to believe that the argument from rationality fails to support the former. And in fact, the argument from rationality works much better as an argument for JMI2. Whereas JMI fails to explain why only rational creatures can have reasons, JMI2 easily explains this fact. For on JMI2, part of what it is for there to be a reason for one to do something is for one to have rational faculties. Since chairs, tables, acorns, dogs, and babies lack rational faculties, JMI2 accounts for why there cannot be reasons for such beings to perform actions. So for all that I have said, the argument from rationality provides us with grounds for accepting some kind of internalism about reasons for action.

⁷⁷ *Ibid.*, 258.

⁷⁸ *Ibid.*, 261.

⁷⁹ For example, she writes: 'Why is it that only creatures who can reason have reasons? Internalism explains this by offering an account of reasons as facts about how we would be motivated if we were properly employing our rational faculties' (*Ibid.*, 273).

I concede that the argument from rationality likely does provide us with some reason to accept JMI2. But I will now argue that there is also decisive reason to reject JMI2, and that consequently, JMI2 will not help my internalist opponents.

Consider two objections against JMI2. The first objection is that JMI2 is vulnerable to well-known conditional fallacy worries.⁸⁰ To see this, consider cases of procedurally irrational agents. Suppose that Mr. Irrational is quite procedurally irrational. For example, he frequently desires to do an act without desiring to perform the means necessary to performing that very act. As a result, most of his endeavors fail miserably. Mr. Irrational is aware that he could receive cognitive psychotherapy at a nearby hospital that would make him much more procedurally rational, and he is aware that this would make him far more successful and happy. Intuitively, there is a reason for Mr. Irrational to undergo cognitive psychotherapy. But it is not the case that Mr. Irrational would be motivated to seek cognitive psychotherapy were he perfectly employing his faculties of procedural rationality. For if he were perfectly employing his faculties of procedural rationality, he would already be maximally procedurally rational, and would therefore have no need for the cognitive psychotherapy. So, the case of Mr. Irrational counts as a counter-example to JMI2.

To set up my second objection to JMI2, consider the fact that our reasons for action very frequently conflict. That is, it is very often the case that there is a reason for us to do a certain action while there is also a reason for us not to do that very action. In fact, nearly every action that we can do is such that there is at least some reason to do it and some reason not to do it. For example, there is a reason for me to go shopping right now (*I need groceries*), and there is a reason for me not to go shopping right now (*I need to finish a paper and I cannot finish it if I go shopping right now*). There is a reason for me to eat another piece of cheesecake (*doing so would be*

⁸⁰ JMI2's vulnerability to conditional fallacy worries is what motivates Markovits (2011) to favor JMI (260).

pleasurable), and there is a reason for me not to eat another piece of cheesecake (*doing so would ruin my diet*). And there is a reason for me to eat my car (*I need iron in my body to survive*), and there is a reason for me not to eat my car (*doing so will kill me*). Reflection on our conflicting reasons suggests that the following thesis is true:

Ubiquity of Conflicting Reasons (UCR): It is extremely common for a person's reasons to conflict – that is, for there to be a reason for her to do an act while there is also a reason not to do that very act.

Any adequate theory of reasons must be able to accommodate UCR.

My second objection to JMI2 is that it cannot accommodate UCR. To explain why this is so, I need to say a bit more about Markovits' conception of procedural rationality. As I mentioned above, Markovits understands being procedurally rational to involve having a set of desires that conforms to the requirements of consistency and coherence. An example of a change in a person's set of a desires that these requirements would favor is the adoption of a desire to take the means necessary to satisfy other desires that she has. But Markovits holds that procedural rationality is not limited to this kind of instrumental reasoning:

Being procedurally rational involves much more than mere instrumental rationality. And as in the case of beliefs, sets of ends can exhibit looser procedural 'virtues' than mere consistency: considerations of coherence and systematic justifiability, as well as inference to the best explanation, can make it more rational for us to abandon certain ends and adopt others.⁸¹

Markovits follows Michael Smith in spelling out the way in which considerations of 'systematic justifiability' can lead to procedurally rational changes in one's ends.⁸² The idea is that a procedurally rational agent has a set of ends that is the product of a procedure that Smith calls the

⁸¹ Markovits (2014: 69).

⁸² *Ibid.*, 132-33.

‘procedure of systematic justification.’⁸³ This is a process of increasing the coherence and unity of one’s set of desires that is analogous to Rawls’ reflective equilibrium. One is supposed to look at one’s set of desires and consider whether it would be on-the-whole more coherent and unified if one were to add to it a further (general or specific) desire. If so, it is part of the procedure of systematic justification to add that desire to one’s set of desires. Similarly, if dropping a certain desire would increase the coherence and unity of one’s set of desires, then the procedure of systematic justification will involve dropping this desire. An agent who completes this procedure of systematic justification ends up with a ‘maximally coherent and unified’ set of desires.⁸⁴

Markovits understands perfect conformity to the requirements of procedural rationality to include reaching this coherent and unified set of desires. As she puts it, ‘the virtue of willing a mutually supportive set of ends is a virtue of procedural practical rationality, as the analogy to the epistemic case once again helps bring out: it’s a matter (at least in the first instance) of my ends’ standing in the right *relations to each other*, not *simply* of my holding or failing to hold a particular end.’⁸⁵

I can now state my second objection to JMI2. Suppose that JMI2 and UCR are both true. It would follow that a lot of actions are such that were we perfectly employing our faculties of procedural rationality, we would both be motivated to perform them and be motivated not to perform them. On the assumption that one can be motivated to do something only by having some desire (broadly construed, to include anything in one’s ‘motivational set’) that is served by doing it, it follows that a lot of actions are such that were we perfectly employing our faculties of procedural rationality, we would both have a desire served by performing them and have a desire

⁸³ Smith (1995: 114).

⁸⁴ *Ibid.*, 116.

⁸⁵ Markovits (2014: 133).

served by not performing them. But it's not true that a lot of actions are such that were we perfectly employing our faculties of procedural rationality, we would both have a desire served by performing them and have a desire served by not performing them. For one who is perfectly employing one's faculties of procedural rationality has a maximally unified and coherent set of desires, and a maximally unified and coherent set of desires would not include a lot of desires that stand in tension with each other. It follows from the above reasoning that it is not the case that both UCR and JMI2 are true. And as I argued, UCR is obviously true upon reflection. So, JMI2 is false.

To sum up this section, Markovits can either offer the argument from rationality as support for JMI or for JMI2. If she offers it as support for JMI, then her argument fails. If she offers it as support for JMI2, then her argument does succeed in the sense that it establishes some grounds for accepting JMI2. But JMI2 is not a plausible version of internalism, for it is vulnerable to conditional fallacy worries and it fails to accommodate UCR. I conclude that the argument from rationality is not helpful support for reasons internalism.

1.3.7 Markovits' Argument from Motivational Disparity

Consider a final argument for JMI that Markovits offers.⁸⁶ This last argument seeks to highlight a certain explanatory advantage of JMI that arises from the following consideration:

Disparity Thesis (DT): Some people are better than others at being motivated to perform acts for which there are reasons. In other words, some people respond to their reasons better than others.

⁸⁶ In what follows, I discuss an argument for JMI that I interpret Markovits to be offering in section 3.2 of *Moral Reason*. I believe that there are at least a couple of distinct arguments for JMI that are run together in this section, but I will focus on what I take to be the most formidable one.

DT certainly seems credible. Some people intentionally cultivate friendships when they are lonely, whereas others never bother to reach out to others. Intuitively, the former people are motivated to act in accordance with their reasons, whereas the latter are not. Similarly, some addicts manage to break out of their addiction, whereas others never do. The former seem to have succeeded in responding to their reasons to get off drugs, whereas the latter are failing to act as they have reason to act. So DT is extremely plausible.

According to Markovits, externalism has no acceptable explanation of DT:

But externalists can offer no explanation for this supposed difference in how well we respond to reasons – no explanation of why some of us have the right motivations and some of us have the wrong ones – that does not itself appeal to the views about what matters that they’re trying to justify.⁸⁷

Here, Markovits does not say that externalists can offer *no* explanation of DT.⁸⁸ But she does suggest that all externalist explanations of DT are problematic or suboptimal in virtue of the fact that they appeal to controversial views about ‘what matters normatively – that is, what we have reason to do or pursue or protect or respect or promotes.’⁸⁹ To illustrate the sort of deficient explanation that is available to the externalist, Markovits observes that externalists ‘can explain why some people have the right motivation by saying, for example, that they’re good people.’⁹⁰ But she claims that this explanation ‘assumes the truth of the normative views that are at issue,’ for *good person* is presumably meant to be analyzed partly in terms of how well the person responds to external reasons.

⁸⁷ *Ibid.*, 56.

⁸⁸ But see Markovits (2016: 521), where she does seem to suggest that externalism offers no explanation ‘of what makes some people better at responding to reasons than others.’

⁸⁹ Markovits (2014: 56).

⁹⁰ *Ibid.*, 56.

In contrast, internalism ‘paints a different and more informative picture of what’s going on when some people are more responsive to genuine reasons than others.’⁹¹ Specifically, internalism can provide the following explanation of why one person is better than another at responding to her reasons: the first person is more procedurally rational than the second person.⁹² To see why this explanation is available to the proponent of JMI, it is helpful to consider how procedural irrationality can undermine our ability to be motivated to act on the reasons that JMI entails that we have. So consider the following case:

Alice deeply desires to work as a lawyer. So, JMI entails that facts that show an act ϕ to stand in a relevant relation to Alice’s desire to work as a lawyer (for example, by showing that ϕ -ing is a necessary means to the fulfillment of this desire) are reasons for Alice to ϕ . The fact that Alice’s grades are mediocre shows that *studying hard for the LSAT* stands in a relevant relation to Alice’s desire to work as a lawyer – specifically, the relevant relation of being a necessary means to Alice’s desire to work as a lawyer. So, JMI entails that the fact that Alice’s grades are mediocre is a reason for Alice to study hard for the LSAT. But (let us suppose) Alice is less than fully procedurally rational in that she fails to desire to study hard for the LSAT, despite the fact that studying hard for the LSAT is a necessary means to fulfilling her deep desire to work as a lawyer. So, on the assumption that JMI is true, Alice would be more motivated in accordance with her reasons were she more procedurally rational.

Cases like this one support Markovits’ claim that ‘internalism holds that being procedurally rational will make a person more likely to respond appropriately to her reasons.’⁹³ And if internalism does hold this, then it is plausible that Markovits is correct in suggesting that the internalist can explain DT by positing that when one person is better at responding to their reasons than another, this is so in virtue of the fact that the former person is more procedurally rational.

It will be helpful to present Markovits’ argument as follows:

(1) The Disparity Thesis is true.

⁹¹ *Ibid.*, 57.

⁹² *Ibid.*, 57-58.

⁹³ *Ibid.*, 58.

- (2) Externalism cannot offer a good explanation of the Disparity Thesis.
- (3) JMI can provide a good explanation of the Disparity Thesis.
- (4) If (1)-(3) are true, then the Disparity Thesis counts as grounds to accept JMI over externalism.
- (5) So, the Disparity Thesis counts as grounds to accept JMI over externalism.

Call this the *argument from motivational disparity*. Premises (1) and (4) seem true, and I will grant them in what follows. As I have already explained, premise (2) is meant to be plausible because of Markovits' observations that (i) all externalist explanations of the differences between people's abilities to respond to their reasons must appeal to controversial views about what matters, and (ii) explanations that appeal to controversial views about what matters are not good explanations. And I have also already explained Markovits' support for premise (3): the proponent of JMI can explain DT by appealing to differences in people's levels of procedural rationality.

I shall object to the argument from motivational disparity by arguing that the internalist faces a certain dilemma that forces her to reject either premise (2) or premise (3). For ease of exposition, I will introduce a toy externalist view that I will call *hedonistic externalism* in order to explore how externalists might explain DT. Hedonistic externalism is the conjunction of the following theses:

- (a) What it is for a fact F to be a reason for agent S to do act X is for it to be the case that F explains why S's X-ing would bring about a good state of affairs.
- (b) A state of affairs SA is good if and only if SA is a state of affairs in which some agent experiences pleasure.

To clarify hedonistic externalism, consider a couple of its implications. It implies that the fact that one's spouse is sad is normally a reason to hug her. For normally, the fact that one's spouse is sad explains why one's hugging her would bring about a state of affairs in which she experiences some emotional pleasure. Hedonistic externalism also implies that the fact that one is hungry is normally a reason to eat something. For normally, the fact that one is hungry explains why one's eating

something would bring about a state of affairs in which one experiences some physical pleasure. Hedonistic externalism is an *externalist* theory since it grounds all reasons for action in evaluative facts and facts about what causes pleasure rather than in facts about our desires.

Since hedonistic externalism is a type of externalism, premise (2) implies that hedonistic externalism cannot offer a good explanation of DT. For (according to Markovits) any explanation of DT offered by the hedonistic externalist will inevitably appeal to controversial views about what matters, and explanations of this sort are not good explanations. Now, the hedonistic externalist has a response to this line of reasoning. To see this, consider a hedonistic externalist who offers the following explanation of DT:

E1: Some people are more disposed to promote people's pleasure than others.

E1 seems to count as a plausible explanation of DT, on the assumption that hedonistic externalism is true. For on hedonistic externalism, E1 clearly implies DT. Furthermore, E1 makes no reference to controversial views about 'what normatively matters'; rather, E1 only refers to dispositions to promote pleasure. So we can imagine the hedonistic externalist challenging premise (2) on the grounds that the hedonistic externalist can offer E1 as an explanation of DT. Since E1 makes no reference to controversial views about what matters, E1 is a good explanation of DT, at least for all that Markovits has said. And of course, analogous explanations of DT could be constructed for more plausible and nuanced externalist theories. A critic of Markovits' might therefore conclude that premise (2) of the argument from motivational disparity is false.

Markovits has an obvious reply to this challenge to premise (2). Upon reflection, the hedonistic externalist's proposed explanation of DT does not consist merely of E1, but rather consists of the conjunction of E1 and:

E2: Hedonistic externalism is true.

One can see that E1 is probably not by itself a sufficient explanation of DT by noticing that it is possible to be reasonably puzzled about why DT is true even if one knows that E1 is true. In contrast, puzzlement about why DT is true should dissipate if one learns that E1&E2 is true. So, Markovits may reasonably suggest that only E1&E2 counts as an acceptable explanation of DT. And since E2 is a controversial view about what matters, the hedonistic externalist has failed to provide a good externalist explanation of DT—that is, an explanation of DT that does not assume the truth of controversial views about what matters.

But now notice a certain structural similarity between the hedonistic externalist's explanation of DT and the internalist explanation favored by Markovits. Markovits offers the following explanation of DT:

I1: Some people are more procedurally rational than others.

But every reason that I can think of for believing that E1 by itself is not a sufficient explanation of DT also seems like a reason to believe that I1 is not by itself a sufficient explanation of DT. For example, it seems possible to be reasonably puzzled about why DT is true even if one has learned that I1 is true. In particular, those of us who don't subscribe to some meta-normative theory (such as JMI) that ties our practical reasons closely to the requirements of procedural rationality can be reasonably puzzled about why DT is true even if we know that I1 is true. So, if we were convinced that E1 cannot explain DT without E2, then we should also believe that I1 cannot explain DT without:

I2: JMI is true.

Markovits therefore faces some pressure to hold that her explanation of DT is really I1&I2, rather than just I1—at least so long as she insists that E1 is not a sufficient explanation of DT.

I can now state my dilemma for the argument from motivational disparity. As I have argued, we should believe that if E1 by itself is an insufficient explanation of DT, then I1 by itself is also an insufficient explanation of DT. It follows that we should believe that either (i) E1 by itself is *not* an insufficient explanation of DT, or (ii) I1 by itself *is* an insufficient explanation of DT. If (i) is true, then for all Markovits has shown, there are externalist explanations of DT that do not assume the truth of controversial views about what matters. And if there are such externalist explanations of DT, then premise (2) of the argument from motivational disparity is undermined; for Markovits' only support for (2) is her claim that there are no externalist explanations of DT of this sort. It follows that if (i) is true, then premise (2) is undermined.

And if (ii) is correct, then this shows that premise (3) is undermined. For suppose that (ii) is correct. Then, the only internalist explanation of DT that Markovits offers is an insufficient explanation of DT, and insufficient explanations are not good explanations. Markovits could move to the position that JMI provides the following explanation of DT: I1&I2. But since I2 is a very controversial view about what matters, this internalist explanation of DT clearly assumes the truth of a controversial view about what matters. And (as Markovits claims, and as I am granting for the sake of the argument) explanations of DT that assume the truth of a controversial view about what matters are bad explanations. So, if (ii) is correct, we have good grounds for doubting premise (3).

It follows from the above reasoning that either premise (2) is undermined or premise (3) is undermined. Either way, we have good grounds for doubting the soundness of Markovits' argument from motivational disparity.

1.4 Conclusion

To sum up this chapter, the argument from epistemic duties provides us with excellent grounds for believing in external reasons for actions, and the most formidable arguments against such reasons – those developed by Kate Manne and Julia Markovits – are unsuccessful. We should therefore believe that both the internalist argument for moral relativism and the internalist argument for moral error theory ought to be rejected. I conclude that moral realism is secure from these influential anti-realist arguments.

Chapter 2: A Defense of the Companions in Guilt Argument for Moral Facts

2.1 Introduction

A *moral fact* is a fact to the effect that something has some moral property. If there are moral facts, then either moral realism's success thesis is true, or none of the extant moral facts are representable via moral judgments. Since this latter disjunct is so implausible, moral realism's success thesis is very likely true if there are moral facts. The purpose of this chapter is to defend an argument for the existence of moral facts, thereby establishing that the success thesis is very likely true.

My strategy will be to defend a companions in guilt argument. This type of argument seeks to show that moral facts are very similar to some other type of fact whose existence is not in question, and that we should therefore believe in moral facts. An increasingly-popular version of the companions in guilt argument appeals to *epistemic facts*—that is, facts to the effect that some object has some epistemic property, such as being epistemically justified or being epistemically obligatory.⁹⁴ I will defend a version of this type of argument, arguing that it withstands the most formidable objections in the literature.

In the next section, I lay out my companions in guilt argument and offer a presumptive case in support of its premises. And in sections 2.3-2.5, I respond to objections to the companions in guilt argument that have recently been pressed in the literature.

2.2 The Companions in Guilt Argument

Discussions of the companions in guilt argument for moral facts typically focus exclusively on what I will call *doxastic epistemic facts*. These are facts to the effect that some doxastic attitude

⁹⁴ Cuneo (2007), Rowland (2013), and Stratten-Lake (2002).

(such as belief that p) has some epistemic property (such as being epistemically justified). Thus, the discussion tends to center on the idea that moral facts and doxastic epistemic facts are very similar—so much so that we ought to believe in the former since we ought to believe in the latter. This focus on doxastic epistemic facts is understandable since – as we saw in the previous chapter – some philosophers hold that doxastic epistemic facts are the only sort of epistemic facts that exist.

Nevertheless, I believe that the best version of this argument appeals to *practical epistemic facts* – facts to the effect that someone epistemically ought to perform some action. As I argue in sections 2.3 and 2.4, moving to this atypical formulation of the argument allows me to evade much of the force of two recent challenges to companions in guilt arguments. So the version of the companions in guilt argument that I favor is as follows:

- (1) **Parity Premise:** If moral facts do not exist, then practical epistemic facts do not exist.
- (2) **Ontological Premise:** Practical epistemic facts do exist.
- (3) Therefore, moral facts exist.

This is the *argument from practical epistemic facts*. I have already defended the Ontological Premise in the previous chapter.⁹⁵ Recall that I argued that we sometimes have certain epistemic duties of action by doing the following: (i) directing attention to cases in which it is intuitively plausible that we have certain epistemic duties of action; (ii) defending the argument from transmission for the claim that we have certain epistemic duties of action; and (iii) rebutting the best argument against the existence of epistemic duties of action. If there are epistemic duties of action, then there are some practical epistemic facts. So I have already offered a presumptive case

⁹⁵ See section 1.2.1.

in favor of the Ontological Premise, and I will further bolster this case in sections 2.4 and 2.5 by rebutting two important challenges to it.

The case in favor of the Parity Premise consists of showing that each promising argument against the existence of moral facts is such as to suggest some parallel argument against practical epistemic facts that is no less plausible than its moral counterpart. If this can be shown, then there would be no good reason to reject moral facts that isn't also a good reason to reject practical epistemic facts, and this would render the Parity Premise quite credible. In this section, I offer a presumptive case in favor of the Parity Premise by arguing that the four most formidable arguments against the existence of moral facts suggest equally-promising arguments against practical epistemic facts.

2.2.1 The Internalist Argument for Moral Error Theory

John Mackie thought that moral facts would be 'queer' were they to exist, and one source of the alleged queerness of moral facts is supposed to be their *objective prescriptivity*. The idea is that moral facts make demands on us from which we cannot escape: they entail that there are certain obligations that we are under, and practical reasons that we have, irrespective of our desires and goals. Mackie argues that such obligations and reasons would be objectionably queer, and on this basis concludes that moral facts are suspect.⁹⁶ This basic strategy has received its most formidable defense by Richard Joyce, whose argument is as follows.⁹⁷ For any person S and action ϕ :

- (1) If S morally ought to ϕ , then S morally ought to ϕ regardless of whether ϕ -ing serves any of her desires.
- (2) If S morally ought to ϕ regardless of whether ϕ -ing serves any of her desires, then there is a reason for S to ϕ regardless of whether ϕ -ing serves any of her desires.

⁹⁶ Mackie (1977: 38-42).

⁹⁷ Joyce (2001: 42).

- (3) So, if S morally ought to ϕ , then there is a reason for S to ϕ regardless of whether ϕ -ing serves any of her desires.
- (4) But it's never the case that there is a reason for a person to do some action regardless of whether doing that action serves any of her desires.
- (5) So, it's never the case that a person morally ought to do some action.

In the previous chapter, I referred to this argument as the *internalist argument for moral error theory*.⁹⁸ Recall that premise (2) follows from moral rationalism, which can be best supported with the moral argument from criticizability.⁹⁹ Premise (4) is an expression of reasons internalism, which can be defended with Bernard Williams' famous action-explanation argument, as well as the arguments defended by Manne and Markovits that we also considered in the previous chapter.¹⁰⁰ And Joyce defends (1) by observing that we ordinarily do not take back our judgment that someone ought to act in a certain way once we discover that acting in that way fails to promote her desires.¹⁰¹ For instance, we would not take back our judgment that the Nazis ought to have refrained from committing genocide if we learned that refraining from genocide would have only frustrated their desires. Premise (1) is therefore credible.

The internalist argument for moral error theory suggests an analogous argument against the existence of practical epistemic facts, which I will call the *internalist argument for epistemic error theory*. For any person S and action ϕ :

- (1) If S epistemically ought to ϕ , then S epistemically ought to ϕ regardless of whether ϕ -ing serves any of her desires.
- (2) If S epistemically ought to ϕ regardless of whether ϕ -ing serves any of her desires, then there is a reason for S to ϕ regardless of whether ϕ -ing serves any of her desires.
- (3) So, if S epistemically ought to ϕ , then there is a reason for S to ϕ regardless of whether ϕ -ing serves any of her desires.
- (4) But it's never the case that there is a reason for a person to do some action regardless of whether doing that action serves any of her desires.

⁹⁸ See section 1.1.

⁹⁹ See section 1.2.3 for an elucidation of the moral argument from criticizability.

¹⁰⁰ Manne (2014), Markovits (2011, 2014, 2016a, 2016b), and Williams (1979).

¹⁰¹ Joyce (2001: 42-43).

(5) So, it's never the case that a person epistemically ought to do some action.

In the previous chapter, I established that premises (1) and (2) are true.¹⁰² So, this argument is sound so long as its fourth premise is true. And its fourth premise is true if the internalist argument for moral error theory is sound, for this fourth premise is one of the premises of the internalist argument for moral error theory. It follows that if the internalist argument for moral error theory is sound, then the internalist argument for epistemic error theory is sound. I conclude that if the internalist argument for moral error theory succeeds in establishing that there are no moral facts, then there is an analogous argument that succeeds in showing that there are no practical epistemic facts. This support for the Parity Premise is especially important because the internalist argument for moral error theory is currently the most prominent argument for the non-existence of moral facts.¹⁰³

2.2.2 The Argument from Irreducible Normativity

In his recent defense of moral error theory, Jonas Olson argues that the strongest version of Mackie's argument from queerness invokes the notion of *irreducibly normative favoring relations*. For a fact F to stand in an irreducibly normative favoring relation to some behavior for person S is for it to be the case that (i) F favors that behavior for S, and (ii) the fact that F favors that behavior is not reducible to some natural fact, such as the fact that engagement in that behavior would promote S's desires. And for our purposes, a *natural fact* is a fact that is characteristically invoked by some natural science, such as physics, chemistry, biology, and psychology.¹⁰⁴

¹⁰² I argue for premise (1) in section 1.2.2, and I argue for premise (2) in section 1.2.3.

¹⁰³ In fact, the internalist argument for moral error theory is so prominent that some philosophers, such as Rowland (2013) and Ingram (forthcoming), define 'moral error theory' as the view that the internalist argument for moral error theory is sound.

¹⁰⁴ This is a common understanding of the natural/non-natural distinction. For example, see Moore (2005: 41), Enoch (2011: 103), Shafer-Landau (2003: 58), Smith (1994: 203).

Olson observes that normative favoring relations are oftentimes not irreducible. For example, suppose that I want to see a particular football game and that the local bar is showing that game this evening. The fact that the bar is showing the game favors my going to the bar. But it's not the case that this is an irreducibly normative favoring relation, for the fact that the bar's showing the game favors my going to the bar is reducible to the fact that the bar's showing the game explains why my going to the bar is a means to the satisfaction of one of my desires. In contrast, Olson argues, it seems plausible that moral facts 'are or entail facts that count in favour of or require certain courses of behaviour, where the favouring relation is irreducibly normative.'¹⁰⁵ For example, the fact that it is morally obligatory for John to donate to Oxfam entails that some fact—perhaps the fact that such a donation would help somebody avoid starvation—favors donating to Oxfam, where this favoring relation is not reducible to some natural fact such as a fact about desire promotion.

If moral facts really do entail irreducibly normative favoring relations, this gives rise to the following error-theoretic argument:

- (1) Moral facts entail that there are facts that favor certain courses of behavior, where the favoring relation is irreducibly normative.
- (2) Irreducibly normative favoring relations are queer.
- (3) So, moral facts entail queer relations.
- (4) If moral facts entail queer relations, moral facts are queer.
- (5) So, moral facts are queer.¹⁰⁶

The conclusion of this argument casts some doubt on the existence of moral facts, for by 'queer,' Olson means something along the lines of 'mysterious in such a way as to render something ontologically suspect.'¹⁰⁷ Olson defends premise (1) by offering cases like the Oxfam case

¹⁰⁵ Olson (2014: 118).

¹⁰⁶ *Ibid.*, 123-124.

¹⁰⁷ *Ibid.*, 87.

described above and by rebutting some objections to it.¹⁰⁸ Premise (4) is extremely intuitive; for if the instantiation of a kind of relation is ontologically suspect, it intuitively follows that anything that entails the instantiation of that kind of relation is similarly suspect. And Olson suggests that premise (2) is also quite intuitive, saying that upon reflection, irreducibly normative favoring relations ‘appear metaphysically mysterious.’¹⁰⁹

This argument from irreducible normativity suggests an analogous argument against the existence of practical epistemic facts, which I will call the *epistemic argument from irreducible normativity*:

- (1) Practical epistemic facts entail that there are facts that favor certain courses of behavior, where the favoring relation is irreducibly normative.
- (2) Irreducibly normative favoring relations are queer.
- (3) So, practical epistemic facts entail queer relations.
- (4) If practical epistemic facts entail queer relations, practical epistemic facts are queer.
- (5) So, practical epistemic facts are queer.

Premise (4) of this argument is clearly just as plausible as the fourth premise of Olson’s argument for moral error theory. And the second premises of these arguments are identical. So, we should believe that the epistemic argument from irreducible normativity is at least as plausible as Olson’s error-theoretic argument if it can be shown that the case for premise (1) of the former argument is at least as strong as the case for premise (1) of Olson’s argument. I will argue that this can indeed be shown.

I will say that for a fact F to stand in a *categorically normative favoring relation* to some behavior for S is for it to be the case that (i) F favors that behavior for S, and (ii) the fact that F favors that behavior does not depend on whether S’s engaging in it would serve her desires. In light of my arguments from Chapter 1, it is reasonable to believe that practical epistemic facts

¹⁰⁸ *Ibid.*, 126-135.

¹⁰⁹ *Ibid.*, 136.

entail that there are facts that favor certain courses of behavior, where the favoring relation is categorically normative. For recall that I defended the claim that epistemic duties of action are categorically applicable, in this sense: *whenever a person epistemically ought to do a certain action, it's the case that she epistemically ought to do that action regardless of whether doing it serves any of her desires.*¹¹⁰ I also defended the claim that epistemic duties are robustly normative, in this sense: *whenever a person epistemically ought to do a certain action, there's a reason for her to do that action.*¹¹¹ It follows from these italicized claims that whenever a person epistemically ought to do a certain action, it's the case that there is a reason for her to do that action regardless of whether doing it serves any of her desires. We should therefore believe that practical epistemic facts entail that there are facts that favor certain courses of behavior, where the favoring relation is categorically normative.

It follows that we should accept premise (1) of the epistemic argument from irreducible normativity so long as it's the case that categorically normative favoring relations are irreducibly normative favoring relations. It is commonly assumed that categorically normative favoring relations would have to be irreducible, and we have good reason to accept this view.¹¹² To see this, let a *normative fact* be a fact to the effect that some fact normatively favors some kind of behavior. For example, it is a normative fact that the fact that putting one's hand in the fire will cause one agony normatively favors (i.e., is a reason for) refraining from putting one's hand in the fire. Pre-theoretically, normative facts seem to involve a kind of force, pull or nudge. If the fact that putting one's hand in the fire will cause one agony genuinely favors refraining from acting in this way,

¹¹⁰ See section 1.2.2.

¹¹¹ See section 1.2.3.

¹¹² The irreducibility of categorically normative favoring relations is oftentimes just assumed without argument. For example, see Ingram (forthcoming: 1), Rowland (2015: 4) and Olson (2011: 64-65), the last of which goes so far as to suggest that 'what non-naturalist realists mean to capture in claiming that moral facts are non-natural is precisely that these facts are or entail categorical reasons.' In what follows, I defend a line of reasoning for the irreducibility of categorical normativity that is similar to arguments offered by Cowie (2016: 123) and Heathwood (2011: 85).

there's an intuitive sense in which one is being nudged, forced, or pulled away from putting one's hand in the fire. These uses of the terms 'nudged,' 'forced,' and 'pulled,' are just metaphors, but I assume that they are nonetheless helpful in referring to a certain feature that we believe normative facts to possess. To express the idea that normative facts have this feature, I will say that normative facts have *normative force*.

Since normative facts intuitively have normative force, we have reason to think that normative facts are reducible to a particular kind of natural fact only if natural facts of that kind involve something very closely analogous to normative force. At least most kinds of natural facts intuitively lack normative force and anything very closely analogous to it. For example, the natural facts that grass is green and that the ocean is deep seem not to nudge us towards certain behaviors in a way analogous to the way in which normative facts nudge us. We therefore have reason to believe that at least most kinds of natural facts are such that normative facts cannot be plausibly reduced to them.¹¹³ There is one kind of natural fact that seems to have something closely analogous to normative force, however: natural facts about desires. For a person to desire to behave in a certain way is for that person to be psychologically forced/pulled/nudged towards behaving in that way. This makes naturalistic reductions of normative facts to facts about desires much more plausible than other kinds of naturalistic reductions.

The foregoing considerations provide us with the following argument for the irreducibility of categorically normative favoring relations. Since natural facts about desires are unique among natural facts in possessing something closely analogous to normative force, we should believe that if normative facts are naturalistically reducible at all, then they're reducible to natural facts about desires. Normative facts that consist in the instantiation of categorically normative favoring

¹¹³ I think that Enoch (2011: 105) has this sort of challenge to naturalistic reductions in mind when he makes use of the 'just-too-different intuition.'

relations cannot be reduced to natural facts about desires; for by definition, categorically normative favoring relations obtain independently of desires. So, normative facts that consist in the instantiation of categorically normative favoring relations are not reducible to natural facts. So, categorically normative favoring relations are irreducibly normative favoring relations.

Recall that above, I have argued that if categorically normative favoring relations are irreducible, then we have good reason to believe that practical epistemic facts entail that there are facts that favor certain courses of behavior, where the favoring relation is irreducibly normative. I conclude that we have good reason to accept premise (1) of the epistemic argument from irreducible normativity. So, we should think that Olson's error-theoretic argument against moral facts succeeds only if a parallel argument against the existence of practical epistemic facts also succeeds.

2.2.3 The Supervenience Argument

Another source of moral skepticism revolves around considerations regarding moral supervenience. Mackie writes:

What is the connection between the natural fact that an action is a piece of deliberate cruelty – say, causing pain just for fun – and the moral fact that it is wrong? It cannot be an entailment, a logical or semantic necessity. Yet it is not merely that the two features occur together. The wrongness must somehow be 'consequential' or 'supervenient'; it is wrong because it is a piece of deliberate cruelty. But just what *in the world* is signified by this 'because'?¹¹⁴

Mackie thinks that whatever relation that obtains between moral facts/properties and natural facts/properties is bound to be queer and mysterious. He infers that there are no moral facts.

Meta-ethicists have identified several distinct challenges to the existence of moral facts that are generated by Mackie's remarks about supervenience. Consider two of them. First, the

¹¹⁴ Mackie (1977: 41).

moral skeptic might argue that the nature of the asymmetric dependence relation that must exist between moral facts and natural facts if the former are to exist is a highly mysterious relation. No plausible explication of the nature of this relation is available. This relation cannot be logical or semantic entailment, if Mackie is to be believed. Nor is it plausible that this relation is causation, for causes precede their effects in time, and it is implausible that a moral fact obtains a moment after the natural fact on which it depends obtains.¹¹⁵ Since the nature of the asymmetric dependence relation between moral facts and natural facts would be mysterious, and since we should reject the existence of such mysterious relations, we should reject the existence of moral facts.¹¹⁶

Second, the error theorist can argue that believers in moral facts face the difficult challenge of explaining a certain fact about the way in which moral and natural facts co-vary. Consider the following principle:

Moral Supervenience: Necessarily, there can be no moral difference between two items without there being some natural difference between them.

There is near-consensus among meta-ethicists today that Moral Supervenience is true, and this is for good reason. To see why, imagine a case in which somebody engages in drunk driving, and fill in the details so that it is obvious that what she did was morally wrong. It is highly intuitive that there could not be a second case of drunk driving that perfectly resembles this one in all natural respects (for instance, with respect to how drunk the driver is, how fast the driver drives, the driver's motivations for driving drunk, and how populated the area driven in is) but that is morally right. In other words, it is highly intuitive that Moral Supervenience holds true in this particular case of drunk driving. And notice that in coming to see that Moral Supervenience holds in this

¹¹⁵ This challenge to the idea that the relationship between moral and natural facts is a causal relationship is raised by Eric Wielenberg (2014), who nonetheless goes on to defend a causal account of this relationship.

¹¹⁶ This is the main error theoretic challenge that Wielenberg (2014) draws from the above passage of Mackie's.

particular case, nothing seems to depend on the special features of drunk driving. We can engage in similar exercises of imagination to justifiably conclude about many particular cases that Moral Supervenience holds true in them. And once we see that this is so and that it is hard to think of intuitive counter-examples to Moral Supervenience, Moral Supervenience becomes quite theoretically attractive. For the truth of this thesis would unify and explain all of the specific necessary connections between moral and natural facts discovered in the various particular cases in which Moral Supervenience was found to hold true.¹¹⁷

Moral Supervenience cries out for explanation. *Non-contingent moral error theory*, according to which it is metaphysically impossible for there to be moral facts, provides us with an easy explanation of it.¹¹⁸ For if it is impossible for there to be moral facts, then it is impossible for there to be moral differences between any two possible items. And from this it trivially follows that it is impossible for there to be a moral difference between two possible items without there being some natural difference. So, proponents of at least one version of moral error theory can easily explain Moral Supervenience.

In contrast, the task of explaining Moral Supervenience is very difficult for proponents of the existence of moral facts. To appreciate why this is so, we can consider the difficulties that arise for two of the most popular non-error-theoretic strategies for explaining Moral Supervenience. First, many try to explain it by adopting:

Moral Naturalism: Moral facts are identical with, or reducible to, natural facts.

If Moral Naturalism is true, then it is quite easy to explain Moral Supervenience. If moral facts are identical to natural facts or are reducible to them, then what it is for there to be a moral difference

¹¹⁷ This argument for Moral Supervenience is offered by McPherson (2012: 211; 2015).

¹¹⁸ This is noted by McPherson (2015).

between two items without a natural difference between them is for there to be a natural difference between two items without a natural difference between them. And it is impossible for there to be a natural difference without a natural difference. So, given Moral Naturalism, Moral Supervenience is not at all mysterious.

This explanation of supervenience is only as tenable as Moral Naturalism, and although Moral Naturalism is a popular thesis, it faces serious challenges. Recall that in the context of defending the irreducibility of categorical reasons, I observed that at least most natural facts intuitively lack normative force and anything closely analogous to normative force. That is, at least most natural facts intuitively fail to ‘nudge’ or ‘pull’ an agent towards some kind of behavior in a way that’s in play when a normative fact obtains. David Enoch goes further, suggesting that *all* natural facts intuitively lack normative force, and many share Enoch’s intuition. Suppose that this intuition is correct. Since all natural facts lack normative force, attempts to reduce moral facts to natural facts will be plausible only if moral facts lack normative force. But moral facts intuitively do have normative force. For example, when you morally ought to donate to charity, this is not just some stale, inert fact about the world. Rather, this moral fact seems to involve your being somehow nudged towards donating to charity. So, we have a reason to doubt that moral facts can be reduced to natural facts.¹¹⁹ It follows that we have a corresponding reason to doubt that Moral Supervenience can be adequately explained by an appeal to Moral Naturalism.

Second, some have tried to offer a moral explanation of Moral Supervenience by appealing to the fundamental moral principle(s). For example, suppose that a simple kind of utilitarianism is true, according to which an act is right if and only if, and because, that action is the one that

¹¹⁹ Enoch (2011: 104-105). Enoch actually presents this argument as an argument against a view that we can call *Normative Naturalism*, according to which *normative* facts are reducible to natural facts. But Enoch also asserts that the moral facts are a subset of the normative facts (2). So, I assume that Enoch wants to reject Moral Naturalism on the basis of the argument I present in the main text.

produces the greatest level of overall pleasure. If this is so, then Moral Supervenience is explained; for then, moral properties would have to supervene on facts about pleasure maximization, and such facts are natural facts. This explanation is an especially popular strategy among those who reject Moral Naturalism, and who consequently cannot avail themselves of the naturalistic explanation considered above.¹²⁰

I agree that appeals to the relevant norms are sometimes sufficient to shed explanatory light on supervenience facts. For instance, Enoch is correct that one can dissolve the mystery that might surround the supervenience of facts about who can legally drink on facts about people's ages by appealing to the legal norms that govern the jurisdiction at hand. But there is a disanalogy between this legal case and the case of Moral Supervenience that makes the moral analogue of this explanation much less illuminating. In the legal case, the fact that the relevant legal norm obtains is obviously itself explicable. We have a ready grasp on why this norm obtains: it has something to do with the structure of our government and how certain politicians voted on some bill. In contrast, when a moral realist appeals to a fundamental moral principle to explain Moral Supervenience, it is not obvious why this fundamental moral principle is true. This is why it is so natural to ask the moral realist the follow-up question, "Alright, but what explains *that*? What explains the fundamental moral principle itself?"

Essentially, the worry for the strategy of appealing to a moral norm to explain Moral Supervenience is that it succeeds in dispelling the mystery of Moral Supervenience only if it is coupled with some explanation of why the moral norm itself obtains, and this latter sort of explanation is itself quite difficult to provide. Perhaps such an explanation can be provided, in which case this strategy may be vindicated. My modest point here is that this explanatory task

¹²⁰ Enoch (2011: 143), Kramer (2013), Olson (2014: 96), and Scanlon (2014: 40), for example.

seems just as difficult as the original explanatory task of explaining Moral Supervenience, so appealing to the moral norm to explain Moral Supervenience sheds little light on why the latter is true.¹²¹

So, some versions of moral error theory easily explain Moral Supervenience whereas it's difficult for believers in moral facts to explain this thesis. Moral skeptics claim that these considerations constitute a second reason to reject the existence of moral facts.

Obviously, in presenting the above two supervenience worries for moral facts, I have only touched on a number of deep and difficult issues, such as (i) our grounds for accepting Moral Supervenience, (ii) moral realists' resources for explaining Moral Supervenience, and (iii) the tenability of Moral Naturalism. But enough has been said to make a presumptive case for thinking that these two supervenience worries for moral facts suggest equally-plausible arguments against the existence of practical epistemic facts.

First, just as moral facts obtain because certain natural facts obtain, practical epistemic facts also seem to obtain because certain natural facts obtain. For example, suppose that it is a fact that you epistemically ought to enroll in a particular professor's critical thinking class. This is surely so because of some natural facts that obtain, such as (i) the fact that one is presently disposed to commit various fallacies, (ii) the fact that this professor is especially talented at helping her students learn to avoid committing these fallacies, (iii) the fact that all of the other critical thinking courses are full, and so on. Of course, any general view about precisely which natural facts are

¹²¹ Enoch seems to respond to this worry by suggesting that not all facts cry out for explanation and that the fundamental moral principle itself (or perhaps some facts about the nature of moral properties) is one such fact (Enoch 2011: 147-48). I agree that if the fundamental moral principle is brutally true and does not cry out for explanation, then one can satisfactorily explain Moral Supervenience by appealing to the fundamental moral principle without providing a supplementary explanation of the fundamental moral principle. But my intuitions (and many others' intuitions, I assume) differ from the ones that Enoch seems to have. The true fundamental moral principle, whatever it is, does seem to cry out for explanation. A reason to think that many people share my intuitions on this point is that many people are inclined to seek out an explanation of the fundamental moral principle.

such that practical epistemic facts depend on them is bound to be controversial. Analogously, any view about precisely which natural facts are such that moral facts obtain in virtue of those natural facts obtaining is bound to be controversial. But regardless of which natural facts are such that practical epistemic facts obtain because of them, we can ask the epistemic analogue of Mackie's famous question: just what *in the world* is signified by the 'because'?

So the same sort of asymmetric dependence relation that obtains between moral and natural facts also obtains between practical epistemic facts and natural facts. If the mysteriousness of the nature of this relation shows that there are no moral facts, then its mysteriousness also shows that there are no practical epistemic facts. So, if the first supervenience challenge to moral facts is sound, then the analogous argument against practical epistemic facts is also sound.

Now consider the second supervenience challenge. Just as a skeptic about moral facts can appeal to a certain fact about the way in which moral and natural facts co-vary, a skeptic about practical epistemic facts can appeal to a certain way in which practical epistemic facts and natural facts co-vary. Consider the following principle:

Practical Epistemic Supervenience: Necessarily, there can be no epistemic differences between two actions without there being some natural difference between them.

Admittedly, Practical Epistemic Supervenience does not enjoy the level of consensus that Moral Supervenience enjoys, but I think that this is just because philosophers have failed to give practical epistemic facts their due attention. For Practical Epistemic Supervenience can be supported with the same two-step argument offered above for Moral Supervenience. First off, imagine a case in which somebody ignores constructive criticism of her ideas, and fill in the details so that it is highly intuitive that she did something epistemically wrong. For example, we can imagine the case of Sam that was discussed in the previous chapter:

Sam receives some critical feedback on a project in progress from his colleague Margaret. Margaret is astute and perceptive; however, she does tend to voice her views with verve, and the force of her comments is in keeping with this trait. Upon reading Margaret's comments, Sam finds himself believing that she doesn't really appreciate his work. He thereupon finds himself annoyed and ultimately angry with her...In a fit of resentment, he forms the intention to ignore the generous criticism that Margaret has offered him and subsequently acts upon it.¹²²

Sam's choice to ignore Margaret's criticisms is epistemically wrong: he epistemically ought not act in that way. And it is intuitive that there could not be a second case of the ignoring of constructive criticism that perfectly resembles this one in all natural respects (for instance, with respect to how astute and perceptive the person offering the criticism is, how aware the ignorer of the criticism is of the criticizer's astuteness, how the criticizer delivers her criticism, and so on), except that the second action is epistemically right. In other words, it is highly intuitive that Practical Epistemic Supervenience holds true in this particular case of the ignoring of criticism.

And notice that in coming to see that Practical Epistemic Supervenience holds in this particular case, nothing seems to depend on the special features of Sam or of the ignoring of criticism. We can engage in similar exercises of imagination to justifiably conclude about many particular cases that Practical Epistemic Supervenience holds true in them. And once we see that this is so, and that it is hard to think of intuitive counter-examples to Practical Epistemic Supervenience, Practical Epistemic Supervenience becomes quite attractive. For the truth of this thesis would unify and explain all of the specific necessary connections between moral and epistemic facts discovered in the various particular cases in which Practical Epistemic Supervenience was found to hold true.

Like Moral Supervenience, Practical Epistemic Supervenience cries out for explanation. *Non-contingent practical epistemic error theory*, according to which it is metaphysically

¹²² Cuneo (2007: 73).

impossible for actions to have epistemic properties, provides us with an easy explanation of it. For if it is impossible for actions to have epistemic properties, then it is impossible for two actions to differ in their epistemic properties. And from this it trivially follows that it is impossible for two actions to differ in their epistemic properties without there being some natural difference between them. So, proponents of at least one version of practical epistemic error theory can easily explain Practical Epistemic Supervenience.

In contrast, the task of explaining supervenience is difficult for proponents of the existence of practical epistemic facts. To appreciate why this is so, we can consider the difficulties that arise for the epistemic analogues of the two popular strategies for explaining Moral Supervenience discussed above. First, one might try to explain Practical Epistemic Supervenience by adopting:

Practical Epistemic Naturalism: Practical epistemic facts are identical with, or reducible to, natural facts.

If Practical Epistemic Naturalism is true, then it is quite easy to explain Practical Epistemic Supervenience. If practical epistemic facts are identical to natural facts or are reducible to them, then what it is for there to be an epistemic difference between two actions without a natural difference between them is for there to be a natural difference between two actions without a natural difference between them. And it is impossible for there to be a natural difference without a natural difference. So, given Practical Epistemic Naturalism, Practical Epistemic Supervenience is not at all mysterious.

The major problem with this strategy for explaining supervenience is that Practical Epistemic Naturalism is questionable, as it shares in the difficulties that afflict its moral analogue. Intuitively, practical epistemic facts have normative force. For example, when you epistemically ought to pay attention to someone's constructive criticism, this isn't just some stale, inert fact about

the world. Rather, the obtaining of this fact seems to involve your being somehow ‘nudged’ or ‘pulled’ towards paying attention to the person’s criticism. Since this is so, we have good reason to think that practical epistemic facts are reducible or identical to some kind of natural fact only if that kind of natural fact has normative force. But as I noted above, it seems that natural facts lack normative force. So we have reason to think that practical epistemic facts are neither identical nor reducible to natural facts. It follows that we have a corresponding reason to doubt that Practical Epistemic Supervenience can be adequately explained by an appeal to Practical Epistemic Naturalism.

Second, one might try to explain Practical Epistemic Supervenience by appealing to fundamental epistemic principle(s). For example, suppose that a simple kind of epistemic consequentialism is true, according to which an act is epistemically obligatory if and only if, and because, that action is the one that produces the greatest amount of true beliefs for the one performing it. If this is so, then Practical Epistemic Supervenience is explained; for then, the epistemic properties of actions would have to supervene on properties having to do with the production of true beliefs, and such properties are natural properties.

The worry for this strategy for explaining Practical Epistemic Supervenience parallels the worry for its moral analogue. Appealing to fundamental epistemic norms to explain Practical Epistemic Supervenience succeeds in dispelling the mystery of Practical Epistemic Supervenience only if it is coupled with some explanation of why the fundamental epistemic norms themselves obtain, and this latter sort of explanation seems difficult to provide. As in the moral case, it may be that such an explanation can be provided, in which case this explanatory strategy may be vindicated. But this explanatory task seems just as difficult as the original explanatory task of

explaining Practical Epistemic Supervenience, so appealing to the epistemic norms to explain this supervenience thesis seems to shed little light on why the latter is true.

So, some versions of practical epistemic error theory easily explain Practical Epistemic Supervenience whereas it is quite difficult for believers in practical epistemic facts to explain this thesis. I conclude that the explanatory disadvantage faced by believers in practical epistemic facts is extremely similar to the explanatory disadvantage faced by believers in moral facts. Since this is so, it's reasonable to think that if the latter kind of explanatory disadvantage should seriously undermine our confidence in moral facts, then the former kind of explanatory disadvantage should seriously undermine our confidence in practical epistemic facts.

2.2.4 The Epistemological Queerness Argument

Mackie's argument from queerness has an epistemological component. He writes:

If we were aware of [moral values], it would have to be by some special faculty of moral perception or intuition, utterly different from anything else in the universe... When we ask the awkward question, how we can be aware of this authoritative prescriptivity, of the truth of these distinctively ethical premises or of the cogency of this distinctively ethical pattern of reasoning, none of our ordinary accounts of sensory perception or introspection or the framing and confirming of explanatory hypotheses or inference or logical construction or conceptual analysis, or any combination of these, will provide a satisfactory answer; 'a special sort of intuition' is a lame answer, but it is the one to which the clear-headed objectivist is compelled to resort.¹²³

In other words, in order to know that some moral proposition is true, we would need to have some kind of special cognitive faculty with which to detect them—a cognitive faculty that is 'queer' in being 'utterly different from anything else in the universe.' We ought not to believe in such queer cognitive faculties, so we ought to reject the claim that we sometimes have moral knowledge. This conclusion is not quite the denial of moral facts' existence, for it is conceivable that there are moral

¹²³ Mackie (1977).

facts, none of which are ever known. But the conclusion that we have no moral knowledge would make the denial of moral facts far more reasonable. For the main reason we have for believing that there are moral facts is that we are confident in the truth of some particular, first-order moral claims, such as the claim that it is wrong to murder children. The conclusion that we have no moral knowledge undermines these grounds for believing in moral facts.

A crucial premise of this error-theoretic argument is that ‘none of our ordinary accounts of sensory perception or introspection or the framing and confirming of explanatory hypotheses or inference or logical construction or conceptual analysis, or any combination of these’ can provide a satisfactory explanation of our ability to detect moral truths. The support for this premise is that moral facts are themselves ‘queer’ in some way that renders it implausible that they are detectable via the normal cognitive faculties.¹²⁴ And the relevant queerness of moral facts is supposed to derive from the kind of normativity that they display.

I can think of two ways that the moral skeptic might try to spell out the relevant kind of queerness. First, she might have in mind the notion that moral facts have *practical clout*.¹²⁵ For moral facts to have practical clout is the following to be true: (i) Moral duties are categorically applicable, in that when a person is morally obligated to perform a certain action, she is morally obligated to do that action independently of whether doing so serves her desires; and (ii) Moral duties are robustly normative, in that when a person is morally obligated to perform a certain action, there is a genuine reason for her to do that action. Second, the moral skeptic might urge that moral facts are not reducible or identical to natural facts, and for this reason cannot be detected

¹²⁴ Thus, Mackie’s epistemological queerness argument seems to depend a version of the metaphysical queerness argument, as Joyce (2015) notes.

¹²⁵ This is Joyce (2006: 62)’s term.

unless we have queer cognitive faculties. These two sources of metaphysical queerness suggest two distinct versions of the epistemological queerness argument against moral facts.

The epistemological queerness argument has an equally-plausible epistemic analogue. In the previous chapter, I argued at length for the conclusion that practical epistemic facts have practical clout, or in other words, that epistemic duties of action are both categorically applicable and robustly normative.¹²⁶ And earlier in this section, I have argued that we have good reason to deny that practical epistemic facts are reducible or identical to natural facts—reasons that parallel our reasons for denying that moral facts are reducible or identical to natural facts. So we have good reason to believe that if the epistemological queerness argument succeeds in establishing that moral facts don't exist, then a closely analogous argument succeeds in showing that practical epistemic facts don't exist.

I conclude that four of the most important arguments against moral facts succeed only if closely-analogous arguments against practical epistemic facts succeed. We therefore have good reason to accept the Parity Premise: If moral facts do not exist, then practical epistemic facts do not exist. But since, as I have sought to establish in Chapter 1, practical epistemic facts do indeed exist, moral facts do, too.

2.3 Heathwood's Objection from Epistemic Naturalism

Some critics have objected to companions in guilt arguments using the following strategy: first, argue that moral facts are more likely than their epistemic companions to be irreducible; and second, argue that this makes it the case that there are better grounds for disbelieving in moral facts than there are for disbelieving in the epistemic facts.¹²⁷ If these two steps are successfully carried out, then the Parity Premise of the companions in guilt argument is undermined. When this

¹²⁶ Again, see sections 1.2.2 and 1.2.3.

¹²⁷ Heathwood (2009) and Cowie (2016) use strategies along these lines.

general strategy is tailored to target the argument from practical epistemic facts that I have been defending, we get the following *objection from epistemic naturalism*:

- (1) Moral facts are more likely than practical epistemic facts to be irreducible.
- (2) If (1) is true, then the Parity Premise is false.
- (3) So, the Parity Premise is false.

Premise (2) is true. For if (1) is true, then the second supervenience argument against moral facts does not suggest an equally-worrisome argument against practical epistemic facts. This is because if (1) is true, then it is more likely that Practical Epistemic Supervenience can be explained by appealing to Practical Epistemic Naturalism than it is that Moral Supervenience can be explained by appealing to Moral Naturalism. Furthermore, if (1) is true, then the epistemological queerness argument against moral facts does not suggest an equally-worrisome argument against practical epistemic facts. This is because if (1) is true, then there are better grounds for believing that the recognition of moral facts would require queer cognitive faculties than there are for believing that the recognition of practical epistemic facts would require queer cognitive faculties. In light of these considerations, we should accept premise (2) of the objection from epistemic naturalism is true.

I must therefore argue against premise (1). Before doing so, it is important to note a complication of the present dialectic. As I mentioned in section 2.2, companions in guilt arguments typically appeal solely to doxastic epistemic facts. As a result, all extant versions of the objection from epistemic naturalism are developed with the specific aim of refuting this typical formulation of the companions in guilt argument. This means that critics of the companions in guilt argument have not argued directly in favor of premise (1), but rather have argued for the analogous premise

that *moral facts are more likely than doxastic epistemic facts to be irreducible*.¹²⁸ In what follows, I first identify the best defense of this analogous premise that has been offered by critics of companions in guilt arguments, which I take to be the defense that has recently been developed by Christopher Heathwood. I then argue that nothing like Heathwood's defense can be constructed to establish premise (1) of the above objection from epistemic naturalism. This will show that the best support for premise (1) suggested by the literature fails, and that consequently, the objection from epistemic naturalism is not (as yet) compelling.

Heathwood's case for the claim that moral facts are more likely than doxastic epistemic facts to be irreducible begins with the endorsement of the following version of the open-question argument. Certain claims seem to us to be self-contradictory (for example, 'He is a brother, but he is not a sibling'), while others seem to us not to be self-contradictory (for example, 'He is a brother, but he is not a father'). When a claim of the form *it is F, but it is not G* strikes us as not being self-contradictory, this is some evidence that it is in fact not self-contradictory, and that 'G' therefore does not mean the same as 'F.' This test for synonymy casts doubt on attempts to analyze moral terms in purely natural terms. For example, consider the proposal that 'morally good' means 'conducive to happiness.' 'The act is conducive to happiness, but it is not morally good' seems not to be self-contradictory, so we have reason to believe that 'morally good' does not mean 'conducive to happiness.' By running through a number of naturalistic analyses of moral terms, it becomes reasonable for one to doubt all naturalistic analyses. We therefore have grounds for

¹²⁸ The argument from practical epistemic facts therefore sidesteps some of the criticism that has been developed by opponents of companions in guilt arguments for moral facts. This is one of the dialectical advantages of appealing to practical epistemic facts rather than doxastic epistemic facts in this context.

believing that no moral term is synonymous with a natural term. And therefore, we have grounds for believing that moral facts are not reducible to natural facts.¹²⁹

The next step of Heathwood's case is to argue that the epistemic analogue of the open-question argument fails to support the irreducibility of doxastic epistemic facts. He asks us to consider the following proposed definition of the epistemic term 'is reasonable for S to believe':

'*p* is reasonable for S to believe' means '*p* is likely, given S's information.'

Applying the above self-contradictory test, we should ask whether the following seems self-contradictory: 'this is likely, given my evidence, but it is not reasonable for me to believe it.' Heathwood invites us to agree that this sentence 'does have an air of incoherence about it in a way that axiological statements—even such patently false ones like 'suffering is intrinsically good—never do.'¹³⁰ Since doxastic epistemic terms are plausibly definable in terms of probabilities and evidence, epistemic open-question arguments against naturalistic reductions of epistemic facts to facts about probabilities and evidence fail.

Together, the success of the moral open-question argument and the failure of the epistemic open-question argument provide us with grounds for believing that moral facts are more likely than doxastic epistemic facts to be irreducible. I take this to be the strongest argument for believing that moral facts and doxastic epistemic facts differ with regards to how likely they are to be irreducible, and I will grant that this argument succeeds.

¹²⁹ Of course, this last step from the non-synonymy of moral and natural terms to the non-reducibility of moral facts to natural facts is quite controversial. This step fails if the semantics of moral terms resemble the semantics of natural kind terms such as 'water,' as Heathwood (2009: 88) acknowledges. I will grant for the sake of the discussion that the non-synonymy of terms implies non-reducibility of facts. In this context, this is a concession that is favorable to my anti-realist interlocutor; for if the non-synonymy of terms does not imply the non-reducibility of the corresponding facts, then the moral open-question argument fails. And if the moral open-question argument fails, then (as is apparent from what follows) Heathwood's objection to the companions in guilt argument fails.

¹³⁰ *Ibid.*, 90.

The important question for our purposes is whether the argument for premise (1) that is naturally suggested by Heathwood's appeal to open-question arguments is also successful. This argument for premise (1) would go as follows. The moral open-question argument succeeds, whereas the open-question argument for the irreducibility of *practical* epistemic facts fails. So, there is strong reason to accept the irreducibility of moral facts for which there is no analogous reason to accept the irreducibility of practical epistemic facts. Therefore, premise (1) of the objection from epistemic naturalism is true.

But this argument for premise (1) fails because the epistemic open-question argument for the irreducibility of practical epistemic facts is at least as plausible as the moral open-question argument. Recall that the epistemic open-question argument for the irreducibility of *doxastic* epistemic facts fails because there are natural terms about probabilities and evidence that are plausibly synonymous with such epistemic terms as 'reasonable belief.' So, the epistemic open-question argument for the irreducibility of practical epistemic facts can be shown to fail for Heathwood-style reasons only if there are natural terms that are plausibly synonymous with such epistemic terms as 'is an action that one epistemically ought to do.' But what sort of natural terms might fit this description? It is quite difficult to see how such epistemic terms could possibly be definable in terms of probability and evidence. There does not seem to be any proposition p that is such that 'this act is epistemically obligatory' is synonymous with ' p is likely, given the evidence.'

A natural thought is that something like 'conducive to the formation of true belief' is a natural term that plays the needed role. This seems like the best candidate for a natural term that's synonymous with practical epistemic terms, for two reasons. First, it would cohere with a lot of our intuitions about the content of our epistemic duties of action. For example, this view about the

meaning of practical epistemic terms has the plausible implication that we epistemically ought to engage in such actions as paying attention to constructive criticism, seeking out experts' opinions, and gathering evidence in certain ways rather than others. Second, this candidate respects the intuitive thought that something's having positive epistemic status indicates that it is somehow conducive to gaining a more accurate grasp of reality. So let us consider the following analysis:

'φ is an act that S epistemically ought to do' means 'φ-ing is conducive to S's formation of true beliefs.'

If this analysis is tenable, then we could conclude that the epistemic open-question argument for the irreducibility of practical epistemic facts is refuted by the sort of objection that Heathwood develops for the case of doxastic epistemic facts. But let's apply Heathwood's self-contradictory test to the above analysis. Consider whether the following statement seems self-contradictory: 'Doing this action is conducive to my forming true beliefs, but it is not something that I epistemically ought to do.' While this statement may be false, it seems no more self-contradictory than the statement, 'the act produces the most amount of happiness, but it is not morally good.'

I conclude that Heathwood's critique of the epistemic open-question argument for the irreducibility of practical epistemic facts is not compelling. This means that the strongest support for premise (1) that is suggested by the literature fails. We should therefore find the objection from epistemic naturalism unpersuasive.

2.4 Ingram's Objection from Epistemic Non-Cognitivism

Stephen Ingram has recently offered a novel objection to companions in guilt arguments that appeals to epistemic non-cognitivism, the epistemic analogue of moral non-cognitivism.¹³¹ According to epistemic non-cognitivism, epistemic judgments are conative, non-representational

¹³¹ Ingram (forthcoming).

attitudes rather than descriptive beliefs. For example, a simple version of epistemic non-cognitivism holds that to judge that a belief is epistemically justified is to approve of the formation of that belief, and that to judge a belief to be irrational is to disapprove of the formation of that belief.¹³²

Companions in guilt arguments are threatened by epistemic non-cognitivism. For if epistemic non-cognitivism is true, then epistemic judgments are not truth apt. If epistemic judgments are not truth apt, then either there are no epistemic facts, or there are, but none are representable by epistemic judgments. Because this latter disjunct is so implausible, we should believe that if epistemic judgments are not truth apt, then there are no epistemic facts. So, if epistemic non-cognitivism is true, then the Ontological Premise of the realist's companions in guilt argument is false. When this type of objection is tailored specifically to target my argument from practical epistemic facts, this results in the following *objection from epistemic non-cognitivism*:

- (1) Non-cognitivism about practical epistemic judgments is true.
- (2) If (1), then there are no practical epistemic facts.
- (3) Therefore, there are no practical epistemic facts.¹³³

A *practical epistemic judgment* is a judgment to the effect that some action is one that someone epistemically ought to do. Premise (2) is true for reasons just explained, so I will argue against (1). But first, I should again remark on the dialectical complication that arises from the fact that the companions in guilt argument that I am defending is non-standard. Because companions in guilt

¹³² See Chrisman (2012) for a helpful overview of epistemic non-cognitivism, and see Cuneo (2007) for extended, critical discussion of its merits.

¹³³ This is not quite the challenge to companions in guilt arguments that Ingram develops. He actually uses epistemic non-cognitivism to attack the alleged parity between moral and epistemic error theory, whereas I am interpreting the challenge posed by epistemic non-cognitivism as a challenge to the existence of the relevant epistemic facts. In doing so, I am following Cuneo (2007). This departure from the letter of Ingram's objection is inconsequential, for if my criticism of the objection from epistemic non-cognitivism succeeds, then it will be apparent that Ingram (forthcoming)'s objection to the moral-epistemic parity succeeds.

arguments normally appeal to doxastic epistemic facts, Ingram does not develop an argument for premise (1), but rather argues for the analogous premise that non-cognitivism about *doxastic epistemic judgments*—that is, judgments to the effect that some doxastic attitude has some epistemic property—is true. My version of the companions in guilt argument therefore neatly evades the specific objection that Ingram develops. But it may be that a close analogue of Ingram’s objection can be developed that does threaten my argument. In what follows, I first explain Ingram’s argument for the claim that non-cognitivism about doxastic epistemic judgments is true. I then argue that nothing like his argument for this claim will succeed in obliging us to accept premise (1) of the above objection from epistemic non-cognitivism. This will show that Ingram’s basic challenge to companions in guilt arguments does not undermine my argument from practical epistemic facts, even if it does undermine the standard formulation of the companions in guilt argument.

Ingram’s argument for non-cognitivism about doxastic epistemic judgments begins with a thought experiment.¹³⁴ Suppose that everyone becomes convinced of error theory about doxastic epistemic matters, and that consequently, they abandon doxastic epistemic discourse. People stop saying that *her belief is irrational* and *he ought to suspend judgment*, for they have concluded that all such doxastic epistemic claims are false. Nonetheless, people naturally continue to have various conative attitudes towards many possible and actual doxastic attitudes. For example, most people will still approve of beliefs that are based on evidence and disapprove of beliefs that are formed entirely as a result of wishful thinking. Imagine that everyone adopts the old doxastic epistemic language to express and communicate the conative attitudes they have towards doxastic attitudes. For example, I express my approval of your evidence-based belief by saying, ‘you are believing

¹³⁴ Ingram’s argument is modeled on an argument for moral non-cognitivism that he extracts from Blackburn (1993).

as you ought.’ And you express your disapproval of beliefs based on wishful thinking by saying, ‘that belief is irrational.’ Essentially, we are being asked to imagine that our doxastic epistemic discourse turns into the sort of discourse that non-cognitivists take it to be already. We can call the sorts of mental states that people express with doxastic epistemic terms in this new setting ‘doxastic shmepistemic judgments.’ And we can call debate about which doxastic shmepistemic judgments to adopt ‘doxastic shmepistemic debate.’

Ingram’s argument goes as follows. Phenomenologically, making doxastic shmepistemic judgments looks and feels just like making doxastic epistemic judgments; and engaging in doxastic shmepistemic debate looks and feels just like engaging in doxastic epistemic debate. Generally, we have reason to believe that things are as they appear. So, we have reason to believe that what it is to make a doxastic epistemic judgment is to make a doxastic shmepistemic judgment, and we have reason to believe that what it is to engage in doxastic epistemic debate is to engage in doxastic shmepistemic debate. In other words, we have reason to believe that the actual world, in which we make doxastic epistemic judgments and engage in doxastic epistemic debate, is really a world in which non-cognitivism about doxastic epistemic judgments is true.

Consider the *similarity premise* in this line of reasoning – that is, the premise that doxastic shmepistemic judgments and doxastic shmepistemic debate look and feel just like the doxastic epistemic judgments and debate that we observe in the actual world. This premise must be defended from some natural worries. First off, when we are deliberating about whether to form a given doxastic epistemic judgment, we experience such deliberation as being aimed at making a correct and non-arbitrary decision. For example, when one is deliberating about whether to make the doxastic epistemic judgment that it is reasonable to doubt the defendant’s guilt, it seems to one that there is a correct decision that one’s deliberation is aimed at discerning. In contrast, it may

seem unlikely that one who is deliberating about which doxastic shmepestemic judgment to make will experience her deliberation as being aimed at making a correct and non-arbitrary decision. After all, since doxastic shmepestemic judgments are conative attitudes, it's not the case that a correct decision regarding which doxastic shmepestemic judgments to adopt can be determined by which of them are true.

Second, it seems that when two parties are engaged in doxastic epistemic debate about whether to adopt a given doxastic epistemic judgment, there is some non-arbitrary 'standard for success and failure' by which we can judge whether the debate is resolved successfully.¹³⁵ For example, if you and I are debating about what we ought to believe about climate change, we experience this debate as being such that not just any convergence in epistemic judgment that might result will count as a successful resolution. In contrast, it may seem unlikely that there is a non-arbitrary standard by which we can judge whether doxastic shmepestemic debate is resolved successfully. After all, since interpersonal disagreement in doxastic shmepestemic judgment consists of differing conative attitudes, it cannot be the case that an objectively correct resolution of the debate consists of convergence on the true doxastic shmepestemic judgments. So, initial reflection suggests that doxastic shmepestemic judgments and debate looks and feels quite different from doxastic epistemic judgments and debate.

Ingram defends his similarity premise from these natural worries by appealing to the popular view that belief 'aims' at truth. On his favored interpretation of the truth-directedness of belief, it consists of the fact that part of what it is for an activity to count as inquiry into how to evaluate beliefs is for that activity to be 'guided by the truth norm.'¹³⁶ Or in other words:

¹³⁵ Ingram (forthcoming: 15).

¹³⁶ *Ibid.*, 15.

Belief-Truth Principle: Part of what it is for an activity to count as inquiry into how to evaluate beliefs is for that activity to be such that it counts as correct or successful to the extent that it results in (i) positively evaluating true beliefs, and (ii) negatively evaluating false beliefs.

Inquiry into which doxastic shmepestemic judgments to make is a kind of inquiry into how to evaluate beliefs.¹³⁷ So, it follows from the Belief-Truth Principle that inquiry into which doxastic shmepestemic judgments to make is correct to the extent that it results in (i) positively evaluating true belief (for example, by forming the doxastic shmepestemic judgment *that one ought to believe p* when p is a true proposition), or (ii) negatively evaluating false beliefs (for example, by forming the doxastic shmepestemic judgment *that one ought not to believe that p* when p is a false proposition).

The foregoing generates the following response to the two natural worries for the similarity premise presented above. Contrary to the first worry, we can account for why one who is deliberating about which doxastic shmepestemic judgment to make will experience her deliberation as being aimed at making a correct and non-arbitrary decision. Given the Belief-Truth Principle, there *is* a correct and non-arbitrary decision for such deliberation to aim at. The correct decision will be whichever decision results in positively evaluating true beliefs or negatively evaluating false beliefs. And contrary to the second worry, there is a non-arbitrary standard by which we can judge whether doxastic shmepestemic debate is resolved successfully. This is simply the truth standard: a doxastic shmepestemic debate is resolved successfully if and only if it results in convergence upon the truth.

I have now offered enough details of Ingram's argument for my purposes. Recall that I want to argue that Ingram's case for non-cognitivism about doxastic epistemic judgments fails to

¹³⁷ I am using the term 'evaluate' quite broadly, such that one can count as having evaluated a belief B both by forming an evaluative belief about B and by forming a conative attitude towards B.

suggest a plausible case for premise (1) of the objection from epistemic non-cognitivism. So now consider how an Ingram-style argument for premise (1) would go. It starts with the following thought experiment. Suppose that everyone becomes convinced of error theory about practical epistemic matters, and that consequently, they abandon practical epistemic discourse. Nonetheless, people naturally continue to have various conative attitudes towards actions about which they used to form practical epistemic judgments. For example, most people will still disapprove of such actions as ignoring constructive criticism and surrounding oneself entirely by people who share one's beliefs. Imagine that everyone adopts the old practical epistemic language to express and communicate the conative attitudes they have towards these sorts of actions. For example, I express my disapproval of your proclivity to refuse to listen to criticism by saying, 'you are acting as you ought not to, from the epistemic point of view.' Essentially, we are being asked to imagine that our practical epistemic discourse turns into the sort of discourse that non-cognitivists take it to be already. We can call the sorts of mental states that people express with practical epistemic terms in this new setting 'practical shmepistemic judgments.' And we can call debate about which practical shmepistemic judgments to adopt 'practical shmepistemic debate.'

Phenomenologically, making practical shmepistemic judgments looks and feels just like making practical epistemic judgments; and engaging in practical shmepistemic debate looks and feels just like engaging in practical epistemic debate. Generally, we have reason to believe that things are as they appear. So, we have reason to believe that what it is to make a practical epistemic judgment is to make a practical shmepistemic judgment, and we have reason to believe that what it is to engage in practical epistemic debate is to engage in practical shmepistemic debate. In other words, we have reason to believe that the actual world, in which we make practical epistemic

judgments and engage in practical epistemic debate, is really a world in which non-cognitivism about practical epistemic judgments is true.

The problem with this argument is that its similarity premise – the premise that practical shmeepistemic judgments and practical shmeepistemic debate look and feel just like practical epistemic judgments and debate that we observe in the actual world – is not plausible. The reasons to find it implausible mirror the reasons to find the similarity premise of Ingram’s original argument to be *prima facie* implausible. First off, when we are deliberating about whether to form a given practical epistemic judgment, we experience such deliberation as being aimed at making a correct and non-arbitrary decision. For example, when you are deliberating about whether to make the practical epistemic judgment that you epistemically ought to restrict your reading to authors who agree with you, it seems to you that there is a correct decision that your deliberation is aimed at discerning. In contrast, it seems unlikely that one who is deliberating about which practical shmeepistemic judgment to make will experience her deliberation as being aimed at making a correct and non-arbitrary decision. After all, it’s not the case that a correct decision regarding which practical shmeepistemic judgments to adopt can be determined by which of them are true, and everybody is aware that this is so in the thought experiment.

Second, it seems that when two parties are engaged in practical epistemic debate about whether to adopt a given practical epistemic judgment, there is some non-arbitrary standard for success and failure by which we can judge whether the debate is resolved successfully. For instance, if you and I are debating about whether we epistemically ought to listen carefully as the climate change denier makes her case, we experience this debate as being such that not just any convergence in epistemic judgment that we might achieve will count as a successful resolution. In contrast, it seems unlikely that there is a non-arbitrary standard by which we can judge whether

practical shmepistemic debate is resolved successfully. After all, since interpersonal disagreement in practical shmepistemic judgment consists of differing conative attitudes, it cannot be the case that an objectively correct resolution of the debate consists of convergence on the true practical shmepistemic judgments. So upon reflection, practical shmepistemic judgment and debate looks and feels quite different from practical epistemic judgment and debate.

Ingram addressed these types of worries in the context of his original argument by appealing to the Belief-Truth Principle. But no analogous move is plausible in the context of defending the above Ingram-style argument for premise (1). For the analogous move would have to proceed as follows. Action ‘aims’ at truth, such that some principle along the following lines is true:

Action-Truth Principle: Part of what it is for an activity to count as inquiry into how to evaluate actions is for that activity to be such that it counts as correct or successful to the extent that it results in (i) positively evaluating actions that promote the formation of true beliefs, and (ii) negatively evaluating actions that promote the formation of false beliefs.

Inquiry into which practical shmepistemic judgments to make is a kind of inquiry into how to evaluate actions. So, it follows from the Action-Truth Principle that inquiry into which practical shmepistemic judgments to make is correct to the extent that it results in (i) positively evaluating actions that promote the formation of true beliefs (for example, by resulting in the formation of the practical shmepistemic judgment *that one ought to enroll in a critical thinking course*), or (ii) negatively evaluating actions that promote the formation of false beliefs (for example, by resulting in the formation of the practical shmepistemic judgment *that one ought not to ignore sources that disagree with what one already believes*). And this means that: (i) we *can* account for why one who is deliberating about which practical shmepistemic judgment to make will experience her deliberation as being aimed at making a correct and non-arbitrary decision, and (ii) there *is* a non-

arbitrary standard by which we can judge whether practical shmepestemic judgment is resolved successfully.

But this line of reasoning faces problems. First off, the idea that there is a constitutive aim of action is far more controversial than the idea that there is a constitutive aim of belief. So the basic assumption underpinning the Action-Truth Principle – namely, that there is an end that is intimately tied to action in the way that truth is tied to belief – will be rejected by many.¹³⁸

Even if we grant that there is a constitutive aim of action, a bit of reflection suggests that the Action-Truth Principle isn't a plausible attempt to identify this aim. We can think of examples of inquiry into how to evaluate actions that is intuitively correct and successful despite not resulting in evaluations that are successful *from the point of view of grasping truth*. For example, inquiry into how we ought to act sometimes strikes us as successful when it results in evaluations that are good from the viewpoints of prudence or morality. Suppose that I deliberate about which charity I ought to donate to and I form the judgment that I ought to donate to Oxfam. We can imagine circumstances in which this deliberation strikes us as successful inquiry despite the fact that the act of donating to Oxfam does not promote the formation of true beliefs.

I conclude that premise (1) of the objection from epistemic non-cognitivism can be rejected. As a result, Ingram's appeal to epistemic non-cognitivism fails to undermine the argument from practical epistemic facts that I have been defending, even if it succeeds in undermining the standard formulation of the companions in guilt argument.

¹³⁸ Among those who reject this assumption is Ingram himself, who claims that 'there is nothing as intimately connected to action as truth is connected to belief' (*Ibid.*, 17). Of course, some philosophers have defended the notion that there is a constitutive aim of action, hoping to thereby secure foundations for practical normativity. See Korsgaard (1996) and Velleman (2000) for prominent examples.

2.5 Streumer's Objection from Global Normative Error Theory

In this section, I will further defend the Ontological Premise by rebutting a quite radical argument against it that some moral error theorists would endorse. Suppose that *global normative error theory* is correct: All normative judgments are false because there are no normative facts. It would follow that there are no practical epistemic facts, for as I have argued, practical epistemic facts entail the existence of normative reasons for action. Furthermore, if global normative error theory is true, then a premise of my argument from transmission – my main argument for the Ontological Premise – is false. For the argument from transmission relied on the premise that we sometimes have epistemic duties of belief, and the existence of such duties would surely be a normative fact. So, a critic might appeal to global normative error theory to show both that my argument for the Ontological Premise fails and that we have decisive grounds for rejecting the Ontological Premise.

This challenge to the Ontological Premise is worthy of consideration because global normative error theory has recently been offered impressive defense.¹³⁹ The most sophisticated defense has been developed by Bart Streumer, so I will focus on critiquing Streumer's defense in what follows. One of the ways that Streumer defends global normative error theory is by offering ingenious rebuttals to several *Moorean arguments* against it—that is, arguments against it that are structurally analogous to G.E. Moore's famous argument against skepticism about the external world.¹⁴⁰ I will argue that there is a Moorean argument against global normative error theory that

¹³⁹ Streumer (2008), Streumer (2011), and Streumer (2013a) defend theses that together entail global normative error theory. Streumer (2013b) and Streumer (forthcoming) defend global normative error theory from a number of objections. Cline (forthcoming) also defends global normative error theory.

¹⁴⁰ Moore (1939). Streumer (2013b) critiques two Moorean arguments against global normative error theory, and Streumer (forthcoming) critiques two more.

evades Streumer's criticisms, thereby undermining the objection to the Ontological Premise that is based on global normative error theory.

Let's start by considering the following case:

Fire

Jeff is sitting in front of a campfire, close enough to place his hands on the hot coals. Were he to do so, he would experience severe agony, which is something that he strongly desires not to experience. Placing his hands on the hot coals would not promote any of his desires, nor would it bring about anything of value. Furthermore, Jeff knows that these are the facts of his situation.

Here is a very intuitive claim: There is a reason for Jeff to refrain from placing his hand on the hot coals. It is hard to think of a more intuitive claim than this.

Now, consider what I will call *the Moorean argument against global normative error theory*:

- (1) Global normative error theory is incompatible with the claim that there is a reason for Jeff to refrain from placing his hand on the hot coals.
- (2) The claim that there's a reason for Jeff to refrain from placing his hand on the hot coals seems much more clearly true than global normative error theory.
- (3) If a claim C and a philosophical theory T cannot both be true, and if C seems much more clearly true than T, we should reject T.
- (4) So, we should reject global normative error theory.

This is almost identical to one of the Moorean arguments to which Streumer (2013b) responds. The only difference is that while the above argument relies on a normative claim about Jeff, the Moorean argument that Streumer considers relies on the normative claim that torturing innocent children for fun is morally wrong. But I believe that my variation of the argument is stronger because: (i) It's even more intuitive that Jeff has a reason to refrain from putting his hand in the fire than it is that it is wrong to torture children for fun; and (ii) While our confidence that it's wrong to torture children for fun is threatened by well-known evolutionary debunking arguments,

it's much less clear that our confidence in Jeff's reason is subject to analogous worries.

Premise (1) is obviously true. Premise (2) is also very plausible, though it could be disputed. Streumer might direct our attention to his arguments in favor of global normative error theory, and insist that once we fully grasp these arguments, premise (2) will no longer be true of us. A full defense of this Moorean argument must involve a consideration of these arguments for global normative error theory. But I will bracket this issue for now because Streumer does not take issue with premise (2).

Streumer attacks premise (3). His criticism of this premise relies on the surprising claim that it is literally impossible to believe that global normative error theory is true. Although I am doubtful that his argument for this claim succeeds, I will grant it for the sake of the argument.¹⁴¹ He objects to (3) as follows. (3) entails the following:

If the normative claim that Jeff has a reason to refrain from putting his hand on the hot coals is incompatible with global normative error theory, and if this normative claim about Jeff seems much more clearly true than global normative error theory, then we should reject global normative error theory.

And this entailment of (3) is false because in this case, 'what explains why C seems much more clearly true than T is not that C actually *is* true, but is instead that we cannot believe T.' I interpret Streumer to be offering a novel type of debunking explanation of our intuitions. He's suggesting that: (i) the explanation for why we intuit that the normative claim about Jeff is much more clearly true than global normative error theory is that global normative error theory does not seem very clearly true to us; and (ii) the explanation for why global normative error theory does not seem very clearly true to us is simply that it is impossible to believe that global normative error theory

¹⁴¹ See Hyun and Sampson (2014) for criticism of Streumer's argument for the unbelievability of global normative error theory.

is true. Since (i) and (ii) are true, the correct explanation of our intuition that the normative claim about Jeff is much more clearly true than global normative error theory does not appeal to the veracity of this intuition or to the truth of the normative claim about Jeff. This intuition is therefore debunked, and it consequently fails to make it the case that we should reject global normative error theory in favor of the normative claim about Jeff. Streumer concludes that the above entailment of premise (3) is false. It follows that premise (3) of the Moorean argument against global normative error theory is false.

But this critique of premise (3) is dubious, because we have strong grounds to reject claim (ii) – the claim that the explanation for why global normative error theory does not seem very clearly true to us is simply that it is impossible to believe that global normative error theory is true.¹⁴² To see why, consider the following intuitive thought. If p and $\sim q$ are both true, then

Explanation: If it is possible that $p \ \& \ q$, then the mere fact that p is not sufficient to explain why $\sim q$.

Explanation accounts for the fact that a common strategy for undermining the claim that p fully explains why $\sim q$ is to point out that it is possible that $p \ \& \ q$. If I suggest that the fact that Jack is late fully explains why Jill is not happy, you might reasonably challenge this purported explanation by pointing out that Jill is normally very patient with Jack's tardiness, and thus that it is entirely possible for Jill to be happy even though Jack is late.

Let us suppose, then, that Explanation is true. Explanation entails that if [(p) we cannot believe in global normative error theory and (q) global normative error theory seems very clearly true to us] is possibly true, then the mere fact that (p) we cannot believe in global normative error theory is not sufficient to explain why ($\sim q$) global normative error theory does not seem clearly

¹⁴² The criticism of claim (ii) that follows is drawn from Hyun and Sampson (2014: 637-639).

true to us. And I see no reason to deny that it is possible for the above conjunction (p & q) to be true. After all, a claim can seem very clearly true to a person even while that person does not in fact believe that claim. A moral error theorist, for instance, might not believe that it is wrong to torture children despite the fact that she has the same moral intuitions as the rest of us. Since a claim can seem very clearly true to a person even when that person *does not* believe that claim, it is reasonable to think that a claim can seem very clearly true to a person even when that person *cannot* believe that claim. Since this is so, we should think that [(p) we cannot believe global normative error theory and (q) global normative error theory seems very clearly true to us] is possibly true. And so, we should deny Streumer's claim that the fact that we cannot believe global normative error theory to be true adequately explains why this error theory does not seem very clearly true to us.

Streumer might respond by conceding that the fact that we cannot believe the error theory does not all by itself fully explain the fact that global error theory does not seem very clearly true to us, but insist that this is at least part of the explanation for why it does not seem very clearly true to us. But even if this is so, this conclusion is compatible with the truth of the relevant normative claims playing a role in a full explanation of why these normative claims seem more clearly true to us than global normative error theory. If they do play such a role, then it remains plausible that premise (3) is true even when C is a particular normative claim and T is global normative error theory.

In his forthcoming book, *Unbelievable Errors*, Streumer modifies his discussion of the Moorean challenge to global normative error theory so that the above criticism does not apply. In particular, he presents and critiques a version of the Moorean argument that appeals to (3)* rather than (3):

(3)* If a claim C and a philosophical theory T cannot both be true, and if we are much more confident that C is true than that T is true, this shows that we should reject T.¹⁴³

Streumer objects to (3)* in the same way that he objected to (3), arguing that (3)* is false when C is a particular normative claim and T is global normative error theory because when this is so, ‘what explains why we are much more confident that C is true than that T is true is not that C is actually true, but is instead that we cannot believe T.’¹⁴⁴ And in a footnote, he claims that nothing like the challenge I have raised against his objection to (3) can be offered against his objection to (3)*.¹⁴⁵ I will grant that this is so for the sake of the argument. But this point is not helpful for the global normative error theorist, for all that it shows is that the critic of global normative error theory should not appeal to the version of the Moorean argument that relies on (3)*. Instead, she should appeal to the version that relies on (3), for as we have seen, Streumer’s debunking argument fails against the latter.

To sum up this section, I have defended a Moorean argument against global normative error theory from Streumer’s criticisms. This helps to defend the Ontological Premise from the objection to it that is based on an appeal to global normative error theory.

2.6 Conclusion

In this chapter, I have argued that moral facts and practical epistemic facts stand or fall together, in the sense that we ought to believe in the former if we ought to believe in the latter. And over the course of this chapter and the previous one, I have argued that we ought to believe in practical epistemic facts. I conclude that we should believe that moral facts exist and that the realist’s success thesis is therefore true.

¹⁴³ Streumer (forthcoming: 215)

¹⁴⁴ *Ibid.*, 215.

¹⁴⁵ *Ibid.*, 215, note 13.

Chapter 3: Pure Moral Deference, Social Goods, and Fair Play

3.1 Two Puzzles of Pure Moral Deference

Suppose that I am wondering whether it is morally permissible to eat meat. Instead of thinking through the matter for myself, I decide to defer to your judgment on the matter. I don't defer to you because you know more about factory farming practices than I do. Nor do I defer to you because I take myself to be suffering from some special cognitive impairment, such as sleep deprivation or drunkenness. Rather, I morally defer simply to avoid the need to engage in first-order moral reasoning myself, much like one might defer to a tax specialist about one's taxes. Call this the *Vegetarian Case*.

Many people report having negative attitudes of some sort towards this type of moral deference. For example, David Enoch writes:

There seems to be something fishy about the idea of moral expertise and indeed testimony, the idea of one person forming her moral judgments merely based on the view of another.¹⁴⁶

Sarah McGrath writes:

There is something off-putting about the idea of arriving at one's moral views by simply deferring... There seems to be something problematic about bypassing the struggle of figuring out what one owes, morally speaking, and leaving this for someone else to determine.¹⁴⁷

And Allan Hazlett, while discussing cases similar to the *Vegetarian Case*, writes:

In these cases, there seems to be something bad—something wrong, or problematic, or objectionable—about believing on the basis of testimony, which is not present in [paradigm cases of deference about non-moral matters]. Generalizing from these cases, there seems

¹⁴⁶ Enoch (2014: 229).

¹⁴⁷ McGrath (2011: 111).

to be some kind of asymmetry, at least in some cases, between moral testimony and non-moral testimony.¹⁴⁸

As Hazlett notes, the widespread discomfort that people feel towards moral deference stands in contrast with our attitudes towards analogous cases of non-moral deference. Nobody finds it off-putting to simply defer to the weatherman about the weather or to a tax specialist about one's taxes. Many people feel differently about moral deference.

These widespread negative sentiments towards moral deference cry out for explanation. I will refer to this *explanandum* as Problematic Deference:

Problematic Deference (PD): Pure moral deference seems problematic to many of us.

A few clarifications are required. First, it is important to note that PD is an observation about what *seems* problematic. The fact that cries out for explanation is the *psychological* claim that many people have certain ill feelings, or negative attitudes, towards pure moral deference. I am not taking the *explanandum* to be the *normative* claim that pure moral deference is genuinely illegitimate in some respect.¹⁴⁹ It is compatible with PD that pure moral deference is not genuinely problematic in any respect. As McGrath observes, it may turn out that 'the seeming strangeness of moral deference is mere appearance, and deference about morality is (or ought to be) as straightforward a matter as deference about local geography.'¹⁵⁰

Second, PD does not say that *everyone* has the negative attitudes towards pure moral deference. Instead, it states that *many of us* find pure moral deference problematic. This restriction

¹⁴⁸ Hazlett (2017: 49).

¹⁴⁹ I am following McGrath (2009, 2011) in taking the *explanandum* to be the psychological claim rather than the normative claim. But some philosophers, such as Sliwa (2012) and Hazlett (2017: 50), have taken the normative claim as their *explanandum*.

¹⁵⁰ McGrath (2011: 116).

is necessary because some people do not share the deep discomfort with pure moral deference that has driven much of the recent meta-ethical literature on moral deference.¹⁵¹

Finally, PD's restriction to *pure* moral deference is meant to reflect the fact that the problematic-seeming instances of moral deference are those that involve treating someone else as possessing 'purely moral information' that one lacks.¹⁵² S_1 's deference to S_2 counts as *pure moral deference* if and only if: (a) S_1 judges that a moral proposition M is true on the grounds that S_2 judges that M is true; (b) S_1 does not take S_2 to be better informed about the empirical facts that are relevant to whether M is true; and (c) there is no reason for S_1 to believe that she is suffering from some special sort of cognitive impairment, such as sleep deprivation or drunkenness. The Vegetarian Case is an example of moral deference that satisfies (a)-(c), and is therefore a case of pure moral deference.

Impure moral deference is moral deference that fails to satisfy (b) or (c). I will refer to impure moral deference that fails to satisfy (b) as *impure moral deference_{info}*, an example of which is as follows:

[Nazi Case]

You want to know which Nazi officials committed the most egregious atrocities, but you are quite ignorant about the morally-relevant, historical details about the conduct of various Nazi officials. You have a friend whose moral sensibilities are similar to your own and who has studied these historical details in great depth. You decide to defer to your friend's judgment about which Nazi officials committed the worst atrocities.¹⁵³

¹⁵¹ McGrath (2009: 323) acknowledges that not everyone in our society has the skeptical attitude towards pure moral deference that she and others report having, speculating that 'perhaps this attitude is common only in the relatively secular, post-Enlightenment West, and is more or less absent from many other cultures.' The fact that many philosophers do not share the negative attitudes towards pure moral deference that Enoch, McGrath, and others report having is confirmed by discussions about moral deference that I have had with many philosophers and non-philosophers over the last five years.

¹⁵² McGrath (2011: 115).

¹⁵³ This case is from McGrath (2011: 113).

This is a case of moral deference, but there is a clear sense in which you are treating your friend primarily as one who is in possession of *historical* information that one lacks. Some of those who report having the ill feelings towards moral deference have clarified that they find impure moral deference_{info} untroubling, or at least much less troubling than pure moral deference.¹⁵⁴

I will refer to impure moral deference that fails to satisfy (c) as *impure moral deference_{impaired}*, an example of which is as follows:

[Drunk Case]

At a bar, you make an insensitive joke that offends a nearby friend. Because you are quite drunk, you cannot think clearly about whether you ought to apologize, so you ask the designated driver whether he thinks that you ought to apologize. He says that you should do so, and you defer to him.

This is a case of moral deference, but some of those who report having the negative attitudes towards moral deference have clarified that they do not find the sort of moral deference depicted in the Drunk Case to be very troubling.¹⁵⁵ This may be because impure moral deference_{impaired} is akin to deference to one's 'unimpaired self,' and that consequently, one who engages in impure moral deference_{impaired} is not treating others as a true moral expert relative to oneself.¹⁵⁶

It will be helpful for what follows to mention a special kind of impure moral deference_{impaired} that frequently comes up in the literature – the moral deference of young children.

Consider the following case of moral deference:

[Child Deference]

Cynthia is four years old. She tries to steal her little sister's favorite toy, but her father catches her in the act. He scolds her and tells her that stealing is wrong. Cynthia believes that her act of stealing was wrong simply on the basis of her father's moral testimony.

¹⁵⁴ For example, this clarification is made by Davia and Palmira (2015: 610), Howell (2014: 391), and McGrath (2011: 113).

¹⁵⁵ For example, this clarification is made by Davia and Palmira (2015: 611) and McGrath (2011: 114).

¹⁵⁶ McGrath (2011: 114) offers this suggestion.

Virtually nobody finds this sort of moral deference off-putting, fishy, or problematic.¹⁵⁷ This is unsurprising in light of what has been said above. Because young children normally have cognitive faculties that are still in the process of developing, children's moral deference is plausibly viewed as a special kind of impure moral deference_{impaired}. Since impure moral deference_{impaired} seems relatively unproblematic, it makes sense that the moral deference of young children seems unproblematic even to those who are otherwise averse to moral deference.

It has recently been argued, most forcefully by McGrath, that PD gives us good reason to reject moral realism. The anti-realist argument goes as follows:

- (1) PD is true.
- (2) Certain versions of moral anti-realism provide a good explanation of PD.
- (3) Moral realism does not provide a good explanation of PD.
- (4) If there are two competing theories, T_1 and T_2 , of domain D , and T_1 provides a good explanation of a significant fact F within D whereas T_2 does not, then F counts significantly in favor of accepting T_1 over T_2 .
- (5) So, PD counts significantly in favor of accepting certain versions of moral anti-realism over moral realism.¹⁵⁸

Call this the *Rebutting Puzzle of Pure Moral Deference*. The contentious premises are (2) and (3). In support of premise (3), McGrath invites us to consider why deference oftentimes seems unproblematic in non-moral domains. The reason why pure deference about taxes, the weather, and physics is oftentimes unproblematic is that these domains are ones in which there are varying levels of expertise in judgment. For example, people who are trained in meteorology are more reliable when it comes to judging future weather, and people who have advanced degrees in physics are more reliable when it comes to judgment matters of physics. This is why it's entirely appropriate and legitimate for some people to just defer to others within these non-moral domains.

¹⁵⁷ A number of philosophers have observed that the moral deference of young children seems unproblematic. See, e.g., Hazlett (2017: 50), Hills (2010: 172; 2011: 254; 2013), Howell (2015), Lord (forthcoming), Mogensen (2015: 3-4), and Sliwa (2012: 175).

¹⁵⁸ McGrath (2009, 2011).

McGrath argues that moral realism implies that some of us have expert moral judgment relative to others. After all, the moral realist holds that the moral facts are ‘neither trivial nor completely unknowable,’ and aren’t such that each person is ideally placed to discern whether a moral utterance from her own mouth is true.¹⁵⁹ If this is what moral facts are like, then we should expect that people vary in the reliability of their moral judgment. Since the presence of expert judgment renders deference on matters of weather, taxes, and chemistry appropriate, we should believe that the existence of moral expertise renders pure deference on moral matters appropriate. For this reason, if moral realism is true, we should expect that pure moral deference seems entirely appropriate. So, according to McGrath, moral realism renders PD mysterious, and consequently, premise (3) quite credible.

In support of premise (2), McGrath argues that several anti-realist views – specifically, non-cognitivism, moral subjectivism, and constructivism – can easily explain PD. I will discuss the details of these anti-realist explanations in the next section. But for now, notice that premise (2) seems *prima facie* plausible in virtue of the fact that some anti-realist views imply that there is little interpersonal variation in the reliability of moral judgment. First, old-fashioned emotivism implies that our moral judgments are not truth apt. If this is so, then it is not the case that some people are more reliable at making true moral judgments than others. Second, simple subjectivism implies that each person is ideally placed to discern whether a moral utterance from her own mouth is true. Were each person the authority on the truth of her own moral utterances, there would be little variation in the reliability of people’s utterances on moral matters. These observations, combined with McGrath’s suggestion that pure deference is appropriate in most non-moral

¹⁵⁹ McGrath (2011: 123).

domains in virtue of the fact that they are domains in which there are varying levels of expertise in judgment, render premise (2) *prima facie* plausible.

When meta-ethicists claim that PD generates a challenge for moral realism, they have something like the Rebutting Puzzle of Pure Moral Deference in mind. But there is a second way in which PD undermines moral realism that has been overlooked. It is commonly thought that a main attraction of moral realism is that it can accommodate the ways in which we are disposed to engage in moral reflection and deliberation. To illustrate, consider David Brink's argument for moral realism, which hinges on the claim that 'realism, and realism alone, provides a *natural* explanation or justification of the way in which we do and can conduct ourselves in moral thought and inquiry.'¹⁶⁰ In support of this claim, Brink points to a number of features of our moral thought and discourse that realism accounts for, such as the following:

- C₁: Moral discourse is typically declarative or assertive in form;
- C₂: Many common moral judgments make references to moral properties; and
- C₃: We act as if there are moral facts—for instance, by arguing with others as if there were right answers to moral issues, by engaging in deliberation when we are uncertain about moral issues, and by believing ourselves to be fallible with respect to moral issues.

Because ordinary moral thought and inquiry is characterized by features like C₁-C₃, it is plausible that realism provides the best explanation of this thought and inquiry. Brink takes this 'to establish a presumptive case in favor of moral realism and to shift the burden of proof to the moral antirealist.'¹⁶¹

It is open to the anti-realist to offer the following *Undercutting Puzzle of Pure Moral Deference* against moral realism. PD is an observation about ordinary moral thought and inquiry.

¹⁶⁰ Brink (1989: 24). See also Brink (1984: 112).

¹⁶¹ Brink (1989: 36). For recent discussions of this type of argument for moral realism, see Kurth (2013), Loeb (2007), and Sinclair (2012).

Specifically, it is the claim that we are strongly disinclined to carry on moral inquiry in a certain way—namely, by outsourcing moral reflection. Far from providing a good explanation of PD, moral realism renders it mysterious for reasons explained above. So, there is an important aspect of ordinary moral thought and inquiry that moral realism does not account for. This undermines the claim that realism provides an excellent account of our ordinary moral thought and inquiry as a whole, even if realism provides a natural explanation of a few features of this thought and inquiry, such as features C₁-C₃. Thus, PD undermines moral realism by undermining an important argument in its favor.¹⁶²

So there are really two challenges to moral realism that are generated by considerations about people's attitudes towards pure moral deference, and my goal in this chapter is to defend moral realism from both of them. In section 3.2, I argue that the four main anti-realist accounts of PD that have been suggested in the literature are inadequate, thereby undermining premise (2) of the Rebutting Puzzle of Pure Moral Deference. In section 3.3, I offer an additional objection to the Rebutting Puzzle of Pure Moral Deference, arguing that its proponents face a dilemma that commits them to the falsity of at least one of its premises. In section 3.4, I develop and defend a realist-friendly explanation of PD. In doing so, I aim to put further pressure on the Rebutting Puzzle of Pure Moral Deference as well as to rebut the Undercutting Puzzle of Pure Moral Deference.

¹⁶² Even if moral realism's inability to explain PD does not by itself substantially undermine Brink's premise that moral realism provides a natural explanation of our ordinary moral thought and discourse, the anti-realist can argue that this inability plays an important role in a cumulative case against this premise. For example, the anti-realist could argue against Brink's premise by arguing that realism cannot explain: (i) PD, (ii) our strong resistance to the use of moral beliefs as evidence against scientific theses, and (iii) many people's inclination to assert that in ethics 'it's all relative.' For discussion of (ii), see Barber (2013); and for discussion of (iii), see Loeb (2007).

3.2 Against Anti-Realist Accounts

3.2.1 Simple Non-Cognitivism

Some have suggested that early forms of non-cognitivism can easily explain PD. For example, consider A.J. Ayer's emotivism, according to which the utterance of a moral sentence serves to express some emotional state. If someone says that it is your moral duty to tell the truth, for instance, her utterance 'may be regarded...as the expression of a certain sort of ethical feeling about truthfulness.'¹⁶³ So moral judgments—those mental states expressed by the utterance of moral sentences—are emotional states of one sort of another. Since emotional states are not truth apt, emotivism implies that moral judgments are not truth apt, and this suggests the following explanation of PD:

No-Truth Account

The purpose of deferring to another person's judgment on some matter is to acquire a true judgment. For example, the purpose of deferring to the weatherman about future weather is to acquire a judgment about future weather that is likely to be true. Therefore, there is a point to engaging in pure moral deference only if doing so is a means to acquiring a true moral judgment. But since moral judgments are not truth apt, it is never the case that pure moral deference is a means to acquiring a true moral judgment. Pure moral deference is therefore always pointless, and the (perhaps implicit) recognition that this is so is the cause of people's discomfort with pure moral deference.^{164,165}

¹⁶³ Ayer (1952: 108).

¹⁶⁴ McGrath (2011: 116-117) and Sliwa (2012: 176) may have to have something like the No-Truth Account in mind when they claim that simple non-cognitivism can easily explain PD.

¹⁶⁵ One may worry that my inclusion of a recognitional component in this account (and in each of the anti-realist accounts that follow) renders it obviously implausible, and that I am therefore being uncharitable to the anti-realist. Few people are simple non-cognitivists, so few people are aware of the defects that afflict pure moral deference if simple non-cognitivism is true. But first, I believe that the anti-realist accounts that McGrath has in mind really do incorporate this recognitional component. This is suggested by McGrath (2009: 329-330)'s discussion of the anti-realist accounts that are based on moral skepticism and moral error theory, for example. So I am not misrepresenting my opponents by including this recognitional component in the anti-realist accounts I discuss in the main text. Second, and more importantly, the anti-realist accounts are not made more plausible if their recognitional components are removed, as I argue in section 3.3.

This explanation of PD is available to any version of non-cognitivism that entails that moral judgments are not truth apt.

But the No-Truth Account faces a serious problem: it incorrectly predicts that impure moral deference will seem just as problematic to people as pure moral deference. For if this account is correct in holding both that the purpose of deference is to gain true judgments and that there are no true moral judgments, then impure moral deference is just as pointless as pure moral deference. And if the recognition of this defect in pure moral deference causes people's discomfort with pure moral deference, then we should expect that the recognition of the same defect in impure moral deference (such as the deference depicted in the Nazi Case, the Drunk Case, and the Child Case) will make people feel similarly uncomfortable. But as I mentioned above, people do not find impure moral deference to be off-putting, fishy, or intuitively problematic. So, there are good grounds for rejecting the No-Truth Account.¹⁶⁶

3.2.2 Sophisticated Non-Cognitivism

There is another reason to find the No-Truth Account unattractive: it is available only to those non-cognitivists who hold that moral judgments are not truth apt. Most contemporary non-cognitivists hold that there is a sense in which moral judgments can be true or false. Simon Blackburn draws on a minimalist understanding of truth, according to which what it is to judge *that p is true* is to judge *that p*.¹⁶⁷ Since these non-cognitivists make moral judgments such as 'stealing is morally wrong,' they are willing to accept claims such as 'it is true that stealing is morally wrong' and 'the judgment that stealing is morally wrong is true.' In this way, contemporary non-cognitivists seek to vindicate the practice of talking as if moral judgments are

¹⁶⁶ Hills (2011: 254) presses a similar criticism of the No-Truth Account, arguing that it incorrectly predicts that the moral deference of young children is problematic.

¹⁶⁷ Blackburn (1993).

genuinely truth apt; and consequently, it is unclear whether the No-Truth Account is compatible with the most sophisticated versions of non-cognitivism. It is therefore tempting to think that the most plausible versions of non-cognitivism enjoy no advantage over realism when it comes to solving the puzzles of pure moral deference.¹⁶⁸

Recently, Guy Fletcher has argued that this is not so. Fletcher defends a non-cognitivist account of PD that is available even to quasi-realist versions of non-cognitivism. His account begins with the following premise:

Sentiments Pure Deference Denial (SPDD): It is very difficult to form desire-like moral sentiments (states such as anger, blame, guilt, and resentment) on the basis of pure moral deference.¹⁶⁹

I will explain Fletcher's case for SPDD shortly, but suppose for a moment that it is true. It would then seem to follow that the sophisticated non-cognitivist can explain PD with the following account:

Non-Cognitivist Sentiments Account

Moral judgments just are desire-like moral sentiments. For example, the judgment that killing is morally wrong just is some sentiment along the lines of anger at killers. Since this is so, and since SPDD is true, it is very difficult to form moral judgments on the basis of pure moral deference. Pure moral deference is therefore very difficult to engage in at all, and many cases in which pure moral deference is purportedly depicted are cases in which an agent is trying to do something that they probably cannot do. The (perhaps implicit) recognition that this is so is the cause of people's discomfort with pure moral deference.

For example, in the Vegetarian Case, someone is described as forming the judgment that it is not permissible to eat meat on the basis of pure moral deference. The proponent of the Non-Cognitivist

¹⁶⁸ Sliwa (2012: 176-177) argues that this is so.

¹⁶⁹ Fletcher (2016: 60). Fletcher actually presents his discussion in terms of what he calls *pure, direct moral deference*, which differs slightly from what I have been calling *pure moral deference*. I will apply the basic idea of Fletcher's proposal to the puzzle of pure moral deference as I have construed it above.

Sentiments Account would explain people's discomfort with this instance of pure moral deference as being the result of their recognition that one of the following is the case: (i) the purported deferrer is not *really* engaging in pure moral deference despite her best efforts to do so; or (ii) the deferrer is managing to do something very rare and peculiar. Since the Non-Cognitivist Sentiments Account does not rely on the thesis that moral judgments are not truth apt, this account is available to some sophisticated non-cognitivists.¹⁷⁰

But the Non-Cognitivist Sentiments Account is problematic because it incorrectly predicts that impure moral deference_{impaired}, such as the deference depicted in the Drunk Case and the Child Case, will seem just as problematic to people as does pure moral deference. To see why this is so, it is necessary to understand Fletcher's defense of SPDD, which goes as follows. There are only two possible ways in which one can, by morally deferring, become strongly incentivized to adopt a particular moral sentiment: (a) one can learn through deference that it would be desirable to have the moral sentiment; or (b) one can learn through deference that the moral sentiment is appropriate. So, if pure moral deference ever strongly incentivizes a person to adopt a new moral sentiment, the act of pure moral deference must be either of type (a) or type (b). But it is psychologically difficult – perhaps even impossible – for humans to adopt a moral sentiment through these two types of deference. As Fletcher explains:

Even utterly trustworthy reports of the *desirability* of some desire-like attitude are not something that agents can use to form and revise desire-like attitudes. Even if it would avert some terrible disaster or get you tenure, you cannot feel anger or resentment towards arbitrary objects or persons at will. And moving to cases of deference concerning the appropriateness of an attitude, if you tell me only that I ought to love a particular person or that shame/anger/resentment towards someone is appropriate, I cannot form these attitudes

¹⁷⁰ Obviously, sophisticated non-cognitivists who deny that moral judgments are moral sentiments cannot rely on the Non-Cognitivist Sentiments Account. For example, Gibbard (1990) argues that moral judgments are just acceptances of norms, and Gibbard (2003) argues that they are just acceptances of plans. Since acceptances of norms and plans are not moral sentiments, Fletcher's account is unavailable to proponents of two attractive version of expressivism.

on the basis of your say-so (remembering again the restriction to cases of pure, direct, deference).¹⁷¹

So, no act of pure moral deference that strongly incentivizes a person to adopt a new moral sentiment is ever likely to lead her to adopt this moral sentiment. Fletcher concludes that we should accept SPDD.

Fletcher's argument for SPDD can easily be re-deployed as an argument for an analogous thesis about impure moral deference_{impaired}. Consider the following principle:

Sentiments Impure Deference Denial (SIDD): It is very difficult to form desire-like moral sentiments (states such as anger, blame, guilt, and resentment) on the basis of impure moral deference_{impaired}.

The modified argument would go as follows. Again, there are only two possible ways in which one can, by morally deferring, become strongly incentivized to adopt a particular moral sentiment: namely, the two ways that are captured by (a) and (b). So, if impure moral deference_{impaired} of the sort depicted in the Drunk Case and the Child Case ever strongly incentivizes a person to adopt a new moral sentiment, the act of impure moral deference_{impaired} must be either of type (a) or type (b). But as Fletcher says, it is psychologically difficult or impossible for humans to adopt a moral sentiment through these two types of deference. So, no act of impure moral deference_{impaired} that strongly incentivizes a person to adopt a new moral sentiment is ever likely to lead her to adopt this moral sentiment. So, we should accept SIDD.

I take no stand on whether this argument for SIDD is convincing, but this argument is clearly just as convincing as the argument for SPDD. So, absent the provision of some alternative source of support for SPDD, we should accept SIDD if we accept SPDD.

¹⁷¹ Fletcher (2016: 60-61).

I can now state my worry for the Non-Cognitivist Sentiments Account. Suppose for a moment that this non-cognitivist account of PD is correct. It follows that we should accept SPDD. Since we should accept SPDD, we should also accept SIDD. Since we should accept SIDD and since moral judgments just are moral sentiments, it follows that we should believe that it is very difficult to form moral judgments on the basis of impure moral deference_{impaired}. We should therefore expect people to find impure moral deference_{impaired} to be intuitively problematic and off-putting; for the recognition of this same kind of difficulty causes people to find pure moral deference to be intuitively problematic and off-putting. The above line of reasoning shows that if the Non-Cognitivist Sentiments Account is correct, then we should expect impure moral deference_{impaired} to seem just as problematic as pure moral deference. But again, impure moral deference_{impaired} does not seem just as problematic as pure moral deference. As I explained above, impure moral deference_{impaired} seems relatively unproblematic even to those who have the negative attitudes towards pure moral deference. We therefore have good grounds for rejecting the Non-Cognitivist Sentiments Account.

3.2.3 Subjectivism

According to *simple subjectivism*, a person's utterance that an action is morally wrong (or morally right) expresses the truth-evaluable proposition that that person disapproves (or approves) of that action.¹⁷² To illustrate, this view implies that if you say 'Capital punishment is morally wrong' and I say 'Capital punishment is morally right,' both of our utterances can be true so long as you disapprove of capital punishment while I approve of it. Some think that the simple subjectivist can offer the following account of PD:

¹⁷² McGrath (2009; 2011).

Simple Subjectivist Account

Because each person is typically best-placed to discern what she approves and disapproves of, simple subjectivism implies that each person is typically best-placed to discern whether her own moral utterances and moral judgments are true. Since one's deference to others is irrational when one is better-placed than others to figure out the matter at hand, simple subjectivism can offer the following explanation of people's sense that pure moral deference is problematic: such deference is irrational, and people's unease with pure moral deference is the result of their (perhaps implicit) recognition that this is so.¹⁷³

Importantly, it is not only the simplest versions of subjectivism that can offer an account along these lines. Any meta-ethical theory according to which a person's moral utterance expresses a truth-evaluable proposition about her own mental states can hold that 'the oddity of pure moral deference turns out to be a special case of the oddity of deferring to someone else about one's own mental states.'¹⁷⁴

But there are two problems with any subjectivist explanation of PD. First, like the previous two anti-realist accounts considered above, it incorrectly predicts that impure moral deference_{impaired} will seem problematic to people. Even though children and intoxicated adults are not as good at introspection as sober adults, it is still the case that many children and intoxicated adults are typically best-placed to know what they approve and disapprove of. Since this is so, simple subjectivism implies that many children and intoxicated adults are typically best-placed to discern whether their own moral utterances and moral judgments are true. Since deference to others is irrational when one is better-placed than others to figure out the matter at hand, simple subjectivism implies that these children and intoxicated adults are being irrational when they engage in moral deference. According to the Simple Subjectivist Account, it is the recognition of this type of irrationality that causes people's discomfort with pure moral deference. So if this

¹⁷³ McGrath (2011: 121).

¹⁷⁴ *Ibid.*, 121. Dreier (1990)'s 'speaker relativism' is an example of a sophisticated version of subjectivism that can provide an account along the lines of the Simple Subjectivist Account.

account is correct, we should expect that the recognition of the same sort of irrationality in cases such as the Drunk Case and the Child Case will make people feel similarly uncomfortable. But again, this is not what we find. Virtually everyone is comfortable with the moral deference of children, and those who are uncomfortable with pure moral deference are comparatively sanguine about the deference depicted in the Drunk Case. So, we have reason to reject the Simple Subjectivist Account.

Second, there are many cases of pure moral deference whose felt fishiness is not explained by the Simple Subjectivist Account. To see this, consider a modified version of the Vegetarian Case:

[Modified Vegetarian Case] You overhear Jake say, ‘It is wrong to eat meat.’ You want to know whether his moral utterance is true. Instead of thinking through the matter yourself, you ask Jill, ‘Was Jake’s moral utterance true?’ Jill says ‘yes,’ and you decide to defer to her judgment. You do not defer to her because she has more empirical information that is relevant to whether Jake’s moral utterance is true: you and Jill are equals in that respect. Nor do you defer to her because you take your judgment to be especially impaired at the moment. Rather, you morally defer to Jill simply to avoid the need to think through the matter yourself, much as you might defer to a tax specialist about your taxes in order to save time.

This is a case of pure moral deference; I assume that the discomfort that people feel towards the deference in the original Vegetarian Case is also felt towards the deference in the Modified Vegetarian Case. Yet the oddity of this latter deference cannot be ‘a special case of the oddity of deferring to someone else about one’s own mental state,’ as the Simple Subjectivist Account must maintain. This is because simple subjectivism implies that your deference in the Modified Vegetarian Case is tantamount to deference about *Jake’s* mental states, not your own. Since there is nothing particularly irrational about deferring about someone else’s mental states, the Simple

Subjectivist Account cannot explain the discomfort that people feel towards the Modified Vegetarian Case. It therefore fails to provide a full explanation of PD.

3.2.4 Constructivism

Constructivism about a normative domain is the view that the normative truths within that domain are fundamentally determined by the result of some duly-specified deliberative procedure.¹⁷⁵ For example, Bernard Williams is a constructivist about the domain of practical reasons, holding that the truths about which actions a person has reason to perform are fundamentally determined by what she would desire to do upon completing a process of sound deliberation.¹⁷⁶ Constructivism about morality similarly holds that moral truth depends on what emerges from our completion of some specified deliberative procedure.¹⁷⁷ McGrath argues that constructivism about morality provides the following explanation of PD:

Constructivist Account

Engaging in pure moral deference involves ‘treating the moral domain as a repository of antecedently existing facts,’ or facts that obtain independently of the operations of our own minds. But constructivism about morality is true, so ‘it is a deep confusion to suppose that there is some independent realm of moral facts or truths that we are attempting to bring into view and to which our moral judgments are answerable.’ The discomfort that people feel towards pure moral deference is caused by their (perhaps implicit) recognition that the deferrers are mistakenly treating the moral domain as a domain of non-constructed facts.¹⁷⁸

I will grant for the sake of the argument that pure deference about morality amounts to treating the moral domain as a domain about which constructivism is false. The Constructivist Account fails

¹⁷⁵ This is a common construal of constructivism – for example, see Darwall, Gibbard, and Railton (1992).

¹⁷⁶ Williams (1979; 1989).

¹⁷⁷ For example, Korsgaard (1996) defends a version of this sort of constructivism.

¹⁷⁸ McGrath (2011: 117).

to comport with the fact that pure deference seems unproblematic in some non-moral, normative domains.¹⁷⁹ To see this, consider the following case of pure deference about fashion:

[Fashion Deference Case]

I am preparing to leave for a dinner party, and I need to decide whether it would be unfashionable to wear a brown jacket given the nature of the rest of my attire. Instead of thinking through the matter myself, I ask my wife whether it would be unfashionable. She says that it would be unfashionable, and I decide to defer to her judgment. I do not defer to her because she knows more about the empirical facts that are relevant to the requirements of fashion. (For example, we are equally knowledgeable about such matters as the precise shade of brown possessed by my jacket, the color of the rest of my articles of clothing, and the size of the jacket.) Further, I do not defer to my wife because I take my own judgment to be especially impaired by conditions such as drunkenness or sleep deprivation. Rather, I defer to her in order to avoid having to think through the matter myself.

Unlike pure moral deference, the pure fashion deference depicted in this case is commonplace and seems unproblematic. Next consider a case of pure deference about etiquette:

[Etiquette Deference Case]

I need to send a work-related email tonight, and I am uncertain whether it is impolite to use the recipient's first name in the opening salutations. My friend was recently in the same situation, so instead of thinking through the matter myself, I ask for his opinion. He says that it would indeed be impolite to use the recipient's first name, and I defer to his judgment. I do not defer to him because he knows more about the empirical facts of the situation. (For example, we are equally knowledgeable about such matters as the recipient's authoritative position relative to my own, the number of previous interactions that I have had with the recipient, and the level of formality that typically characterizes email exchanges within my industry.) Further, I do not defer to my friend because I am cognitively impaired at the moment. Rather, I defer to him simply because the email needs to be sent tonight and I do not have time to carefully come to a decision on my own.

Again, this sort of pure deference about the requirements of etiquette does not provoke the sort of negative reactions that many philosophers report having towards pure moral deference.

¹⁷⁹ By a 'normative' domain, I simply mean a domain that has to do with rules or standards of correctness. In Parfit (2011: 308-310)'s terminology, I am referring to domains that are normative in the 'rule-implying' sense.

Suppose that the Constructivist Account is correct. Then all of the following are true: (a) pure deference about morality amounts to treating the moral domain as a domain of facts that obtain independently of our mental processes; (b) this way of treating the moral domain counts as a mistake since constructivism about morality is true; and (c) the recognition that (a) and (b) are true is the cause of people's discomfort with pure moral deference. If we should believe (a), then we should also believe that pure deference about fashion and etiquette amounts to treating these normative domains as domains of facts that obtain independently of our mental processes. But some kind of constructivism surely provides the best account of the domains of fashion and etiquette: nobody thinks that fashion and etiquette facts are robustly mind-independent in the way that moral realists believe moral facts to be robustly mind-independent. So, if we should believe (a), then we should believe that pure deference about fashion and etiquette amounts to *mistakenly* treating these normative domains as domains of non-constructed facts. According to (c), this is the same type of defect whose recognition causes people's discomfort with pure moral deference. It follows from the above reasoning that if we should accept the Constructivist Account, then we should expect that the instances of pure deference about etiquette and fashion that are depicted in the Etiquette Deference Case and the Fashion Deference Case strike people as problematic in the same way that instances of pure moral deference strike them as problematic. Because this is not so, we have good reason to reject the Constructivist Account.

I have argued that the four main anti-realist accounts of PD that have been proposed in the literature are inadequate. This provides us with strong grounds for rejecting premise (2) of the Rebutting Puzzle of Pure Moral Deference.

3.3 A Dilemma for the Rebutting Puzzle of Pure Moral Deference

I now want to defend a second, independent criticism of the Rebutting Puzzle of Pure Moral Deference, so let us now grant for the sake of the argument that my arguments in the previous section fail. As I emphasized in section 3.1, not all people have the negative attitude towards pure moral deference that we have been discussing. Call the group of people who do have this negative attitude *the anti-deferrers*. And call the versions of anti-realism whose truth implies that pure moral deference is flawed in some respect *the PD-friendly versions of anti-realism*.¹⁸⁰ Proponents of the Rebutting Puzzle of Pure Moral Deference face the following *Dilemma for Rebutters*:

- (1) Either all anti-deferrers believe in some PD-friendly version of anti-realism, or this is not so.
- (2) If all anti-deferrers believe in some PD-friendly version of anti-realism, then the third premise of the Rebutting Puzzle of Pure Moral Deference is false, and moral realism does, after all, provide a good explanation of PD.
- (3) If it's not the case that all anti-deferrers believe in some PD-friendly version of anti-realism, then the second premise of the Rebutting Puzzle of Pure Moral Deference is false, and anti-realist views fail, after all, to provide good explanations of PD.
- (4) So, either the second premise or the third premise of the Rebutting Puzzle of Pure Moral Deference is false.

To see why we should accept premise (2), consider the following observation: One way to plausibly explain why a person has a negative attitude towards X is to show that she has some belief whose truth would imply that X is genuinely problematic in some respect. For example, if Alice feels uneasy about driving her car, and if she (perhaps incorrectly) believes that her car has been rigged to explode when it is driven, then it is not hard to think of an explanation of Alice's unease. Her unease is plausibly explained by the fact that she holds the belief that driving the car would be a fatal mistake. Similarly, if Bert finds Chad off-putting, and if Bert believes Chad to be

¹⁸⁰ Perhaps all versions of anti-realism are PD-friendly versions of anti-realism, but I can remain neutral on whether this is so.

an inconsiderate person, then it is easy to provide a plausible explanation of Bert's negative attitude towards Chad: Bert's negative attitude is plausibly explained by the fact that he believes Chad to be inconsiderate.

Premise (2) is very plausible in light of these considerations. For suppose that all anti-deferrers believe in some PD-friendly version of anti-realism. Moral realists could then provide the following plausible explanation of PD: anti-deferrers are uncomfortable with pure moral deference because each of them possesses a belief—specifically, a belief in some kind of moral anti-realism—whose truth would imply that pure moral deference is flawed in some respect. For example, a moral realist can explain the simple non-cognitivist's negative attitudes towards pure moral deference by appealing to the fact that her beliefs imply that pure moral deference is pointless. And a moral realist can explain the constructivist's negative attitudes towards pure moral deference by pointing out that her beliefs imply that such deference mistakenly treats the moral domain as a domain of realistic facts. So, we should accept premise (2) of the Dilemma.

Now consider premise (3). Suppose that it is not the case that all anti-deferrers believe in some PD-friendly version of anti-realism. That is, some of the anti-deferrers believe in moral realism, or believe in some version of anti-realism whose truth does not imply that pure moral deference is flawed in some respect, or are agnostic about which meta-ethical theory is correct. As the four anti-realist accounts discussed in the previous section illustrate, anti-realists' natural strategy for explaining PD – and the only anti-realist strategy for explaining PD that has been proposed in the literature – involves two steps: first, identify some version of anti-realism that implies that pure moral deference suffers from some flaw F; and second, posit that people's discomfort with pure moral deference is the result of their recognition of F. Clearly, the only versions of anti-realism that can potentially succeed at carrying out this strategy are the PD-

friendly versions of anti-realism. So, at least until some novel anti-realist strategy for explaining PD is discovered, we should accept:

- (A) Certain types of moral anti-realism provide a good explanation of PD only if some PD-friendly version of anti-realism provides a good explanation of PD.

And the following two claims are true: A good explanation of PD must explain the negative attitudes of *all* the anti-deferrers; and some of the anti-deferrers do not believe in any PD-friendly version of anti-realism. So, we ought to accept:

- (B) Some PD-friendly version of anti-realism provides a good explanation of PD only if some PD-friendly version of anti-realism explains the negative attitudes of anti-deferrers who do not believe in any PD-friendly version of anti-realism.

Finally, we should accept:

- (C) No PD-friendly version of anti-realism explains the negative attitudes of anti-deferrers who do not believe in any PD-friendly version of anti-realism.

To see why we should accept (C), it will be helpful to start by focusing on why a couple of specific PD-friendly versions of anti-realism fail to explain the negative attitudes of those anti-deferrers who are convinced moral realists. First, consider why simple non-cognitivism fails to explain a moral realist's negative attitudes towards pure moral deference. Recall that simple non-cognitivism's account of PD is the No-Truth Account, which posits that people's discomfort with pure moral deference is caused by their recognition of the fact that pure moral deference is pointless in virtue of never being a means to acquiring true moral judgments. This account can successfully explain the moral realist's discomfort with pure moral deference only if the moral realist does in fact recognize that pure moral deference is pointless in this way. But surely the moral realist does not believe that pure moral deference is pointless in this way, and hence does

not recognize that it is pointless in this way. So, the No-Truth Account fails to explain the moral realist's negative attitudes towards pure moral deference.

Second, consider why simple subjectivism fails to explain a moral realist's negative attitudes towards pure moral deference. Recall that simple subjectivism's account of PD is the Simple Subjectivist Account, which posits that people's discomfort with pure moral deference is caused by their recognition of the fact that pure moral deference is irrational in virtue of the fact that it constitutes deference to others about one's own mental states. The moral realist surely does not believe that pure moral deference is irrational in this way, and hence does not recognize that it is. So, the Simple Subjectivist Account fails to explain the moral realist's negative attitudes towards pure moral deference.

Analogous reasoning can be constructed to show of any PD-friendly version of anti-realism that its account of PD cannot provide a good explanation of the negative attitudes of anti-deferrers who do not accept any PD-friendly version of anti-realism. The purported flaws of pure moral deference that figure in these anti-realist accounts of PD are unlikely to be recognized by those anti-deferrers who do not accept any PD-friendly version of anti-realism. So, we should accept (C).

Against this defense of (C), it might be objected that I have been making an unwarranted assumption all along about the explanatory strategies that are available to anti-realists who seek to explain PD. I have been assuming that the anti-realist's account must include a recognitional component—that it must explain the discomfort with pure moral deference by appealing to *the recognition* of some flaw that afflicts pure moral deference. But perhaps anti-realist accounts should be understood as claiming that people's discomfort with pure moral deference is caused by *the fact that it is flawed in such-and-such respects* rather than by *people's recognition of the fact*

that it is so flawed. On this approach, the No-Truth Account holds that the moral realist's negative attitudes are directly explained by the fact that pure moral deference is pointless in a way that the moral realist has not managed to notice; and the Simple Subjectivist holds that the moral realist's negative attitudes are directly explained by the fact that pure moral deference is irrational in a way that she has overlooked. Understood in this way, the PD-friendly versions of anti-realism could in principle explain the negative attitudes of anti-deferrers who reject all PD-friendly versions of anti-realism.

Anti-realists are free to remove the recognitional component of their accounts in this way, but the resulting anti-realist accounts are deeply implausible. Typically, the fact that X has flaw F does not explain a person's negative attitude towards X if she is completely unaware that X has F. For example, suppose that I am averse to drinking the milk that is in my refrigerator. If this milk has gone sour, and if I have no idea that this is so since I have not seen or tasted it in days, then it is unlikely that the milk's having gone sour explains why I am averse to drinking the milk today. And if Alice is secretly racist and Bert has no idea that this is so, then it is unlikely that Alice's racism explains Bert's aversion to Alice. Analogously, if pure moral deference has a certain flaw F, it is unlikely that this fact explains the negative attitudes towards pure moral deference of people who are completely unaware that pure moral deference has F. This is why anti-realist accounts that lack the recognitional component are implausible. Because they are implausible, anti-realists should not appeal to them in order to evade the force of my argument for (C).

So, we should accept (C). And it follows from the conjunction of (A), (B), and (C) that it is not the case that certain types of moral anti-realism provide a good explanation of PD. Recall that we have reached this conclusion on the supposition that it is not the case that all anti-deferrers believe in some PD-friendly version of anti-realism. So, we should believe that if it is not the case

that all anti-deferrers believe in some PD-friendly version of anti-realism, then it's not the case that certain types of moral anti-realism provide a good explanation of PD. In other words, we should accept premise (3) of the Dilemma for Rebutters.

I conclude that the Rebutting Puzzle of Pure Moral Deference is unsound, for either its second premise is false or its third premise is false.

3.4 A Realist-Friendly Account

I have defended two objections to the Rebutting Puzzle of Pure Moral Deference, but the Undercutting Puzzle of Pure Moral Deference remains unscathed. In order to rebut this latter challenge to moral realism, I will seek to defend an account of PD that is available to the moral realist. Before delving into specifics, it will be helpful to describe the general shape of the realist explanation I endorse.

We should distinguish between unified accounts and disjunctive account of PD. A *unified account* works by directing attention to a single type of flaw that is present in all of the cases of pure moral deference about which people have negative attitudes, and positing that people's discomfort with pure moral deference is caused by their recognition of this flaw. All of the anti-realist accounts considered in section 3.2 are unified accounts, for each of them seek to identify a single flaw that afflicts pure moral deference in virtue of the truth of some version of anti-realism. A *disjunctive account* works by directing attention to a set of flaws that is such that each of the problematic-seeming cases of pure moral deference exhibits at least one of the flaws in this set, and positing that people's discomfort with pure moral deference is caused by their recognition of one of these flaws.

Some think that a unified account is better than a disjunctive account, all else equal.¹⁸¹ While I agree that this is so, I endorse a disjunctive account of PD for two reasons. First, providing a plausible disjunctive account would suffice for my purpose, which is to rebut both the Undercutting Puzzle of Pure Moral Deference and the Rebutting Puzzle of Pure Moral Deference. Second, the literature on moral deference includes many diverse attempts to identify the source of the apparent badness of pure moral deference. Some attempts appeal to epistemic flaws of pure moral deference; other attempts appeal to moral defects of pure moral deference; and still other attempts appeal to defects that are neither epistemic nor moral in flavor.¹⁸² Some have appealed to defects of the act of deference while others have appealed to defects of the deferrers.¹⁸³ Reflection on these facts make it plausible that there are really multiple defects of pure moral deference the recognition of which is responsible for people's discomfort with it. For this reason, a disjunctive account seems preferable to a unified account, all things considered.

In order to develop a full, realist-friendly, disjunctive account of PD, one must direct attention to a set of flaws that satisfies two conditions: (a) each problematic-seeming case of pure moral deference is afflicted by at least one of the flaws in this set; and (b) each of the flaws in this set is a *realist-friendly flaw*, that is, a flaw that's such that its afflicting pure moral deference is not in tension with the truth of moral realism. Call the set of flaws that satisfies these two conditions the *realist-friendly set of flaws*. Some of the flaws that belong in this set have already been

¹⁸¹ For example, Enoch (2014) opines that 'at least if other things are equal, it would be preferable on methodological grounds to have a unified explanation of [the discomfort with moral deference], one that applies wherever the discomfort applies' (254).

¹⁸² McGrath (2009) offers an epistemic, 'Socratic' account of PD; Howell (2014) offers a moral account of PD; and Fletcher (2016) offers an account that is neither clearly moral nor clearly epistemic.

¹⁸³ Enoch (2014) argues that acts of moral deference are problematic in virtue of making the deferrer forego the opportunity to realize a valuable kind of moral achievement. Howell (2014) argues that it is sometimes the case that the discomfort with pure moral deference is caused by one's recognition of an absence of virtue in the deferrer prior to her choice to defer.

identified by realists who have contributed to the recent debate about moral deference.¹⁸⁴ In what follows, I seek to contribute to the development of a satisfactory realist account by directing attention to two further important defects of pure moral deference that also warrant inclusion in the realist-friendly set of defects.

3.4.1 The Social Goods Account

Most realist-friendly accounts of PD argue that moral deference is problematic in virtue of the negative consequences that moral deference has for the individual deferrer. For example, it has been argued that moral deference strikes people as problematic because it undermines the deferrer's moral character,¹⁸⁵ moral understanding,¹⁸⁶ appreciative moral knowledge,¹⁸⁷ ability to perform morally worthy actions,¹⁸⁸ and epistemic rationality.¹⁸⁹ A different strategy is to argue that moral deference is problematic in virtue of the negative consequences that moral deference has for the deferrer's society.¹⁹⁰ I consider this latter strategy to be quite underexplored, and I will pursue it in what follows.

To begin, notice that pure moral deference is an unusual phenomenon in our society. Most people form their moral beliefs non-deferentially. If you try to think of the last time that you formed a moral belief by deferring to a trusted friend or colleague, chances are that you will find it difficult to think of a recent example. In our society, most adults habitually think through moral matters on their own, reflect on the reasons for and against various positions on moral issues, and believe

¹⁸⁴ I find many of the realist-friendly explanations of PD developed in the literature quite attractive as partial accounts of PD, especially those developed in Davia and Palmira (2015), Groll and Decker (2014), Hazlett (2017), Hills (2013), and Sliwa (2012).

¹⁸⁵ Howell (2014).

¹⁸⁶ Hills (2013).

¹⁸⁷ Lord (forthcoming).

¹⁸⁸ Markovits (2012) and McGrath (2011).

¹⁸⁹ Davia and Palmira (2015).

¹⁹⁰ Hazlett (2017) is the first to propose this shift to a 'social account.' The account that I develop below builds upon his social account.

what seems true by their own lights. I will call such activity *autonomous moral inquiry*, and I will refer to the group of people in our society who always engage in autonomous moral inquiry *the community of autonomous moral inquirers*.

In calling this practice *autonomous moral inquiry*, I do not mean to suggest that it is an entirely personal or individualistic activity. Autonomous moral inquiry is frequently an interpersonal, cooperative activity in which we seek others' moral opinions and consider the merits of their views in an effort to come to our own moral conclusions. The important sense in which the practice of autonomous moral inquiry is autonomous is that its participants always form their moral beliefs non-deferentially. If I am trying to figure out what to believe about abortion, for example, I may well ask you about what you think of the matter. Listening to you explain your views is likely to be helpful. You may raise some points that I had not considered. But if your arguments for your views do not ultimately seem persuasive to me, I am unlikely to adopt your moral views on abortion just because I trust your judgment.

The foregoing reflections render the following thesis plausible:

Autonomous Moral Inquiry Thesis: In our society, there is a large community of autonomous moral inquirers.

The first step towards developing my realist-friendly account is to defend this thesis. One might challenge this thesis by directing attention to the practices of religious communities. In such communities, moral deference to some religious text (such as the Bible) or to some religious authority (such as the Pope) seems to be common. But this is not clear once we distinguish between moral deference and non-moral deference about morally-relevant matters. If a religious person believes that it is wrong to ϕ because she believes that God forbids ϕ -ing, and if she believes that God forbids ϕ -ing because the Bible says so, this is not really a case of *moral* deference. Rather,

it is a case of deference about what God forbids, and the religious person forms her moral belief by making an inference from a claim about what God forbids.

Further, when the Bible says that it is wrong to ϕ , there is normally effort on the part of religious folks to discern *why* ϕ -ing is wrong. And once a conclusion is drawn about the underlying reasons for the wrongness of ϕ -ing, it seems plausible that the religious person's belief that it is wrong to ϕ is based directly on her beliefs about these underlying reasons. For example, a religious person might initially believe that gossiping is wrong because the Bible says so, but upon reflection, come to believe that gossiping is wrong in virtue of its tendency to rupture relationships. Once she achieves this insight, it seems that her belief that gossiping is wrong is not wholly deferential. So, even in religious groups where the practice of pure moral deference is probably more common than elsewhere, moral deference is not as prevalent as it initially appears to be.

A second challenge to the Autonomous Moral Inquiry Thesis goes as follows. In our society, there exists a large amount of agreement on a lot of moral matters. For example, nearly everybody believes that slavery is wrong. It may be argued that the best explanation of this convergence of moral opinion is that we are all taught that slavery is wrong from a young age. Thus, it seems that moral deference to conventional moral wisdom is extremely common. But first, even though people's beliefs about the wrongness of slavery might initially be formed by deference, it is typical for people to come to appreciate some of the underlying reasons why slavery is wrong. Once they do, their belief that slavery is wrong may be based on this moral insight into the underlying reasons, rather than on deference to others. And second, it is important to keep in mind that there is a major way that society influences our moral beliefs that does not involve moral deference, although it may initially seem to. When everybody around us holds a certain moral view, this affects how things *seem* to us. A popular moral belief can start to seem obviously true,

and we may well adopt the moral belief on the basis of the seeming. So long as the grounds of the moral belief is the seeming, this is not really a case of moral deference. When I say that moral deference is rare and that most adults in our society are members of the community of autonomous moral inquirers, I do not mean to deny that our society influences the moral intuitions that ground many of our moral beliefs.

The next step towards developing my realist-friendly account of PD is to establish:

Benefits Thesis: Each of us is greatly benefited by the activities of the large community of autonomous moral inquirers.

There are at least five ways in which we are benefited by the existence of a large community of autonomous moral inquirers. First, the existence of this sort of community facilitates our acquisition of justified moral beliefs, and consequently, it helps us to gain moral knowledge. It does this in two ways. To introduce the first of these ways, consider the following remarks of John Stuart Mill's:

He who knows only his own side of the case knows little of that. His reasons may be good, and no one may have been able to refute them. But if he is equally unable to refute the reasons on the opposite side, *if he does not so much as know what they are*, he has no ground for preferring either opinion. The rational position for him would be suspension of judgment.¹⁹¹

As Mill observes, one is justified in holding a belief only if one is aware of what can be said against it. If one has no idea what can be said against one's belief, then for all one knows, the belief is opposed by devastating arguments. The collective efforts of the community of autonomous moral inquirers therefore facilitates our possession of justified moral belief, for these efforts make us far more aware of what can be said against our moral beliefs than we otherwise would be. After all,

¹⁹¹ Mill (1978: 35), my emphasis.

the task of discovering and formulating all of the important arguments that can be made for and against each position on each moral issue is an extremely large and difficult task. It stands to reason that success at this task is promoted when a vast number of people are applying their minds to it, thinking through the moral issues themselves, and trying to discern the reasons for and against various views. If most people in our society simply deferred on moral matters, our collective knowledge of the arguments for and against various moral beliefs would be much impoverished. This collective knowledge would be limited to what the smaller group of non-deferrers could come up with by themselves. So, the community of autonomous moral inquirers generates extensive knowledge regarding what can be said against various moral beliefs, and therefore facilitates our acquisition of justified moral belief and moral knowledge.

The second way in which the community of autonomous moral inquirers helps us to acquire justified moral beliefs and moral knowledge is by providing us with an important method for testing our moral beliefs' plausibility. As Justice Oliver Wendell Holmes Jr. observed, 'the best test of truth is the power of the thought to get itself accepted in the competition of the market.'¹⁹² The existence of a large community of autonomous moral inquirers generates a lively, public marketplace of ideas in which we can test out our moral views, and thereby gain a better grasp on how likely they are to be true.

Second, the existence of a large community of autonomous moral inquirers benefits us by facilitating our acquisition of moral understanding. One way in which it does this is clear from what has been said above. The widespread practice of autonomous moral inquiry produces a very large body of knowledge about the important arguments for and against various moral beliefs. Many such arguments reveal the underlying reasons why various actions are right or wrong. For

¹⁹² Quoted from Feinman (2014).

example, when one learns about Don Marquis' famous 'future like ours' argument against abortion, one gains not only a piece of evidence for the wrongness of abortion, but also a grasp on one of the underlying reasons why abortion is sometimes morally problematic (on the assumption that Marquis' argument is sound).¹⁹³ Since an appreciation of the arguments for and against belief in a moral claim often brings with it an appreciation of the reasons why that moral claim is true or false, the large community of autonomous moral inquirers provides us with access to a lot of information about such underlying moral reasons. And the provision of this sort of information is conducive to our acquiring moral understanding. So, the community of autonomous moral inquirers helps us to acquire moral understanding.

There is a second important way in which this community helps us to acquire moral understanding. Consider the following abilities: (i) the ability to identify the important moral considerations for and against an action; (ii) the ability to responsibly weigh competing moral considerations; (iii) the ability to infer a moral claim from the facts that make that moral claim true; (iv) the ability to appreciate and follow others' explanations of the truth of their moral beliefs; and (v) the ability to explain in one's own words the reasons that favor one's own moral beliefs. How can one acquire and hone these abilities? In general, one can acquire and hone an ability by observing others exercising that ability well. One may get better at painting by closely observing skillful painters. And one can become better at philosophy by surrounding oneself with good philosophers and paying attention to how they do philosophy. It therefore stands to reason that one way to get better at abilities (i)-(v) is to observe others exercising them well. And because of the community of autonomous moral inquirers, there are many people who can exercise these abilities well. If this community did not exist—that is, if most people normally deferred to some moral

¹⁹³ Marquis (1989).

guru—most people would not be well-practiced at these abilities. But because most people think about moral matters themselves rather than simply deferring, many people in our society are well-practiced at abilities (i)-(v). Consequently, we have access to many good models of the skillful exercising of these abilities, and the community of autonomous moral inquirers have the effect of facilitating our acquisition of these abilities.

Some have argued that these abilities are necessary for, and perhaps even constitutive of, moral understanding.¹⁹⁴ Even if this is too strong of a claim, it is at least true that the possession of abilities (i)-(v) is conducive to the acquisition of such understanding. So, we should conclude that the existence of a large community of autonomous moral inquirers facilitates our acquisition of moral understanding.

A third way in which the community of autonomous moral inquirers benefits us is by helping us to perform morally worthy actions. On one popular and plausible account, one performs a morally worthy action only if one performs a right action for the reasons why that action is right.¹⁹⁵ For example, if one donates to famine relief, one's act of charity is morally worthy only if it is done out of concern for the well-being of the famine victims. If the act of charity is done out of self-interest or out of a *de dicto* concern for doing what is right, then one's act does not count as morally worthy. On this account of morally worthy action, we need moral understanding of why right acts are right in order to be well-placed to perform morally worthy actions. As I have argued above, such moral understanding is facilitated by the activities of the community of autonomous moral inquirers. So, on the assumption that this plausible account of morally worthy actions is correct, the community of autonomous moral inquirers promotes our ability to perform morally worthy actions.

¹⁹⁴ Hills (2009: 102).

¹⁹⁵ See Markovits (2010) for defense of this view.

Fourth, this community helps to decrease the rate at which our moral rights are violated. To see how this is so, first notice that a person's possession of justified moral beliefs and moral understanding is conducive to her reliably acting rightly. One who has a lot of moral understanding—a firm, systematic grasp on the reasons why various acts are right and why other acts are wrong—will tend to be skilled at discovering which acts that she can perform are morally permissible. And one who has a lot of justified moral beliefs is more likely to have true beliefs about which actions available to her are morally permissible than people who are comparatively morally ignorant. Most people desire to act rightly. So, the possession of justified moral beliefs and moral understanding is conducive to one's reliably acting rightly.

Because the possession of justified moral beliefs and moral understanding is conducive to one's reliably acting rightly, anything that facilitates people's acquisition of justified moral beliefs and moral understanding also facilitates people's reliably acting rightly. As I have already argued, the community of autonomous moral inquirers facilitates people's acquisition of justified moral beliefs and moral understanding. So, this community also facilitates people's reliably acting rightly. In general, when people are reliably acting rightly, this makes it less likely that they will wrong you and thereby violate your moral rights. So, the community of autonomous moral inquirers makes it less likely than it otherwise would be that people will violate your moral rights.

Finally, the existence of a large community of autonomous moral inquirers provides us with opportunities for a rich panoply of intellectual and aesthetic pleasures. Because most people engage in first-order moral inquiry rather than simply morally deferring, and because an obvious component of responsible first-order moral inquiry is seeking out moral discussion with others, it is reasonable to believe that the existence of a large number of autonomous moral inquirers

generates much of the first-order moral dialogue and debate that we see in our society.¹⁹⁶ And consider the vast array of goods that are produced by first-order moral dialogue and debate and our disposition to engage in them: stimulating conversation with friends and acquaintances, opportunities to enroll in college ethics courses and to attend ethics lectures, engaging political discourse, and fine literature the production of which is driven by the goals of moral exploration and moral persuasion.¹⁹⁷ We should therefore conclude that the community of autonomous moral inquirers indirectly provides us with the enjoyment of many great intellectual and aesthetic goods.

In light of the foregoing discussion, it is quite plausible that the consequences of pure moral deference are oftentimes sub-optimal in certain respects. When a person decides to engage in pure moral deference rather than to be a member of the community of autonomous moral inquirers, she foregoes the opportunity to help sustain the community that contributes to the production of the five great benefits canvassed above. That is, she foregoes the opportunity to:

- (a) help others gain justified moral beliefs and moral knowledge,
- (b) help others gain moral understanding,
- (c) help others to perform morally worthy actions,
- (d) help to decrease the rate at which others' moral rights are violated, and
- (e) help to produce certain intellectual and aesthetic goods.

There is therefore a way in which the choice to engage in pure moral deference can lead to sub-optimal social consequences. This suggests the following account of PD:

Social Goods Account

Those who engage in pure moral deference are failing to engage in autonomous moral inquiry. As a result, they forego the opportunity to contribute to the production of a number of great benefits that are enjoyed by many people. So, the consequences of pure moral

¹⁹⁶ Hazlett (2017) also notes this connection between our disinclination to morally defer and the existence of widespread first-order moral dialogue.

¹⁹⁷ A lot of literature fits this description. For example, *Oliver Twist* and 'A Modest Proposal' aim to morally criticize society's treatment of the poor; *Uncle Tom's Cabin* denounces slavery; and *The Brothers Karamazov* explores the problem of evil and the meta-ethical implications of atheism.

deference are oftentimes sub-optimal, and people's discomfort towards pure moral deference is sometimes caused by their (perhaps implicit) recognition of this fact.

The Social Goods Account provides some explanatory relief for the realist, and should be incorporated into the realist's disjunctive account of PD.

3.4.2 The Fair Play Account

A second account becomes apparent once we combine some of the foregoing ideas with a plausible version of the principle of fair play. A. John Simmons helpfully articulates the basic idea of this principle as follows:

In its most general possible form, the principle of fair play asserts that those who benefit from the good-faith sacrifices of others, made in support of a mutually beneficial cooperative venture, have a moral obligation to do their parts as well (that is, to make required sacrifices) within the venture.¹⁹⁸

To appreciate the intuitive appeal of the principle of fair play, consider its application to the following thought experiment of Michael Huemer's:

You are in a lifeboat with several other people. You are caught in a storm, and the boat is taking on water, which needs to be bailed out. Other passengers take up containers and start bailing. The other passengers' efforts are clearly sufficient to keep the boat afloat; thus, no large negative consequences will result if you refuse to bail. Nevertheless it seems obvious that you should help bail water. Intuitively, it would be unfair to let others do all the work.¹⁹⁹

In this scenario, your fellow shipmates make a good-faith sacrifice (in the form of physical exertion) in support of a mutually beneficial cooperative venture (the venture of bailing out the water that threatens to sink the boat). As a result, you receive the great benefit of not drowning.

¹⁹⁸ Simmons (2001: 29). The classic statements of this principle are offered by Hart (1955) and Rawls (1964).

¹⁹⁹ Huemer (2013: 87).

The principle of fair play therefore implies that you have a moral obligation to make similar sacrifices to support the cooperative venture. As Huemer notes, this seems intuitively correct.

Despite the plausibility of the basic idea of the principle of fair play, everyone agrees that the wholly general version of the principle articulated above is subject to many counter-examples. Philosophers who wish to make use of the principle of fair play seek to avoid counter-examples by carefully restricting the range of cases to which it applies. Here is the duly-restricted version of the principle of fair play that I accept:

Restricted Principle of Fair Play: If a group of people are all engaging in a certain type of activity, and if in doing so they collectively produce a public good G, and if a person P is benefited by G and could not receive G independently from the activity of this group of people, and if P would judge that receiving G is worth the costs of bearing a fair share of the burdens of the production of G, then P has a strong *prima facie* moral duty (a ‘duty of fair play’) to help contribute to this group’s production of G.

By a *public good*, I mean a good that is such that it is impossible or very costly for those who contributed to its production to prevent non-contributors from enjoying it. For example, clean air is a public good. If a group of environmentalists succeed in dramatically reducing air pollution, thereby producing the benefit of clean air, it is not feasible for them to prevent non-contributors from enjoying this benefit.

The restrictions incorporated into the above version of the principle of fair play enable it to avoid the counter-examples raised against it by the principle’s critics. Consider three important examples. First, Robert Nozick famously argued that a wholly-general version of the principle of fair play has the implausible implication that others can foist on you a duty to assist in the operation of a ‘public address system’ just by giving you some entertainment from that system that you cannot easily avoid, even though you do not judge the entertainment to be worth the costs of

assisting in the public address system's operation.²⁰⁰ The above version of the principle of fair play does not have this implausible implication, for it applies only when the beneficiary judges the benefits she receives from the collective efforts of the group to be worth the costs of bearing a fair share of the burdens of their production.

Second, Richard Arneson argues that an overly-simple version of the principle of fair play has the implausible implication that a neighborhood gift-giving association could make you obligated to give gifts to others just by giving you valuable gifts for which you never asked.²⁰¹ My favored version of the principle avoids this worry, for the benefit that the gift-giving association confers on you is not a public good: the gift-giving association could very easily refrain from conferring this benefit on you.

Third, A. John Simmons has offered the following thought experiment as a counter-example to 'non-voluntarist' versions of the principle of fair play, or versions that do not restrict its application to cases in which the beneficiary 'voluntarily accepts' the benefits of the cooperative activities of the group:

Suppose there is a severe drought in my rural neighborhood, where we are all dependent for water on our wells, wells that are now drying up. I am hard at work, successfully digging a new, much deeper well in my backyard to supply my family. But my neighbors, instead of doing the same, opt to dig a long trench along our neighborhood road and beyond, diverting water from a river several miles away, so that all will have access to running fresh water in front of their homes. If I decline to participate in my neighbors' scheme, have I breached an obligation of fair play by benefitting as a free rider? It seems plain to me that I have not.²⁰²

Although my version of the principle of fair play is non-voluntarist, it does not have the implausible implication that you are morally obligated to participate in the cooperative scheme

²⁰⁰ Nozick (1974: 93).

²⁰¹ Arneson (1982: 618).

²⁰² Simmons (2001: 34).

depicted in this thought experiment. For my principle requires the beneficiary to receive a public good that is such that she could not receive it independently from the activity of the cooperative activities of the group. In Simmons' trench-building scenario, you do have the means to obtain the relevant benefit—a source of water—without relying on the cooperative activities of your neighbors.²⁰³

The intuitiveness of the basic idea of the principle of fair play, along with my version's ability to avoid counter-examples, provides us with good grounds for accepting this version of the principle. I can now state my fair play account:

Fair Play Account

Many people in our society are members of the community of autonomous moral inquirers. In engaging in the activity of autonomous moral inquiry, these people collectively produce certain public goods that are such that all of us (i) receive these public goods, (ii) cannot receive these public goods independently of the activities of the community of autonomous moral inquirers, and (iii) would judge that receiving these public goods is worth the costs of bearing a fair share of the burdens of their production. Since this is so, and since the Restricted Principle of Fair Play is true, all of us have a strong *prima facie* duty of fair play to help contribute to the community of autonomous moral inquirers' production of this public good. Since these public goods are the effects of a lot of people engaging in autonomous moral inquiry, the way to contribute to the production of these goods is to engage in autonomous moral inquiry ourselves. So, we have a strong *prima facie* duty of fair play to think moral matters through ourselves and to form our moral beliefs non-deferentially. Moral deferrers are in violation of their strong *prima facie* moral duty to be autonomous moral inquirers. People's discomfort towards pure moral deference is sometimes caused by their (perhaps implicit) recognition of this fact.

This account is quite plausible in light of the foregoing discussion of the social benefits of autonomous moral inquiry. As I have argued, the community of autonomous moral inquirers collectively produce the following goods that we all enjoy: (a) a greatly-increased stock of justified moral beliefs and moral knowledge, (b) a greatly-increased stock of moral understanding, (c) an enhanced ability to perform morally worthy actions, (d) some protection from rights violations,

²⁰³ My strategy for avoiding Simmons' challenge is borrowed from Klosko (2005).

and (e) a rich variety of intellectual and aesthetic enjoyments. (a)-(e) are public goods, for it would be very costly for the members of the community of autonomous moral inquirers to prevent non-contributors from enjoying them. It is quite difficult to see how one who wished to receive (a)-(e) without the help of the community of autonomous moral inquirers could go about doing so. And any reasonable person would judge that (a)-(e) are sufficiently beneficial to be worth the costs of bearing a fair share of the burdens of their production—these costs consisting of the time and mental exertions involved in engaging in autonomous moral inquiry. So, it is reasonable to accept the Fair Play Account's central claim that pure moral deferrers are in violation of a strong *prima facie* duty of fair play to engage in autonomous moral inquiry. This provides additional explanatory resources to the realist in her project of explaining PD.

3.5 Conclusion

In summary, the Rebutting Puzzle of Pure Moral Deference that was first introduced by Sarah McGrath fails, for (i) no plausible anti-realist account of PD has ever been developed, and (ii) even bracketing this problem, the Rebutting Puzzle faces a dilemma that shows that at least one of its premises is false. And the Undercutting Puzzle of Pure Moral Deference fails because *contra* McGrath, realists can provide an adequate explanation of people's discomfort towards pure moral deference. I conclude that moral realism is secure from the challenges raised by the recent debate about moral deference.

Bibliography

- Alston, William (1988). "The Deontological Conception of Epistemic Justification," *Philosophical Perspectives* 2: 257-299.
- Arneson, Richard (1982). "The Principle of Fairness and Free-Rider Problems," *Ethics* 92: 616-633.
- Ayer, A.J. (1952). *Language, Truth and Logic*. New York: Dover Publications, Inc.
- Barber, Alex (2013). "Science's Immunity to Moral Refutation," *Australasian Journal of Philosophy* 91: 633-653.
- Blackburn, Simon (1993). *Essays in Quasi-Realism*. Oxford: Oxford University Press.
- Booth, Anthony (2006). "Can There be Epistemic Reasons for Action?" *Grazer Philosophische Studien* 73: 133-144.
- Boyd, Richard (1988). "How to Be a Moral Realist," *Essays on Moral Realism*, ed. G. Sayre-McCord. Ithaca, NY: Cornell University Press. 181-228.
- Brink, David (1984). "Moral Realism and the Sceptical Arguments From Disagreement and Queerness," *Australasian Journal of Philosophy* 62: 111-125.
- Brink, David (1989). *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- Brink, David (1992). "A Puzzle About the Rational Authority of Morality," *Philosophical Perspectives* 6: 1-26.
- Broome, John (2013). *Rationality Through Reasoning*. Oxford: Wiley Blackwell.
- Brunero, John (forthcoming). "Recent Work on Internal and External Reasons," *American Philosophical Quarterly*.
- Chrisman, Matthew (2016). "Epistemic Normativity and Cognitive Agency," *Noûs* 51: 1-21.
- Cline, Brendan (forthcoming). "The Tale of a Moderate Normative Skeptic," *Philosophical Studies*.
- Cowie, Christopher (2016). "Good News for Moral Error Theorists: A Master Argument Against Companions in Guilt Strategies," *Australasian Journal of Philosophy* 94: 115-130.
- Cuneo, Terence (2007). *The Normative Web: An Argument for Moral Realism*. Oxford: Oxford University Press.

- Darwall, Stephen (2003). "How Should Ethics Relate to (the Rest of) Philosophy?: Moore's Legacy," *Southern Journal of Philosophy* 41: 1-20.
- Darwall, Stephen, Allan Gibbard and Peter Railton (1992). "Toward Fin de Siecle Ethics: Some Trends," *Philosophical Review* 101: 115-189.
- Davia, Cory and Michele Palmira (2015). "Moral Deference and Deference to an Epistemic Peer," *Philosophical Quarterly* 65: 605-625.
- Dreier, James (1990). "Internalism and Speaker Relativism," *Ethics* 101: 6-26.
- Enoch, David (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Enoch, David (2014). "A Defense of Moral Deference," *Journal of Philosophy* 111: 229-258.
- Feinman, Jay. (2014). *Law 101: Everything You Need to Know About American Law*, Fourth Edition. Oxford: Oxford University Press.
- Feldman, Richard (2002). "Epistemological Duties," *The Oxford Handbook of Epistemology*, ed. Paul K. Moser. Oxford: Oxford University Press. 362-384.
- Finlay, Stephen and Mark Schroeder (2012). "Reasons for Action: Internal vs. External," *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, URL = <<https://plato.stanford.edu/entries/reasons-internal-external/>>.
- Fletcher, Guy (2016). "Moral Testimony: Once More With Feeling," *Oxford Studies in Metaethics*, Volume 11, ed. Russ Shafer-Landau. Oxford: Oxford University Press. 45-73.
- Gert, Bernard and Joshua Gert (2016). "The Definition of Morality," *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, URL = <<https://plato.stanford.edu/entries/morality-definition/>>.
- Gibbard, Allan (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge: Harvard University Press.
- Gibbard, Allan (2003). *Thinking How to Live*. Cambridge: Harvard University Press.
- Grimm, Stephen (2009). "Epistemic Normativity," *Epistemic Value*, eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 243-264.
- Groll, Daniel and Jason Decker (2014). "Moral Testimony: One of These Things is Just Like the Others," *Analytic Philosophy* 55: 54-74.
- Hall, Richard and Charles Johnson (1998). "The Epistemic Duty to Seek More Evidence," *American Philosophical Quarterly* 35: 129-139.

- Hart, H.L.A (1955). "Are There Any Natural Rights?" *Philosophical Review* 64: 175-191.
- Harman, Gilbert (1975). "Moral Relativism Defended," *Philosophical Review* 84: 3-22.
- Harman, Gilbert (2000). *Explaining Value and Other Essays in Moral Philosophy*. Oxford: Oxford University Press.
- Hazlett, Allan (2013). *A Luxury of Understanding*. Oxford: Oxford University Press.
- Hazlett, Allan (2017). "Towards Social Accounts of Testimonial Asymmetries," *Noûs* 51: 49-73.
- Heathwood, Christopher (2009). "Moral and Epistemic Open-Question Arguments," *Philosophical Books* 50: 83-98.
- Heathwood, Christopher (2011). "Desire-Based Theories of Reasons, Pleasure, and Welfare," *Oxford Studies in Metaethics*, Volume 6, ed. Russ Shafer-Landau. Oxford: Oxford University Press.
- Heil, John (1983). "Doxastic Agency," *Philosophical Studies* 43: 355-364.
- Hills, Alison (2010). *The Beloved Self: Morality and the Challenge From Egoism*. Oxford: Oxford University Press.
- Hills, Alison (2011). "Moral Epistemology," *New Waves in Metaethics*, ed. Michael Brady. New York: Palgrave-Macmillan. 249-263.
- Hills, Alison (2013). "Moral Testimony," *Philosophy Compass* 8: 552-559.
- Horgan, Terry and Mark Timmons (2000). "Nondescriptivist Cognitivism: Framework for a New Metaethic," *Philosophical Papers* 29: 121-153.
- Howell, Robert. (2014). "Google Morals, Virtue, and the Asymmetry of Deference," *Noûs* 48: 389-415.
- Hubin, Donald (2003). "Desires, Whims, and Values," *The Journal of Ethics* 7: 315-335.
- Huemer, Michael (2005). *Ethical Intuitionism*. New York: Palgrave Macmillan.
- Huemer, Michael (2013). *The Problem of Political Authority: An Examination of the Right to Coerce and the Duty to Obey*. New York: Palgrave Macmillan.
- Hyun, Alexander and Eric Sampson (2014). "On Believing the Error Theory," *Journal of Philosophy* 111: 631-640.

- Ingram, Stephen (forthcoming). "Epistemology Schmepeistemology: Moral Error Theory and Epistemic Expressivism," *Inquiry*.
- Johnson, Robert (1999). "Internal Reasons and the Conditional Fallacy," *Philosophical Quarterly* 49: 53-71.
- Joyce, Richard (2001). *The Myth of Morality*. Cambridge: Cambridge University Press.
- Joyce, Richard (2006). *The Evolution of Morality*. Cambridge: MIT Press.
- Joyce, Richard (2015). "Mackie's Arguments for the Moral Error Theory," *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, URL = <<https://plato.stanford.edu/entries/moral-anti-realism/moral-error-theory.html>>.
- Kelly, Thomas (2003). "Epistemic Rationality as Instrumental Rationality: A Critique," *Philosophy and Phenomenological Research* 66: 612-640.
- Klosko, George (2005). *Political Obligation*. Oxford: Oxford University Press.
- Kornblith, Hilary (1983). "Justified Belief and Epistemically Responsible Action," *Philosophical Review* 92: 33-48.
- Korsgaard, Christine (1996). *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Kramer, Matthew (2009). *Moral Realism as a Moral Doctrine*. Oxford: Blackwell.
- Kurth, Charlie (2013). "What Do Our Critical Practices Say About the Nature of Morality?" *Philosophical Studies* 166: 45-64.
- Loeb, Don (2007). "The Argument from Moral Experience," *Ethical Theory and Moral Practice* 10: 469-484.
- Lord, Errol (forthcoming). "How to Learn about Aesthetics and Morality Through Acquaintance and Testimony," *Oxford Studies in Metaethics*, Volume 12, ed. Russ Shafer-Landau. Oxford: Oxford University Press.
- Mackie, John (1977). *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- Manne, Kate (2014). "Internalism About Reasons: Sad But True?" *Philosophical Studies* 167: 89-117.
- Markovits, Julia (2011). "Why Be an Internalist About Reasons?" *Oxford Studies in Metaethics*, Volume 6, ed. Russ Shafer-Landau. Oxford: Oxford University Press. 255-279.
- Markovits, Julia (2012). "Acting for the Right Reasons," *Philosophical Review* 119: 201-242.

- Markovits, Julia (2012). "Saints, Heroes, Sages, and Villains," *Philosophical Studies* 158: 289-311.
- Markovits, Julia (2014). *Moral Reason*. Oxford: Oxford University Press.
- Markovits, Julia (2016a). "Precis of Moral Reason," *Philosophy and Phenomenological Research* 92: 518-530.
- Markovits, Julia (2016b). "Reply to Sobel and Kearns," *Philosophy and Phenomenological Research* 92: 549-559.
- Marquis, Don (1989). "Why Abortion is Immoral," *Journal of Philosophy* 86: 183-202.
- McGrath, Sarah (2009). "The Puzzle of Pure Moral Deference," *Philosophical Perspectives* 23: 321-344.
- McGrath, Sarah (2011). "Skepticism about Moral Expertise as a Puzzle for Moral Realism," *Journal of Philosophy* 108: 111-137.
- McPherson, Tristram (2012). "Ethical Non-Naturalism and the Metaphysics of Supervenience," *Oxford Studies in Metaethics*, Volume 7, ed. Russ Shafer-Landau. Oxford: Oxford University Press.
- McPherson, Tristram (2015). 'Supervenience in Ethics,' *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, URL = [<https://plato.stanford.edu/entries/supervenience-ethics/>](https://plato.stanford.edu/entries/supervenience-ethics/).
- Mill, John (1978). *On Liberty*. Indianapolis: Hacking Publishing Company.
- Millgram, Elijah (1996). "Williams' Argument Against External Reasons," *Noûs* 30: 197-220.
- Mogensen, Andreas (2015). "Moral Testimony Pessimism and the Uncertain Value of Authenticity," *Philosophy and the Phenomenological Research*. DOI: 10.1111/phpr.12255.
- Moore, G.E. (2005). *Principia Ethica*. New York: Barnes and Noble Publishing, Inc.
- Moore, G.E. (1939). "Proof of an External World," *Proceedings of the British Academy* 25: 273-300.
- Nozick, Robert (1974). *Anarchy, State, and Utopia*. New York: Basic Book, Inc.
- Olson, Jonas (2011). "In Defense of Moral Error Theory," *New Waves in Metaethics*, ed. Michael Brady. New York: Palgrave-Macmillan.

- Olson, Jonas (2014). *Moral Error Theory: History, Critique, Defence*. Oxford: Oxford University Press.
- Parfit, Derek (2011). *On What Matters*. New York: Oxford University Press.
- Portmore, Douglas (2011). "Consequentialism and Moral Rationalism," *Oxford Studies in Normative Ethics*, ed. Mark Timmons. Oxford: Oxford University Press.
- Railton, Peter (1986). "Moral Realism," *Philosophical Review* 95: 163-207.
- Rawls, John (1964). "Legal Obligation and the Duty of Fair Play," *Law and Philosophy*, ed. Sidney Hook. New York: New York University Press.
- Rowland, Richard (2013). "Moral Error Theory and the Argument from Epistemic Reasons," *Journal of Ethics and Social Philosophy* 7: 1-24.
- Rowland, Richard (2015). "Rescuing Companions in Guilt Arguments," *The Philosophical Quarterly* 66: 161-171.
- Scanlon, T.M. (2014). *Being Realistic About Reasons*. Oxford: Oxford University Press.
- Schroeder, Mark (2007). *Slaves of the Passions*. Oxford: Oxford University Press.
- Shafer-Landau, Russ (2003). *Moral Realism: A Defence*. Oxford: Oxford University Press.
- Shafer-Landau, Russ (2006). "Ethics as Philosophy: A Defense of Ethical Nonnaturalism," *Metaethics After Moore*, eds. Terry Horgan and Mark Timmons. Oxford: Oxford University Press. 54-62.
- Shafer-Landau, Russ. (2009). "A Defence of Categorical Reasons," *Proceedings of the Aristotelian Society* 109: 189-206.
- Sinclair, Neil (2012). "Moral Realism, Face-Values and Presumptions," *Analytic Philosophy* 53: 158-179.
- Simmons, A. John (2001). *Justification and Legitimacy: Essays on Rights and Obligations*. Cambridge: Cambridge University Press.
- Skorupski, John (1999). *Ethical Explorations*. Oxford: Oxford University Press.
- Sliwa, Paulina (2012). "In Defense of Moral Testimony," *Philosophical Studies* 158:175-195.
- Smith, Michael (1994). *The Moral Problem*. Oxford: Basil Blackwell.
- Smith, Michael (1995). "Internal Reasons," *Philosophy and Phenomenological Research* 55: 109-131.

- Stratton-Lake, Phillip (2002). "Introduction," *The Right and the Good*, ed. David Ross. Oxford: Clarendon. xxv-xxvi.
- Streumer, Bart (2008). "Are There Irreducibly Normative Properties?" *Australasian Journal of Philosophy* 86: 537-561.
- Streumer, Bart (2011). "Are Normative Properties Descriptive Properties?" *Philosophical Studies* 154: 325-348.
- Streumer, Bart (2013a). "Do Normative Judgements Aim to Represent the World?" *Ratio* 26: 450-470.
- Streumer, Bart (2013b). "Can We Believe the Error Theory?" *Journal of Philosophy* 110: 194-212.
- Streumer, Bart (forthcoming). *Unbelievable Errors*. Accessed on 7-21-2016.
- Suikkanen, Jussi (2011). "The Possibility of Love Independent Reasons," *Essays in Philosophy* 12: 32-54.
- Talbot, Brian (2014). "Truth Promoting Non-Evidential Reasons for Belief," *Philosophical Studies* 168: 599-618.
- Timmons, Mark (1999). *Morality without Foundations: A Defense of Ethical Contextualism*. Oxford: Oxford University Press.
- Tubert, Ariela (2016). "Sound Advice and Internal Reasons," *Pacific Philosophical Quarterly* 97: 181-199.
- Velleman, David (2000). *The Possibility of Practical Reason*. Oxford: Oxford University Press.
- Vogelstein, Eric (2013). "Moral Normativity," *Philosophical Studies* 165: 1083-1095.
- Wielenberg, Erik (2014). *Robust Ethics: The Metaphysics and Epistemology of Godless Normative Realism*. Oxford: Oxford University Press.
- Williams, Bernard (1979). "Internal and External Reasons," *Moral Luck*. Cambridge: Cambridge University Press, 1981: 101-113.
- Williams, Bernard (1989). "Internal Reasons and the Obscurity of Blame," *Making Sense of Humanity, and other philosophical papers*. Cambridge: Cambridge University Press, 1995: 35-45.
- Wolterstorff, Nicholas (2005). "Obligation, Entitlement, and Rationlity," *Contemporary Debates in Epistemology*, eds. M. Steup and E. Sosa. Oxford: Blackwell Publishing.