

Knowledge and the Many Norms on Action

Dissertation

Presented in Partial Fulfillment of the Requirements for the Degree Doctor of  
Philosophy in the Graduate School of The Ohio State University

By

James Fritz

Graduate Program in Philosophy

The Ohio State University

2019

Dissertation Committee

Tristram McPherson, Advisor

Declan Smithies

Justin D'Arms

Copyrighted by  
James Christopher Fritz  
2019

## Abstract

This dissertation develops and defends a view of the relationship between knowledge and norms on action. The view I defend is an *impurist* one; on this view, whether a person knows that *p* can depend on factors that are unrelated to the truth or likelihood of *p*. On my impurism, normative facts about actions and options—including, for instance, the *costs of relying on a belief in action*—are among the non-truth-related factors that can make a difference to knowledge. My view is distinctive, in part, because of the further claim that *moral* facts about actions and options are among the normative considerations that can make a difference to knowledge. In a slogan: epistemic norms are sensitive to moral considerations. In an even shorter slogan: there is *moral encroachment* in epistemology.

The first half of the dissertation defends the claim that there is moral encroachment in epistemology. In chapter 1 (“Pragmatic Encroachment and Moral Encroachment”), I argue that moral encroachment is at least as well-motivated as a more familiar view: the view that there is *pragmatic* encroachment in epistemology. In chapter 2 (“Uncertainty, Belief, and Ethical Weight”), I draw on insights from moral psychology to provide an original argument for moral encroachment.

The second half of the dissertation starts from the assumption that there is moral encroachment in epistemology. It is devoted to describing in detail how moral encroachment works. In chapter 3 (“Moral Encroachment and Reasons of the Wrong Kind”), I distinguish between a radical and a moderate version of moral encroachment. I raise a problem for the radical version: it threatens to erase the distinction between epistemic and non-epistemic norms on belief. This problem, I note, does not afflict the moderate version of moral encroachment. In chapter 4 (“Knowledge and the Many Norms on Action”), I show that it is no trivial task to explain why knowledge is sensitive to both *moral* norms and norms of *practical coherence*. In fact, I argue, existing versions of impurism generally lack the resources to provide the needed explanation. I also make some suggestions about the best ways to for an impurist to meet this challenge.

It’s clear that impurists both can and should add texture to their picture of norms on action. They can make significant headway by endorsing two of my primarily conclusions. First: there are multiple distinct families of normative considerations (including, for instance, *moral* considerations and considerations of *practical coherence*) that are equally well-placed to make a difference to knowledge. And second: only *some* of the moral considerations that bear on belief are apt to make a difference to knowledge. These conclusions point the way toward a clearer picture of an impurism worth defending.

Dedication

For Allie

## Acknowledgments

Chapter 1 of this dissertation, “Pragmatic Encroachment and Moral Encroachment,” initially appeared in *Pacific Philosophical Quarterly*. It appears with the permission of John Wiley and Sons (license number 4598230513436). I thank the publisher for this permission.

I owe an enormous debt of gratitude to the faculty and graduate students of the Department of Philosophy of the Ohio State University. Nearly every faculty member and student who was a member of the department while I attended the university helped me to become a better philosopher, to refine this dissertation, or both.

For helpful conversation, I also want to thank a number of philosophers outside the Ohio State University, including: Rima Basu, Michael Hannon, Elizabeth Jackson, Brian Kim, Matthew McGrath, Sarah Moss, Nathaniel Sharadin, Keshav Singh, Michael Pace, and Alex Worsnip.

I’m extremely grateful to have had the support of David Enoch and David Plunkett. Without their help, I would have missed irreplaceable opportunities to improve myself as a philosopher. I’m also extremely grateful to have worked with Julia Jorati as both a student of philosophy and as a teacher-in-training.

I can't successfully put into words how thankful I am to have worked with my wonderful dissertation committee: Justin D'Arms, Declan Smithies, and my advisor, Tristram McPherson. These are kind, supportive, and insightful philosophers, and they have taken every opportunity to support my philosophical and professional development.

Finally, I want to acknowledge the sacrifice and the love from my family that made this dissertation possible. To my siblings and siblings-in-law, to Connie and Chris Holoman, to Shannon and Chris Fritz, and to Allie: I love you. Thank you.

## Vita

2008	B.A., Philosophy and English, University of North Carolina
2010	M.A.T., Secondary English Education, American University
2013-present	Graduate Teaching Associate, Department of Philosophy, The Ohio State University

## Publications

“Pragmatic Encroachment and Moral Encroachment.” 2017. *Pacific Philosophical Quarterly* vol. 98, issue 51, pp. 643-661.

“What Pessimism about Moral Deference Means for Disagreement.” 2018. *Ethical Theory and Moral Practice* vol. 21, issue 1, pp. 121-136.

“Conciliationism and Moral Spinelessness.” 2018. *Episteme* vol. 15, issue 1, pp. 101-118.

“Moral Steadfastness and Metaethics,” with Tristram McPherson. 2019. *American Philosophical Quarterly* vol. 56, issue 1, pp. 43-56.

## Fields of Study

Major Field: Philosophy

Subfields: Metaethics, Epistemology



## Table of Contents

Abstract .....	iii
Dedication .....	v
Acknowledgments .....	vi
Vita .....	viii
Introduction .....	1
Chapter 1: Pragmatic Encroachment and Moral Encroachment .....	4
Chapter 2: Uncertainty, Belief, and Ethical Weight .....	32
Chapter 3: Moral Encroachment and Reasons of the Wrong Kind .....	62
Chapter 4: Knowledge and the Many Norms on Action .....	98
Bibliography .....	131

## Introduction

Suppose that a demon comes to you in the middle of the night. He credibly promises to give \$1 million to a good charity, on the condition that you form the belief that the number of stars in the universe is even. The demon has put you into a predicament. It certainly would be fortunate, morally speaking, if you somehow came to earn the demon's bribe—that is, if you somehow came to believe that the number of stars in the universe is even. Nevertheless, if you form that belief, you seem to be violating an important norm on belief-formation; after all, you don't have any evidence that favors the conclusion that the stars are even in number over the conclusion that the stars are odd in number.

Cases like the ones above provide evidence that there is a distinction to be drawn between different sorts of norms on belief. We can evaluate beliefs for how *desirable* they are, and for whether they are *good to have*. But we can also evaluate beliefs in a narrower way. We can evaluate them, for instance, with an eye solely to whether they are *rational*, whether they are *based on sufficient evidence*, or whether they amount to *knowledge*. When we do so, we generally ignore some good-making features of beliefs—including, for instance, whether those beliefs would have good

consequences (like greater funding for a charity). Call norms of this latter sort *epistemic norms*.

How, precisely, should we understand the distinction between epistemic and non-epistemic norms on belief? One initially tempting approach is straightforward: epistemic norms are sensitive only to considerations that have to do with the truth or likelihood of belief, and therefore, are never sensitive to practical or moral considerations.

Despite the initial appeal of this simple approach, practical and moral considerations have been creeping back into epistemology. Recent years have seen several defenses of the view that there is *pragmatic encroachment* in epistemology: the view that knowledge of *p* can depend on practical facts—including, even, practical facts that do not bear on the truth or likelihood of *p*. Some defenses of pragmatic encroachment, for instance, suggest that the *costs of error* about *p* for an agent can make a difference to whether that agent has knowledge.

Though pragmatic encroachment has enjoyed a great deal of discussion over the last ten years, the notion that there is *moral encroachment* in epistemology has been largely neglected. If there is moral encroachment in epistemology, then knowledge of *p* can depend on moral facts—including, even, moral facts that do not bear on the truth or likelihood of *p*. Consider an example: perhaps, in some cases, whether a person *knows* that a train will be leaving the station on time depends on how *morally* important it is that she arrive at her destination on time.

My goal in this dissertation is to develop and defend a view of epistemology that involves moral encroachment. I'll leave it to the chapters that follow to bring out the contours of the view. This brief introduction serves to orient the reader to the topic at hand, and to the discussion to follow.

It may be helpful to think of the dissertation as divided into two halves. The first half, comprised of Chapters 1 and 2, defends the view that there is moral encroachment in epistemology. Chapter 1 does this work by arguing that moral encroachment is at least as well-motivated as a more familiar thesis: pragmatic encroachment. Chapter 2 draws on insights from moral psychology to provide an original argument in favor of moral encroachment.

The second half of the dissertation, comprised of Chapters 3 and 4, takes it for granted that there is moral encroachment; its goal is to illuminate just how moral encroachment works. Chapter 3 draws conclusions regarding *which* sorts of moral considerations bear on knowledge. Chapter 4 sketches a way for moral encroachment to live peacefully alongside pragmatic encroachment in a unified, appealing picture of epistemic normativity.

## Chapter 1

### Pragmatic Encroachment and Moral Encroachment

This dissertation is concerned with a recently popular movement in epistemology, one which sometimes goes by the name ‘impurism.’<sup>1</sup> On impurist views, whether a person has knowledge depends on features of her environment that are not paradigmatically epistemic ones.

Which features of a belief, a believer, or her environment are ‘paradigmatically epistemic?’ Certain examples spring to mind: for instance, truth, degree of likelihood, and safety from error. As Ichikawa and Steup (2016) note, these are factors that seem to have “some sort of intimate connection with truth.” I won’t, in this dissertation, attempt to spell out precisely what that intimate connection is. My arguments require only the assumption that two sorts of facts are not paradigmatically epistemic: facts about what matters *to a person* (like the fact that *it would be in Patricia’s interest to open a 401k*), and moral facts (like the fact *that it is morally permissible for Patricia to give money to charity*). This is a weak assumption. Traditional theories of knowledge

---

<sup>1</sup> I take this term from Fantl and McGrath (2009, 35). See Ichikawa and Steup (2016, section 11) for a notion of “subject-sensitivity” that is similar in some respects to impurism, and for a useful contrast with Jason Stanley’s related claim of interest-relativity.

do not posit systematic connections between knowledge and practical interests or systematic connections between knowledge and morality.

It's worth noting a complication: *some* moral facts, or facts about what matters to an agent, provide evidence for belief. (For instance, perhaps the fact that *Rex was cruel to Willa* provides evidence that *Rex is acquainted with Willa*.) All parties to the debate over impurism should acknowledge that, by playing a role of this sort, practical and moral facts can serve as paradigmatically epistemic factors. More precisely, then, the weak assumption that informs this dissertation is the following: when facts about practical interests and moral facts do not bear on whether *p*, they are not among the paradigmatically epistemic factors relevant to whether a belief in *p* is knowledge. If this weak assumption is right, then one way to be an impurist is to make the following claim: whether a person knows that *p* can vary with practical and moral features of her environment—including ones that do not bear on whether *p* is true.

This chapter investigates a crucial question for impurism: *which* of the many factors that are not paradigmatically epistemic can make a difference to knowledge? The most prominent early defenses of impurism answer this question by citing facts about what matters *to the believer*: in order to bear on *S*'s knowledge, they suggest, a non-paradigmatically-epistemic feature of her environment must be connected to things that matter *to S* in some way.<sup>2</sup> Call this the thesis of pragmatic encroachment. By

---

<sup>2</sup> See, e.g., Fantl & McGrath (2007, 559n1), Stanley and Hawthorne (2008, 583), and Stanley (2005, 92-3).

contrast, defenders of impurism have tended to ignore a more radical hypothesis: features that matter *morally*, whether or not those features matter to the subject, can affect what a subject knows. Call this the thesis of moral encroachment. In a slogan, then, most prominent defenders of impurism accept *pragmatic* encroachment on knowledge without paying attention to *moral* encroachment on knowledge.<sup>3</sup>

This chapter argues for a striking conclusion about impurism: to the extent that arguments for impurism are successful in showing that knowledge is sensitive to practical norms, those arguments are equally successful in showing that knowledge is sensitive to moral norms. Insofar as there are good arguments for pragmatic encroachment, there are also good arguments for moral encroachment.

### **Section 1: Practical Normativity and Moral Normativity**

This chapter focuses on a contrast between two ways in which impurism could be defended. First, it might be (and usually is) defended through an appeal to pragmatic encroachment. On this first approach, knowledge depends on facts that are not paradigmatically epistemic facts precisely because it depends on facts about what matters *to the subject*. Second, it might be defended through an appeal to moral

---

<sup>3</sup> At the time of this chapter's initial publication, the only major defense of moral encroachment could be found in Pace (2011). It's perhaps worth noting that Pace's argument is significantly different from the argument of this chapter. While Pace argues for moral encroachment without making any appeal to the existing pragmatic encroachment literature, this chapter engages directly with that literature and reveals its connection to moral encroachment. Since the publication of this chapter in 2017, several other defenses of moral encroachment have appeared: see Moss (2018a), Schroeder (2018), Basu and Schroeder (2019), Basu (forthcoming), Bolinger (forthcoming).

encroachment. On this second approach, knowledge depends on facts that are not paradigmatically epistemic facts precisely because it depends on facts about what matters *morally speaking*.

I do not claim that the distinction between what matters *to the subject* and what matters *morally speaking* is exhaustive, precise, or maximally illuminating.<sup>4</sup> It's entirely possible that there are other, more productive ways of distinguishing between sorts of normativity that bear on a subject's action. The crucial point, for this chapter, is that extant arguments for impurism can be adapted to support the claim that knowledge is sensitive to what matters *morally* just as well as they can be used to support the claim that knowledge is sensitive to what matters *to the knower*. Though these arguments are generally used to defend a surprising connection

---

<sup>4</sup> Just what does it mean for something to matter *to a subject*? The arguments below are compatible with several different answers to this question. It's not totally clear whether the current defenders of impurism are more easily interpreted as focusing on very subjective practical facts (like the fact that some action would further a project of mine, regardless of whether that project's realization would be good for me) or more objective practical facts (like the fact that some outcome would be in my interests, regardless of whether or not I actually desire or aim at that outcome). I will not attempt to settle this interpretive question; either notion is compatible with the arguments presented here.

That said, I set aside here at the outset views that tie morality and practical normativity very closely. Consider, for instance, views on which the fact that it would be morally wrong for a person to  $\phi$  entails that it matters a great deal *to that person* that she not  $\phi$ . On views of this sort, variations in the moral facts will sometimes amount to changes in the practical facts. A defender of pragmatic encroachment who takes such a view, then, can explain all the cases below by appeal to practical facts; she can take the position that moral facts have no *distinctive* force in epistemology. But she nevertheless must grant one of the most *prima facie* surprising theses suggested in this chapter: epistemic facts are interestingly connected to moral facts.

That claim—that epistemic facts are interestingly connected to moral facts—is much less obvious, and a much more daring claim, on the (plausible) assumption that what matters to a subject can diverge dramatically from what matters morally. This chapter aims to show that, *even granting* significant separation between what matters to a subject and what matters morally, the arguments for pragmatic encroachment can be adapted to support moral encroachment. (Thanks to Tristram McPherson for useful discussion here.)



between knowledge and *practical* normativity, they can be adapted to defend an even more surprising connection between knowledge and *moral* normativity.

So just how do extant arguments for impurism work? The most prominent premise in these arguments posits a particular sort of connection between knowledge and action. Let's call this premise in the arguments for impurism the *knowledge-action link*. Consider the following examples:

If you know that  $p$ , then  $p$  is warranted enough to justify you in  $\phi$ -ing, for any  $\phi$ . (Fantl and McGrath 2009, 66)

If a subject knows that  $p$ , then she is in a good enough epistemic position to rely on  $p$  in her practical reasoning. (Brown 2008, 245)

Where one's choice is  $p$ -dependent, it is appropriate to treat the proposition that  $p$  as a reason for acting iff you know that  $p$ . (Hawthorne and Stanley 2008, 578)

It's worth noting that the final formulation of the knowledge-action link above, from Hawthorne and Stanley, is particularly strong; it posits a *biconditional* connection between knowledge and sufficient warrant for action. Though all knowledge-action links have their critics (e.g. Brown 2008), this biconditional version is particularly controversial.<sup>5</sup> For the purposes of this chapter, we can set it aside and focus on the weaker, less controversial version of the knowledge-action link: knowledge is a sufficient condition for epistemically unproblematic action. In what follows, I'll work with the following formulation:

---

<sup>5</sup> For criticism of the biconditional knowledge-action link, see Fantl and McGrath (2009, 124-5); and Smithies (2012, 269-70).

**Knowledge-Action Link:** If S knows that *p*, then S is warranted enough to act (believe, prefer) as if *p*.

Why would anyone accept even this weaker form of the knowledge-action link? I'll take a moment, now, to briefly mention two of the reasons for which defenders of impurism have found the knowledge-action link attractive.<sup>6</sup>

First, impurists often note that we defend and encourage action by talking about knowledge. When I judge that you are being too reticent, I can encourage you to act by saying, "There's no need to keep looking for more evidence. You know that it'll work" (cf. Fantl and McGrath 2007, 562-3). Likewise, if I want to defend my action from distinctively epistemic criticism, I can do so by insisting that I knew my action would work. This sort of practice suggests that people generally consider knowledge that *p* a sufficient condition on the appropriateness of acting as if *p*.

Impurists also appeal to the close connection between our judgments about knowledge and our judgments about negligence and blame. If you let your child play near a dog, and I accuse you of negligence, you can escape blame by showing me that you know the dog is safe. This pattern of behavior, again, suggests that people tacitly consider knowledge that *p* a sufficient condition on the appropriateness of acting as if *p* (Hawthorne and Stanley 2008, 572-3). Several related observations about

---

<sup>6</sup> Both of the considerations below can be, and have been, used to defend the biconditional knowledge-action link as well (Hawthorne and Stanley 2008, 571-3), but I'll be solely concerned with the way in which they support the claim that knowledge is a sufficient condition for epistemically unproblematic action.

everyday thought and talk, then, give the knowledge-action link some *prima facie* plausibility.

Now, note that there is a gap between the knowledge-action link and impurism: that is, the claim that whether S has knowledge depends on features of S's environment that are not paradigmatically epistemic. To reach that conclusion, the defenders of impurism must adopt an additional premise that connects a claim about sufficiently warranted action to features of S's environment that are not paradigmatically epistemic. Let's call that premise the *action-environment link*. Here is a rough schema for the typical arguments that use a knowledge-action link and an action-environment link to support impurism (Cf. Brown 2013, 245):

**Knowledge-Action Link:** If S knows that p, then S is warranted enough to act (believe, prefer) as if p.

**Action-Environment Link:** Feature F of S's environment (where F is not a paradigmatically epistemic feature of an environment) makes a difference as to whether S is warranted enough to act (believe, prefer) as if p.

Therefore, **Impurism:** Whether S knows that p depends on features of S's environment that are not paradigmatically epistemic features.

This chapter will make a comparison between two ways of filling out the action-environment link. One, which has been adopted by many prominent defenders of impurism, is the *practical* action-environment link.<sup>7</sup> Adopting this variant of the action-environment link results in the following argument for Impurism:

---

<sup>7</sup> Since most defenders of impurism do not make the distinction between pragmatic encroachment and moral encroachment, they are usually not explicit about their adherence to the former thesis rather than the latter. Nevertheless, the focus of their writing does consistently suggest a focus on matters of

**Knowledge-Action Link:** If S knows that p, then S is warranted enough to act (believe, prefer) as if p.

**Practical Action-Environment Link:** Facts about what matters *to S* (which are not paradigmatically epistemic features of an environment) make a difference as to whether S is warranted enough to act (believe, prefer) as if p.

Therefore, **Impurism:** Whether S knows that p depends on features of S's environment that are not paradigmatically epistemic features.

The other variant of the action-environment link that's relevant for this chapter is the *moral* action-environment link. It can be used, in much the same way, to support impurism:

**Knowledge-Action Link:** If S knows that p, then S is warranted enough to act (believe, prefer) as if p.

**Moral Action-Environment Link:** Facts about what matters *morally* (which are not paradigmatically epistemic features of an environment) make a difference as to whether S is warranted enough to act (believe, prefer) as if p.

Therefore, **Impurism:** Whether S knows that p depends on features of S's environment that are not paradigmatically epistemic features.

As these argument-schemas demonstrate, defenders of moral encroachment and pragmatic encroachment alike can accept the knowledge-action link. In fact, as we'll see in section 4, they can accept the knowledge-action link for just the same reasons. So arguments in favor of the knowledge-action link, though they occupy a great deal

---

importance *to the subject* rather than matters of objective importance; see, e.g., Fantl & McGrath (2007, 559n1), Stanley and Hawthorne (2008, 583), and Stanley (2005, 92-3).

of the current literature on impurism, do not favor pragmatic encroachment over moral encroachment.

To the contrary, the difference between current arguments for pragmatic encroachment and analogous arguments for moral encroachment lies most importantly in the difference between versions of the action-environment link. So this chapter will be primarily concerned with arguments for the practical action-environment link and arguments for the moral action-environment link. Section 2 and section 3 investigate two of the most prominent ways in which defenders of impurism have argued for the practical action-environment link. In each case, I argue that analogous lines of reasoning can be used, with just as much plausibility, to defend the moral action-environment link.

## **Section 2: The Argument from High-Stakes Gambles**

The most explicit support for the practical action-environment link in the current literature comes from Jeremy Fantl and Matthew McGrath. It is an argument about rational action and high-stakes gambles.

The argument from high-stakes gambles starts from the assumption that we know some propositions *fallibly*—that is, we sometimes know that  $p$  even though there is a non-zero epistemic chance for us that not- $p$  (Fantl and McGrath 2009, 11).<sup>8</sup> Suppose,

---

<sup>8</sup> This gloss relies on a notion of fallibilism that Fantl and McGrath (2009) call “Strong Fallibilism.” In their (2007), Fantl and McGrath appeal to a considerably weaker notion of fallibilism (559). On either

for example, that in everyday, low-stakes contexts, you know the proposition *that Plato taught Aristotle* fallibly. Given the knowledge-action link, this means that you are warranted enough to act as if Plato taught Aristotle. After all, if the knowledge-action link holds, any proposition that you know (fallibly or not) is warranted enough for you to act on it.

Now, suppose that you are presented with a high-stakes gamble: if it is true that Plato taught Aristotle, you get a piece of candy, but if it is false, you get tortured to death. Fantl and McGrath ask, 'Would it be rational for you to stake your life on the proposition that Plato taught Aristotle? It seems to us it would not' (2009, 13). This reaction seems plausible enough. Moreover, it seems plausible that this sort of phenomenon generalizes. As long as there is an epistemic chance, however slim, that things will go very badly for you, you can be rationally required to avoid that chance as long as the outcome would be sufficiently bad.

But if this is right, then the mere fact that you are offered a high-stakes gamble about  $p$  can change whether your belief that  $p$  is warranted enough to support rational action. We now have an argument for the practical action-environment link: whether you are warranted enough to act on your beliefs depends on the facts about things that matter to you.

---

reading of fallibilism, the argument from high-stakes gambles can be modified to yield the moral version of the action-environment link.

Here, briefly, is the argument from high-stakes gambles:

- (1) Consider some proposition  $p$  that you know fallibly.<sup>9</sup> In everyday, low-stakes cases, you are warranted enough to act as if  $p$ .
- (2) There is some high-stakes gamble as to whether  $p$  such that you would not be rational to take that high-stakes gamble.
- (3) (2) is true because, when presented with the high-stakes gamble, you are not epistemically warranted enough with respect to  $p$  to take the gamble.
- (4) The only difference between the high-stakes gamble scenario and everyday cases is a difference in the facts about things that matter *to you*—specifically, a difference in the value *to you* of possible outcomes of your available actions.<sup>10</sup>
- (5) So, **Practical Action-Environment Link**: Facts about what matters to  $S$  (which are not paradigmatically epistemic features of an environment) make a difference as to whether  $S$  is warranted enough to act (believe, prefer) as if  $p$ .

Fantl and McGrath do not call attention to (3), but it is a crucial step in this form of argument. It's worth pausing to see why.

In order to support a viable version of impurism, the action-environment link must be a claim about a particular sort of problem that gets in the way of rational action.

The problem must be a distinctively epistemic one: one's epistemic position with respect to  $p$  is not strong enough to warrant acting as if  $p$ . The defender of impurism

---

<sup>9</sup> Some infallibilists may resist the argument from high-stakes gambles in a fairly flatfooted way: we are not rational to take high-stakes gambles on most propositions precisely because there is (always) an epistemic chance for us that those propositions are not true, and that we therefore cannot know those propositions. This approach threatens to rule out an enormous amount of everyday knowledge. Other infallibilists may respond to the argument in a different way; they may incorporate impurism into their account of epistemic possibility. This amounts to the admission that whether or not there is an epistemic chance for me that  $p$  depends on features of my environment that are not paradigmatically epistemic features. Thanks to an anonymous referee for suggesting more detail about infallibilism, and to Declan Smithies for helpful discussion.

<sup>10</sup> This formulation is somewhat loose; I grant that, in all the cases to be discussed below, differences in practical and moral facts will always be accompanied by differences in the non-normative facts on which they supervene.

must take care to tie *only* this sort of rational defect in action to knowledge. Otherwise, she will run the risk of implying, implausibly, that no one can ever act irrationally on the basis of a known proposition. Surely, some people can act irrationally on the basis of knowledge. A person can, for instance, rely on known propositions to act in a way that foreseeably subverts his welfare—say, by buying a horribly addictive drug. The defender of impurism should say that such actions, though rationally defective, are not rationally defective in virtue of the weakness of the subject’s warrant for his beliefs. They are rationally defective because they are based on practically irrational aims.<sup>11</sup>

Fantl and McGrath, then, must be suggesting that the person who takes a (sufficiently) high-stakes gamble is irrational in a way that is distinctively connected to his *epistemic position*. By doing so, they can use (1)-(4) to support a version of the practical action-environment link that can interact in the right way with the knowledge-action link.

It will be useful to consider the way that this line of thinking applies to a more realistic case. Consider the following one, from Fantl and McGrath (2007):

... if I find out that the police are (for the first time ever) about to ticket illegally parked cars on my quiet, rural street, my stakes in whether my car is legally parked rise. This makes a difference to whether I’m rational to act as if my car is legally parked, even if my strength of

---

<sup>11</sup> On some views, aims can only be practically irrational in virtue of incoherence with other elements of the subject’s psychology; on others, aims can also be practically irrational in virtue of their failure to track objective facts about what matters to the subject. The argument in the main text is compatible with either explanation.



epistemic position doesn't change..., my belief doesn't change... , etc.  
(2007, 560)

The belief that my car is parked legally outside seems like a paradigm case for the sort of belief that, in everyday cases, I am warranted enough to take for granted. But once I am placed in a high-stakes scenario like the one that Fantl and McGrath sketch, my epistemic position is no longer strong enough to warrant me in acting as if my car is parked legally. Indeed, I might need to go check and make sure that the car is parked legally. The defender of impurism uses cases like these as evidence for the practical action-environment link: whether my beliefs are warranted enough to justify depends on the facts about what matters *to me* (in this case, the facts about the value to me of possible outcomes of my available actions).<sup>12</sup>

I'll now argue that Fantl and McGrath's line of reasoning can also be used to support an action-environment link that appeals to moral facts. Call the modified argument the *moral* argument from high-stakes gambles:

- (6) Consider some proposition *p* that you know fallibly. In everyday, low-stakes cases, you are warranted enough to act as if *p*.
- (7) There is some high-stakes situation in which there would be a moral problem with your acting as if *p*.
- (8) (7) is true because, in the high-stakes situation, you are not epistemically warranted enough to act as if *p*.
- (9) The only difference between the high-stakes situation and everyday cases is a difference in the moral facts—specifically, a difference in your moral obligations and permissions.
- (10) So, **Moral Action-Environment Link**: Facts about what matters *morally* (which are not paradigmatically epistemic features of an

---

<sup>12</sup> The argument from high-stakes gambles also applies neatly to the bank cases that pervade so much of the discussion about impurism (see especially Stanley 2005, 1-15).

environment) make a difference as to whether S is warranted enough to act (believe, prefer) as if p.

To see the appeal of the moral argument from high-stakes gambles, let's consider a modified version of the example about whether your car is parked legally. Suppose that you do not learn that the police are about to ticket illegally parked cars. You learn, instead, that there is a maniacal traffic officer making his way down your street, and that he is nearly at his wit's end when it comes to dealing with illegally parked cars. If he sees that your car is parked legally, he will calm down and spend the rest of the night in peace. But if he sees that your car is parked illegally, he will fly into a homicidal rage and kill five innocent people. Luckily, you have the time to go check and make sure that your car is parked legally. In such a case, it would be morally problematic if you failed to go check.<sup>13</sup>

This case illuminates the appeal of premises (6)-(9). It shows, in short, that when you are offered a sufficiently high-stakes *moral* gamble as to whether *p*, *p* can thereby become insufficiently warranted to justify you in acting. Imagine that you stay in your easy chair rather than going to check on your car. There would be something problematic about this behavior.<sup>14</sup> And the problem seems best explained in the

---

<sup>13</sup> Interestingly, Fantl and McGrath (2002, 85) describe a case remarkably similar to this one. Even more interestingly, they use the case in the course of arguing for an action-environment link. But they do not acknowledge (and might not accept) the striking implication of the thought experiment: knowledge can depend on moral facts in just the same way that it depends on facts about what matters to the subject.

<sup>14</sup> Just what would be morally problematic about your staying in your easy chair in the high-stakes case? One intuitively appealing answer is that it is a morally wrong action. But this answer is not available on all views about wrong action. Certain objective consequentialists, for instance, might say that you are right to stay in your easy chair because, as a matter of fact, this is the action that will have

following way: if you were to remain in your easy chair, you would be relying on a belief that was insufficiently warranted to make it the case that your action is morally unproblematic. But if this is right, then changing certain moral facts about your environment—including ones that do not bear on whether  $p$  is true—can change whether you are warranted enough to act as if  $p$ . In other words, the moral action-environment link is true.

It's important to acknowledge a wrinkle in this new line of reasoning. It's tempting to suppose that, even in this new example where lives hang in the balance, the facts about what matters *to the subject* can adequately explain the difference between permissible action in the ordinary case and in the high-moral-stakes case. After all, most of us care a great deal about protecting human lives where we can easily do so. Protecting others' lives, in other words, does not simply matter morally; it also usually matters *to us*. If this is right, then perhaps the problem with a person's warrant in the case offered above can be explained entirely through appeal to pragmatic encroachment.

To see why this approach won't work, consider a further modification of the case. Imagine a different subject whose car is parked outside. He does care a little bit about

---

the best consequences. Such consequentialists will be likely to claim that the moral problem in this scenario is a problem with you as an agent, not with your action (cf. Moore 2006). But on any plausible first-order normative theory, there will be *some* moral problem that arises when you stay in your easy chair precisely because of the weakness of your epistemic position. I say more about this in the main text below.

whether his car is parked legally, simply because he lightly prefers to follow rules. But he is radically apathetic about other human lives; it does not matter *to him* even a little bit whether any other person lives or dies. In an everyday case, this person can easily know that his car is parked legally. But what about a case where the maniacal traffic officer comes to this person's street? Well, in this situation, the facts about what matters to him are just the same as the facts about what matters to him in an ordinary, low-moral-stakes case. After all, others' lives simply do not matter to him. Nevertheless, there would certainly be a serious moral problem with his acting as if his car were parked legally—by, say, remaining seated in his easy chair. An appeal to merely *pragmatic* encroachment cannot explain this case in the same way that it explains high-stakes gambles like Fantl & McGrath's.

Just what is the problem with the apathetic man's failing to go check his car? Well, we might cite any of several problems. Perhaps, for instance, his inactivity reveals his repugnant character. But any satisfactory explanation of the problem with this man's action must also note an epistemic failing: his belief that his car is parked legally is not warranted enough. To see this, consider a final modification of the maniacal traffic officer case. In this modified version, the man, while retaining his repugnant character, has arbitrarily strong evidence for the belief that his car is parked legally. Perhaps, for instance, he is standing right next to his car, seeing that it is parked legally, and hearing several policemen say, "yep, that car sure is parked legally." In such a case, at least one moral problem with the man dissipates. Even if he does

continue to fall short morally in some respects, he does not fall short morally by calling off the search for further evidence that his car is parked legally. Imagine the person trying to defend his inactivity from moral criticism by saying, “Nothing I do will make a difference! My car is parked legally.” In this most recent case, the apathetic man’s epistemic position is strong enough that this is a legitimate excuse. Not so in the original case; if the man is sitting in his easy chair, with only a vague memory of the way he parked his car, he cannot morally excuse his action by appeal to his belief that the car is parked legally. We must explain the difference between these two cases by appealing to a distinctive sort of moral problem—one that arises in virtue of weakness in epistemic warrant. In the original case, the apathetic car owner’s belief that his car is parked legally is insufficiently warranted to morally excuse his inactivity.

This point bears emphasis, because without it, the thesis of moral encroachment might seem like an obvious non-starter. There are a multitude of ways in which a person’s action or reasoning might be morally problematic, and not all of those moral shortcomings are plausibly rooted in that person’s epistemic warrant for her beliefs. Imagine, for instance, that you attempt to poison your wealthy uncle in order to get his inheritance. You carry out your nefarious plan, all the while relying on fallible beliefs (like the belief *that this poison is strong enough to kill an adult man*). Your action, and reasoning, would be morally problematic. But this sort of moral problem does not lie in your degree of epistemic warrant for any of your beliefs. After all, even

if you attempted to poison your uncle on much stronger or much weaker evidence of your poison's strength, you would still be engaged in an evil project. This sort of moral problem, then, bears no obvious connection to the epistemic warrant for your beliefs. The problem with the morally apathetic car owner's action, by contrast, is a problem that apparently does have a strong connection to the epistemic warrant for his beliefs. It's only the latter sort of problem that provides support for the thesis of moral encroachment.

As a result, the argument on offer here does not have the implausible result that no one can ever act in a way that is morally problematic on the basis of sufficiently warranted premises. There is a nice symmetry between this conclusion and a conclusion about pragmatic encroachment mentioned earlier in the section. Above, I noted that the defender of pragmatic encroachment must accept that it's possible to act irrationally from known premises. She can get this result by distinguishing between actions that are irrational in virtue of an epistemic problem (like weakness of warrant) and actions that are irrational for other reasons (like practically irrational aims). In the same way, the defender of moral encroachment must accept that it's possible to act immorally from known premises. She can get this result by distinguishing between actions that are immoral in virtue of an epistemic problem (like weakness of warrant) and actions that are immoral for other reasons (like bad desires or values). Only certain cases of morally problematic action—like the

apathetic car owner's—betray a distinctively epistemic problem, and thereby provide support for the thesis of moral encroachment.<sup>15</sup>

We've now seen an argument, the moral argument from high-stakes gambles, that can support the moral action-environment link. What's more, it can do so in just the same way that Fantl and McGrath's argument from high-stakes gambles supports the practical action-environment link. Section 4 of this chapter will consider and reject an attempt to accept Fantl and McGrath's argument while rejecting the moral argument from high-stakes gambles. But before we consider that objection, let's turn to another argument for impurism that can be used to support the conclusion of moral encroachment.

### **Section 3: The Argument from Unacceptable Practical Reasoning**

Fantl and McGrath's argument from high-stakes gambles is a particularly explicit attempt to defend the action-environment link. But some defenders of impurism may

---

<sup>15</sup> An anonymous referee asks for more detail about just which environments are the ones in which moral facts affect warrant for belief. While I remain uncommitted to any particular story about the precise relationship between moral facts and epistemic warrant (and, indeed, uncommitted even to the thesis of moral encroachment), it will be useful to sketch one such story in this footnote. In doing so, I provide a sort of possibility proof, showing that there are viable pictures of moral encroachment that preserve the crucial analogy with viable pictures of pragmatic encroachment.

The defender of moral encroachment might claim that *S*'s epistemic position with respect to *p* is insufficient if the most subjectively morally choiceworthy action available to *S* is distinct from the most subjectively morally choiceworthy action available to *S* on the assumption that *p*. The analogous principle about pragmatic encroachment is: *S*'s epistemic position with respect to *p* is insufficient if the rationally best action available to *S* is distinct from the rationally best action available to *S* on the assumption that *p*.

Note that this principle makes no reference at all to the way that *S* would in fact act if *P* were the case (or if she believed that *p* were the case). So it is in no danger of falsely implying that the fact that a person will in fact act immorally if *p* (or if he believes that *p*) has any tendency to undermine his epistemic warrant for *p*.

have good reason to argue for that link along different lines. Hawthorne (2004) enumerates some benefits of adopting impurism without calling any particular attention to stakes, and Hawthorne and Stanley (2008, 588) worry outright that stakes-sensitivity may not adequately explain the connection between practical environment and knowledge. In light of this, we should consider a line of argument in support of the practical action-environment link that does not require an appeal to stakes. This section shows how such an argument can be extracted from the examples of unacceptable practical reasoning that appear in Hawthorne (2004) and Hawthorne and Stanley (2008).

Consider a contrast between two chains of practical reasoning that start from the same premise. First, suppose you are in a bookstore thinking about whether to buy a visitor's guide to Blackpool. You reason as follows:

I will be going to Blackpool next year.  
So I will be able to make good use of a Blackpool visitor's guide next year.  
So I ought to buy a Blackpool visitor's guide.

This seems like it could be a perfectly acceptable chain of practical reasoning.

Next, suppose that on your way out of the bookstore, you are offered life insurance. It would be unacceptable for you to reason thus:

I will be going to Blackpool next year.  
So I won't die beforehand.  
So I ought to wait until next year before buying life insurance.  
(Hawthorne 2004, 175).



Intuitively, the problem with this second bit of practical reasoning lies in the fact that you are insufficiently epistemic warranted in the first premise (Hawthorne 2004, 175-6). Your epistemic position is no longer strong enough for you to take it for granted that you will be going to Blackpool next year. But we can stipulate that nothing changes between the time of the first bit of reasoning and the second bit of reasoning except that you are offered life insurance. The fact that you are engaged in practical reasoning about whether to buy life insurance, then, has made a difference to whether you are warranted enough to rely on one of your beliefs in your practical reasoning.

Contrasts like the one above suggest an argument for an action-environment link that is quite similar to (1)-(5):

- (1') Consider some proposition  $p$  that you know fallibly. In certain practical environments, you are warranted enough to rely on  $p$  as a premise in practical reasoning.
- (2') But in other practical environments, it is rationally unacceptable for you to rely on  $p$  as a premise in practical reasoning.
- (3') (2') is true because, in the latter sort of practical environment, you are not epistemically warranted enough to rely on  $p$ .
- (4') The only difference between the former and latter bits of practical reasoning is a difference in the facts about things that matter *to you*—specifically, a difference in which of the things that matter to you are the topic of your current practical reasoning.
- (5') So, **Practical Action-Environment Link**: Facts about what matters to  $S$  (which are not paradigmatically epistemic features of an environment) make a difference as to whether  $S$  is warranted enough to rely on  $p$  as a premise in practical reasoning.

I'll now argue that an argument much like (1')-(5') can be used to support the moral action-environment link.

- (6') Consider some proposition  $p$  that you know fallibly. In certain practical environments, you are warranted enough to rely on  $p$  as a premise in practical reasoning.
- (7') But in other practical environments, it is morally problematic for you to rely on  $p$  as a premise in practical reasoning.
- (8') (7') is true because, in the latter sort of practical environment, you are not epistemically warranted enough to rely on  $p$ .
- (9') The only difference between the former and latter bits of practical reasoning is a difference in the moral facts relevant to your decision.
- (10') So, **Moral Action-Environment Link**: Facts about what matters *morally* (which are not paradigmatically epistemic features of an environment) make a difference as to whether  $S$  is warranted enough to rely on  $p$  as a premise in practical reasoning.

To see the appeal of this new argument for the moral action-environment link, return to the example of the morally apathetic car owner. When no lives are at stake, there is nothing wrong with his using the following reasoning to justify remaining in his easy chair:

My car is parked legally.  
 So no one will be killed on account of the way that my car is parked.  
 So I ought to stay in my easy chair.

But a change in the moral facts in the man's environment suffices to make this reasoning morally unacceptable. Once the man is aware that the maniacal traffic officer poses a potential threat to the lives of innocents, he cannot use this reasoning to justify remaining in his easy chair. Moreover, this reasoning is morally unacceptable for precisely the same reason that the reasoning about life insurance is rationally unacceptable; in both cases, the first premise is insufficiently warranted. Just as you are under rational pressure not to take it for granted that you will go to Blackpool next year once you are offered life insurance, the car owner is under moral

pressure not to take it for granted that his car is parked legally once he learns that lives hang in the balance.

The same considerations that support (8) in the modified argument from high-stakes gambles also motivate (8'). If the morally apathetic car owner gains extremely strong evidence that his car is parked legally, at least one moral problem with his using the reasoning above dissipates. So at least one moral problem with the practical reasoning above is a problem with the strength of the man's warrant for believing that his car is parked legally.

In the reasoning about the car, then, we have an example of practical reasoning that might be either morally acceptable or morally problematic depending on the moral facts in the reasoner's practical environment. This suggests that whether a person is warranted enough to rely on a belief in practical reasoning depends on moral facts. In other words, we can borrow yet another standard line of thinking from the defenders of impurism to generate another argument for the moral action-environment link.

#### **Section 4: An Objection**

It's worth calling attention to one major difference between standard arguments for the action-environment link and the modified versions that I've been offering. This is

a difference between two senses in which I've used the phrase 'warranted enough to act as if  $p$ .' Standard arguments for impurism are concerned with the question of whether a subject is in a position to rationally perform the action that would be rationally appropriate on the assumption that  $p$ .<sup>16</sup> Call this the question of whether a subject is warranted enough to *rationally* act as if  $p$ . The modified arguments for impurism that I've offered, by contrast, are concerned with a slightly different question. In cases like the maniacal traffic officer scenario, the relevant question is whether a subject can, without behaving morally problematically, perform an action that is morally appropriate on the assumption that  $p$ . Call this the question of whether a subject is warranted enough to *morally* act as if  $p$ . These are two importantly different questions, and they concern two importantly different properties.

This difference between ways to understand the phrase 'warranted enough to act as if  $p$ ' suggests a way for the defender of pragmatic encroachment to avoid committing herself to moral encroachment. She might argue, first, that I have used the phrase to pick out properties that are not coextensive. This would mean that, in some cases, a person's epistemic position is strong enough that she is warranted enough to rationally act as if  $p$ , but weak enough that she is not warranted enough to morally act

---

<sup>16</sup> The ensuing discussion could be translated from a discussion about whether a person is warranted enough to act as if  $p$  to a discussion about whether a person is warranted enough to use  $p$  in her practical reasoning. So the same objection, and the same reply, could be made with respect to the argument in section 2 and the argument in section 3.

as if *p*. I'll grant the imagined objector this first step; I suspect that it's true, and it's part of the reason to be interested in moral encroachment in the first place.

Second, the imagined objector might argue that knowledge requires that one's belief is 'warranted enough' in the first sense, but not the second. She might argue, in other words, that a person can know that *p* even though he is not in a position to act as if *p*, morally speaking. Perhaps this is the right thing to say about the morally apathetic car owner in the high-moral-stakes case. Since the fact that innocent people might die does not matter *to him*, there is no encroachment on his knowledge. He knows that his car is parked legally, and he can rationally stay in his easy chair. It's just that his epistemic position isn't strong enough for his remaining in his easy chair (an action that would be morally appropriate on the assumption that his car is parked legally) to be *morally* unproblematic.

To see the problem for this objection, we need to widen our focus and reconsider the typical arguments that impurists offer for the knowledge-action link. After all, if we are to read 'warranted enough to act as if *p*' in the action-environment link as 'warranted enough to rationally act as if *p*,' but not as 'warranted enough to morally act as if *p*,' then we must read the same phrase in the same way when it shows up in the knowledge-action link. Otherwise, the argument from the knowledge-action link and the action-environment link to impurism will not be valid. So, do the arguments for the knowledge-action link support the restricted reading—that is, that if *S* knows that *p*, then *S* is warranted enough to rationally act as if *p*, though not necessarily

warranted enough to morally act as if *p*? As it turns out, no. The reasons to think that there is a tie between knowledge and *rationally* acceptable action are also reasons to think that there is a tie between knowledge and *morally* acceptable action.

Recall the two lines of support for the knowledge-action link that were discussed in section 1 above. Its defenders cite our common practice of defending and encouraging actions by talking about knowledge. They also note the close connection between our judgments about knowledge and our judgments about negligence and blame.

To the extent that these considerations motivate acceptance of a link between knowledge and the rationality of action, they also motivate acceptance of a link between knowledge and the moral appropriateness of action. We can defend the moral appropriateness of action by making claims about knowledge. As we've already noted, the car owner in the high-moral-stakes case can attempt to portray his action as morally unproblematic by saying, "What's all the fuss? I know my car is parked legally." And if a person is worried about whether she has enough evidence to be morally blameless in acting on a given belief, we can encourage her to act by claiming that she knows that her action will work. This suggests that the average person considers knowers of *p* to be warranted enough to morally act as if *p*—where, again, morally acting as if *p* means taking an action that would be morally appropriate on the assumption that *p*.

Further, the conceptual connection between knowledge, negligence, and blame for *immoral* action seems just as strong as the conceptual connection between

knowledge, negligence, and blame for *irrational* action. Recall the example of the parent who responds to a charge of negligence, and attempts to escape blame, by citing her knowledge that the dog playing with her child is a safe one. At first glance, this practice seems to be *even more* strongly connected to the concern that there might be a moral problem with her action than to the concern that there might be a rational problem with her action. In asserting that she knows the dog is safe, the mother is primarily concerned to communicate that her epistemic position with respect to the dog does not make her action morally problematic. Again, this is some evidence that the everyday speaker treats knowledge that *p* as a sufficient condition for being warranted enough to morally act as if *p*.

The arguments for the knowledge-action link, then, do not support a reading of 'warranted enough to act as if *p*' as 'warranted enough to act as if *p* rationally, but not necessarily morally.' To the contrary, those arguments can be also used to support the conclusion that, when a person knows that *p*, the weakness of her epistemic position cannot stand in the way of the moral appropriateness of her acting as if *p*. To the extent that we have a good case for a connection between knowledge and rational action, we also have a good case for a connection between knowledge and morally unproblematic action. So the arguments for impurism suggest that, when a person's belief that *p* has the level of warrant necessary for knowledge, its level of warrant cannot present a problem either for the rationality of action or for the moral appropriateness of action.

This means that the defender of impurism does not have a principled way of endorsing the arguments from high-stake gambles or unacceptable practical reasoning without also endorsing the moral versions of these arguments. Our beliefs are often insufficiently warranted to allow us to rationally take certain high-stakes gambles. But our beliefs are also often insufficiently warranted to us to permissibly act on them when something very morally important hangs in the balance. To the extent that the former phenomenon illustrates pragmatic encroachment on knowledge, the latter phenomenon illustrates moral encroachment on knowledge.

In conclusion: to the extent that there are good arguments for the thesis of pragmatic encroachment on knowledge, there are also good arguments for thesis of moral encroachment on knowledge.



## Chapter 2

### Uncertainty, Belief, and Ethical Weight

This chapter continues to make a case for moral encroachment in epistemology. It also moves beyond the work of Chapter 1 in a significant way: rather than adapting existing arguments from the pragmatic encroachment literature, I offer an original argument for the conclusion that there is moral encroachment in epistemology. I'll make this case by drawing on insights about moral psychology: I argue, roughly, that moral encroachment follows from a plausible connection between our epistemic obligations and our ethical obligations.

Here is the plan for the chapter. In section 1, I introduce a notion of ethical weight and use it to formulate the first premise in my argument: Weight-Sensitivity for Action. I argue that Weight-Sensitivity for Action is extremely plausible. In section 2, I introduce and motivate a second premise: the Belief-Action Link. In section 3, I show that Weight-Sensitivity for Action and Belief-Action Link entail my surprising conclusion: Weight-Sensitivity for Belief. Section 4 responds to objections, and section 5 discusses the implications of my conclusion.

## Section 1: Weight-Sensitivity for Action and Weight-Sensitivity for Belief

Some norms on action take into account an agent's epistemic position. Others do not.

To see this, imagine a case:

**Gift Bags** You must choose to take home one of two gift bags. Your evidence suggests that bag A contains the better gift, but your evidence is misleading; the better gift is in bag B.

Which bag ought you choose? Well, in one sense, you ought to choose bag B, since it has the better gift. Call this the *objective* sense of 'ought.' But it's very plausible that there is another sense in which you ought to pick bag A. This sense of 'ought' is sensitive to your epistemic position; given your misleading evidence, bag A is the bag to pick.<sup>17</sup> This is the sense of 'ought' that I will call the *rational* 'ought'.<sup>18</sup> On my usage, the action that a person *rationally ought* to perform is the action that she ought to perform relative to her epistemic possibilities.

My primary goal in this section is to illustrate a very plausible point: the rational action for a person to take depends not only on the *actual* features of her choice situation but also on the *possible* features of her choice situation. More precisely, I am arguing for:

---

<sup>17</sup> Some objectivists (e.g. Graham 2010) will reject the claim that there are any genuine obligations that are sensitive to epistemic position. Section 4.1 explains why even those who have this theoretical commitment should embrace my argument.

<sup>18</sup> Here, I do not intend to make a claim about the uniquely best way to use the term 'rational' or to refute incompatible analyses of the concept *rational*. I use the term 'rational,' rather, to pick out a notion that is both clearly of theoretical importance and familiar within philosophical discourse about rational action.

**Weight-Sensitivity for Action** When a person's epistemic possibilities are divided, the action that she rationally ought to perform can depend on the degree to which each of those possibilities is ethically weighty for her.

I'll pause to say a bit about the way to understand my terminology. First, this claim concerns conditions in which epistemic possibilities are divided. Divisions in epistemic possibilities are not to be confused with merely psychological uncertainty. They are, rather, conditions that make it epistemically appropriate for a person to be uncertain.

Second, a bit more on the rational 'ought.' I've pointed out that the rational 'ought,' unlike the objective 'ought,' is relativized to a person's epistemic possibilities. But this is not quite enough to settle the question of what flavor of normativity is in play. Plausibly, there are many norms on action that take epistemic position into account. There's a norm that says what would be most *prudent*, given your epistemic position; there's a norm that says what would be *morally* best, given your epistemic position; perhaps there's even a norm that says which move would best promote victory in chess, given your epistemic position. In this chapter, I'll set these norms on action aside. The 'ought' that concerns me is the one weighs *all* these sorts of considerations—prudential, moral, and so on—together. This is sometimes called the all-things-considered 'ought,' 'ought' *simpliciter*, or, in Philippa Foot's memorable turn

of phrase, the “free floating and unsubscripted” ‘ought.’<sup>19</sup> In what follows, I’ll adopt the label “all-things-considered” to pick out this sort of normativity.

Now, distinguish between two sorts of all-things-considered obligations. One sort of obligation is not relativized to the agent’s epistemic position; it takes into account *all* of one’s reasons for action (prudential, moral, and otherwise). This is what I’ve called *objective* obligation. The other sort is relativized to one’s epistemic position; though it weighs up different sorts of reasons, including prudential reasons, moral reasons, and so on, it only takes into account the ones that are within one’s epistemic reach. This is what I’ve called *rational* obligation.<sup>20</sup>

Finally, Weight-Sensitivity for Action is a claim about *ethical weight*, a term of art that I define as follows:

**Ethical Weight** An epistemically possible state of affairs involves *ethical weight* (or ‘is ethically weighty’) to the extent that, if that state of affairs obtains, one’s options differ in their objective choiceworthiness.<sup>21</sup>

To get a better grip on this term’s role, consider the following case:

**Game Show** You are the contestant in a game show in which you must select only one box to take home. Your evidence suggests a 50% chance

---

<sup>19</sup> This formulation can be found in Foot (1997, 320n15). It’s perhaps worth noting that some argue that the notion of an ‘all-things-considered’ or ‘free and unsubscripted’ ought is confused. See, for example, Feldman (2000, 692-3), Tiffany (2007), and Baker (forthcoming). I won’t attempt to address these worries in the dissertation; for some attempts to do so, see Thomson (2001, 46) and McPherson (2018).

<sup>20</sup> Compare Kieseewetter (2017), who suggests that rationality is a matter of responding correctly to one’s *available* reasons, and Lord (2018), who suggests that rationality is a matter of responding correctly to one’s *possessed* reasons. Both availability and possession, as these authors use them, are (at least partly) epistemic notions—Kieseewetter gives an account of availability in terms of evidence (2017, 162-6), and Lord gives an account of possession in terms of knowledge (2018, ch. 3-4).

<sup>21</sup> For a similar notion, see Ross (2006, 258-9).

that *Possibility 1* obtains: both boxes A and B contain five thousand dollars. Your evidence also suggests a 50% chance that *Possibility 2* obtains: box A contains a candy bar and box B is empty. It would be better for you to receive five thousand dollars than a candy bar, and better for you to receive a candy bar than to receive nothing.

Game Show illustrates a difference in the ethical weight of two possibilities. *Possibility 1* is not ethically weighty at all; if that possibility obtains, your two options (pick box A or pick box B) are equally objectively choiceworthy. Loosely speaking: if that possibility obtains, it doesn't matter which box you pick. (No matter what, you'll be going home with five thousand dollars.) *Possibility 2*, on the other hand, is somewhat ethically weighty for you; if that possibility obtains, it does matter somewhat which box you pick. On *Possibility 2*, one option offers you a candy bar, while the other offers you nothing at all, and by stipulation, a candy bar is more choiceworthy than nothing.

We're now ready to see the plausibility of Weight-Sensitivity for Action. It claims that what you ought to do, rationally speaking, does not solely depend on *how likely* credible possibilities are; rather, a mere change to the *ethical weight* of those credible possibilities can make a difference to what you ought to do. The Game Show case helps to illustrate the shape of this connection between ethical weight and rational obligation. In Game Show, you rationally ought to pick box A. Why? Well, roughly, because *Possibility 1*'s lack of ethical weight means that you need not take it into account. *Possibility 1*, no matter how epistemically likely, implies nothing about which choice is best. So the only possibility that is relevant for your decision is *Possibility 2*; in that scenario, your choice between the two boxes matters. In other words,

*Possibility 2*, unlike *Possibility 1*, has ethical weight. So you rationally ought to act in the way that would be best if *Possibility 2* obtained.

There's a familiar sort of reasoning, then, that showcases a connection between ethical weight and rational obligations. Now, this is not to say that ethical weight is the sole determinant of a person's rational obligations; among other things, a rational agent is also sensitive to the *likelihood* of credible possibilities.<sup>22</sup> But, in at least some scenarios, a change to the ethical weight of an epistemic possibility is sufficient to bring about a change in one's rational obligations. Suppose, for instance, that we modify the Game Show example by removing the candy bar from the equation. This eliminates the ethical weight of *Possibility 2*; if that possibility obtains, then both boxes are empty, and therefore equally objectively choiceworthy. And this is enough to change the landscape of your rational obligations. In the new, modified case, you are no longer rationally obligated to pick box A; rather, it is rationally permissible for you to choose either box. A difference in the ethical weight of one's epistemic possibilities, then, can suffice to make a difference to one's rational obligations. In other words, Weight-Sensitivity for Action is true.

---

<sup>22</sup> It is worth emphasizing that my argument does not rely on the assumption that the best way to model rational decision-making under uncertainty will advert directly to considerations about ethical weight. I leave open, for instance, the possibility that rational obligations are always determined by a simple principle: a person always rationally ought to maximize expected value. Even on a model such as this, however, the key claim in the text is true: sometimes, a change in a proposition's ethical weight is sufficient to bring about a change in rational obligation. Thanks to an anonymous referee for encouraging me to make this clear.

These brief remarks have shown why Weight-Sensitivity for Action is highly plausible. The following claim, on the other hand, does not enjoy much *prima facie* plausibility:

**Weight-Sensitivity for Belief** When a person's epistemic possibilities are divided with respect to  $p$ , the doxastic states that she rationally ought to have with respect to  $p$  can depend on the degree to which each of those possibilities is ethically weighty for her.

At first glance, Weight-Sensitivity for Belief seems straightforwardly false. How could the question of whether I rationally ought to believe  $p$  depend on whether  $p$  implies that my choices matter? The fact that my choices would matter a great deal if  $p$  were true (or false) does not seem to bear on whether  $p$  actually is true. Nor does it seem to bear on whether  $p$  is likely to be true given my evidence. Despite this initial oddness, I'll be arguing for Weight-Sensitivity for Belief below.

What's more, I'll be arguing for Weight-Sensitivity for Belief while maintaining the traditional distinction between epistemic and non-epistemic reasons for belief. Paradigmatic non-epistemic reasons for belief include direct threats or bribes for belief, and many hold that they cannot make a difference to whether a belief is rational.<sup>23</sup> Suppose, for instance, that an eccentric millionaire offers to give me a huge monetary reward if I believe that Madrid is the capital of Australia. This is a paradigmatic non-epistemic reason for belief. Though there is some sense in which it might be best for me to respond to the bribe by believing that Madrid is the capital of Australia, there is also clearly some important norm on belief-formation that forbids

---

<sup>23</sup> For this traditional perspective, see Kelly (2002); for an important alternative, see Rinard (2017).

my doing so. It's the latter sort of norm that concerns me in this chapter; when I make claims about rational belief, I mean to be picking out this *epistemic* sort of rationality. Importantly, it is a sort of rationality that is not sensitive to considerations like direct threats or bribes for belief.

Why does this distinction between epistemic and non-epistemic reasons matter? Well, if one thinks that the rationality of a belief can be influenced by bribes or threats for belief, one may not see what is counterintuitive and interesting about Weight-Sensitivity for Belief. According to Weight-Sensitivity for Belief, the rationality of belief depends on one's choice situation. But if non-epistemic reasons like bribes and threats influence the rationality of belief, then this is old news. My argument in this chapter will aim at a more exciting conclusion: that even our more austere notion of epistemic rationality—the one that sets aside paradigmatically non-epistemic reasons—is nevertheless subtly sensitive to ethical weight. (To put this point in terms familiar from chapter 1: in setting aside paradigmatically non-epistemic reasons, I am *not* setting aside *all but the* paradigmatically epistemic reasons, and therefore, I am leaving open the possibility of impurism. I say much more about the distinction between epistemic and non-epistemic reasons, and how it interacts with impurism, in Chapter 3.)

## **Section 2: The Belief-Action Link**

The previous section showcased the initial plausibility of Weight-Sensitivity for Action and the initial implausibility of Weight-Sensitivity for Belief. This section



introduces and defends a weak, plausible principle connecting belief and action.<sup>24</sup> In section 3, I will use this principle to argue from Weight-Sensitivity for Action to Weight-Sensitivity for Belief.

Here is the principle I'll be discussing:

**Belief-Action Link:** If you rationally ought to believe of a particular action,  $\varphi$ ,<sup>25</sup> that you objectively ought to perform it, then it is rationally permissible for you to  $\varphi$ .

To better grasp this principle, imagine that you believe that you objectively ought to mow the lawn. Further imagine that your belief is rationally appropriate. Now try to imagine that, in the same case, it is rationally impermissible for you to mow your lawn. Something seems to have gone wrong; if your belief is appropriate given your epistemic position, then surely your epistemic position doesn't forbid you to act as your belief suggests you should!

We can supplement these intuitive grounds for the Belief-Action Link with an argument in its favor:

---

<sup>24</sup> The Belief-Action Link offers only a fairly weak, plausible connection between normative belief and action. But some have defended stronger, more contentious connections. See, for instance, the "ought infallibilism" defended in Way and Whiting (2016): "if you justifiably believe that you ought to  $\varphi$ , then you ought to  $\varphi$ ." See also Fantl and McGrath's "KB inference" (2007, 565): "S knows that A is the best thing she can do... So, S is rational to do A." The KB inference differs from the Belief-Action Link in two notable ways. First, rather than making any claims about all-things-considered norms on action, the KB inference limits itself to claims about decision-theoretic rationality. Second, it offers a sufficient condition on rational *obligation* to act; the Belief-Action Link offers only a sufficient condition on rational *permission* to act. For more on this, see footnote 26 below.

<sup>25</sup> For this principle to be plausible, your normative belief must be a belief about a particular action *de re*, not *de dicto*. A belief that you objectively ought to do some action, understood *de dicto*, can certainly come apart from rational obligation. For instance, I can believe that I objectively ought to perform *the action, whatever it is, that maximizes utility*, but be clueless as to which action would maximize utility.

- (1) If you rationally ought to believe of a particular action,  $\varphi$ , that you objectively ought to perform it,  $\varphi$ , then you rationally ought to believe of  $\varphi$  that you rationally ought to perform it.
- (2) If you rationally ought to believe of a particular action,  $\varphi$ , that you rationally ought to perform it, then it is rationally permissible for you to  $\varphi$ .

So, **Belief-Action Link:** If you rationally ought to believe of a particular action,  $\varphi$ , that you objectively ought to perform it, then it is rationally permissible for you to  $\varphi$ .<sup>26</sup>

I'll now offer some grounds for accepting the premises of this argument, taking them in reverse order.

Premise (2) is, roughly, the claim that epistemic norms and norms on rational action never conspire to require akrasia—that is, an agent who acts in accordance with these norms is never required to believe that she rationally ought to  $\varphi$  while also not  $\varphi$ ing. It's quite plausible that a clear-eyed agent and rational agent is never *permitted* to be akratic, and it's even more plausible that a clear-eyed and rational agent is never *required* to be akratic.

Premise (1) claims, roughly, that what it's rationally appropriate to believe about your *objective* obligations determines what it's rationally appropriate to believe about your *rational* obligations.

---

<sup>26</sup> Thanks to Declan Smithies for this way of presenting the case for the Belief-Action Link.

Now, there are reasons to worry about a tight connection between beliefs about objective obligations and beliefs about rational obligations. For instance, the following example might seem to cause trouble for any such connection:

Jill is a physician... [whose] patient, John,... has a minor but not trivial skin complaint. She has three drugs to choose from... Drug A is very likely to relieve the condition but will not completely cure it. One of drugs B and C will completely cure the skin condition; the other though will kill the patient, and there is no way that she can tell which of the two [is which]. (Jackson 1991, 462-3)

This case might seem to challenge premise (1). It shows that, in certain cases, there's nothing rationally inappropriate about an agent who believes that she objectively ought *not* perform some particular action (administer drug A) while simultaneously believing that she rationally ought to perform it. If that's right, is there any reason to think that there's something odd about an agent who believes that she objectively *ought* to perform some particular action while believing that she rationally ought *not* perform it?

Yes, there is. Beliefs about particular objective obligations play an importantly different role than do beliefs in cases like Jill's. Outright beliefs about particular objective obligations, loosely speaking, provide fixed points within deliberation about what to do. Once I have settled the question of what my particular objective obligations are, there is no question left, from my perspective, as to how I ought to guide my action: I ought to do the objectively required action.<sup>27</sup> From within the

---

<sup>27</sup> It might be tempting to suppose that this conclusion faces counterexamples when a person's choices have consequences for her future actions. Take, for instance, the following revision to the case in the

deliberator's perspective, it would be bizarre to ask the question "I'm aware of exactly what I ought to do, *objectively* speaking, but how ought I act, *rationally* speaking?"<sup>28</sup>

Premise (1), then, is only concerned with cases—notably *unlike* Jill's—in which an agent is rationally obligated to form beliefs about her particular objective obligations. In any such situation, it suggests, an agent is in a position to conclude that her epistemic position supports following those objective obligations.

Premises (1) and (2) help to show why the Belief-Action link is highly plausible. But, as the next section shows, accepting the Belief-Action Link alongside Weight-Sensitivity for Action leads to a surprising result.

### **Section 3: The Argument for Weight-Sensitivity for Belief**

We've now taken a close look at all the claims that make up this chapter's core argument. Let's now consider the argument itself.

---

main text. Suppose that Jill faces a choice between ordering a package that includes solely drug A and ordering a package that includes drugs B and C. She knows that she will not be able to revise her order, to acquire the drugs that she does not order, or to gain any further information about the effects of the drugs. It may be tempting to see this case as one in which Jill rationally ought to believe that she *objectively* ought to order the package including drugs B and C (after all, the maximally effective drug is in that package), and in which she also rationally ought to believe that she *rationally* ought to order the package including drug A. If that's right, premise (1) fails; Jill's rational belief about her objective obligations does not settle the matter of what to believe about her rational obligations.

But I see no appeal to the claim that Jill rationally ought to believe that she objectively ought to order the package including drugs B and C. On my usage, the facts about what a person objectively ought to do are the facts about what she ought to do in light of *all* the facts. There is an important difference between what a person ought to do in light of all the facts and what a person would be obliged to do if she knew that, throughout her life, she would actually be aware of all the facts. *If* Jill could count on later coming to know which of drugs B and C is deadly, she would be objectively obliged to order drugs B and C. But, crucially, she knows that she will later act under conditions of incomplete information. In light of all the facts, *including* this fact about the incompleteness of her future information, she ought to order drug A. (Thanks to Tristram McPherson for this objection.)

<sup>28</sup> Compare Sepielli (2009, 55).

Premise 1: **Weight-Sensitivity for Action** When a person's epistemic possibilities are divided, the action that she rationally ought to perform can depend on the degree to which each of those possibilities is ethically weighty for her.

Premise 2: **Belief-Action Link**: If you rationally ought to believe of a particular action,  $\varphi$ , that you objectively ought to perform it, then it is rationally permissible for you to  $\varphi$ .

So, **Weight-Sensitivity for Belief** When a person's epistemic possibilities are divided with respect to  $p$ , the doxastic states that she rationally ought to have with respect to  $p$  can depend on the degree to which each of those possibilities is ethically weighty for her.

More loosely stated: if what you rationally ought to do is dependent on ethical weight, and what you rationally ought to believe has implications for what you rationally ought to do, then what you rationally ought to believe is also dependent on ethical weight.

To illustrate why this argument works, I'll discuss a pair of cases. First, consider:

**Wayne** Wayne is late to work. He has just dropped his daughter off at kindergarten and is exiting the parking lot. As he turns a corner, the back wheels of his car hit an unexpected bump. He cannot quite see what caused the bump. He (appropriately) has a credence of .95 that the bump came from his hitting a tree branch. But he (also appropriately) has a .05 credence that the bump came from his hitting a child's leg.

How, *rationally* speaking, ought Wayne act? If Weight-Sensitivity for Action is true, the answer depends on the ethical weight of credible possibilities. The possibility that he has hit a child's leg is a very ethically weighty one for Wayne; in that scenario, his options differ greatly in their objective choiceworthiness. If Wayne has hit a child,

then it would be much better for him to stop and help the child than for him to ignore the bump and drive to work. The possibility that he has hit a branch, by contrast, is much less ethically weighty for Wayne.

Weight-Sensitivity for Action suggests that, given the weight of the possibility that he hit a child, Wayne could be rationally required to remain in the parking lot rather than driving straight to work. Combine this fact with the Belief-Action Link. Suppose that Wayne rationally ought not drive straight to work. That means, when  $\varphi$  is Wayne's driving straight to work, the Belief-Action Link's consequent is false. So its antecedent must be false as well: it must be false that Wayne rationally ought to believe that he objectively ought to drive straight to work. And it seems that the Belief-Action Link is wholly sensible in ruling out that possibility. Imagine that Wayne stops his car and gets out to look around, all the while believing, *Objectively, I ought not do this.*<sup>29</sup> This behavior lacks an important sort of coherence; surely, it cannot be the appropriate way for Wayne to fulfill the requirements on his belief and action. Wayne's case, then, shows how Weight-Sensitivity for Action and the Belief-Action Link can work together to place constraints on what he is obligated to believe.

Let's move on to a second case:

**Diana** Diana is trapped in a room where she is watching a camera feed of Wayne's exit from the parking lot. The camera's perspective is nearly identical to Wayne's, so she has the same epistemic possibilities that

---

<sup>29</sup> Brown (2008, sec. 7) describes a case in which, she suggests, there is nothing odd about maintaining a belief while acting in order to respect the possibility that it is false. Wayne's case, which involves an action-guiding normative belief, brings out the problem with applying this strategy too broadly: it threatens to license, as rational, a wide range of akratic actions.

Wayne has: she appropriately has .95 credence that he hit a tree branch and .05 credence that he hit a child's leg. She knows that Wayne is late for work and that she cannot affect what happens next.

Diana's case is a particularly important one, precisely because we cannot treat it in the same way that we've just treated Wayne's. Weight-Sensitivity for Action and the Belief-Action Link combined to tell us something important about Wayne's beliefs: he is not obligated to believe that he objectively ought to drive straight to work. By contrast, Weight-Sensitivity for Action and the Belief-Action Link tell us nothing about what Diana ought to believe. The reason is that, unlike Wayne, Diane faces no chance of *akrasia*. If Wayne believes that he objectively ought to drive straight to work, he must do so, on pain of incoherence. Not so for Diana. Since she faces no choices related to Wayne's situation, any action she takes will cohere perfectly well with the belief that Wayne objectively ought to drive straight to work. So the Belief-Action Link and Weight-Sensitivity cannot rule out the possibility that Diana is required to believe the proposition in which (appropriately) she has .95 credence: that Wayne ought to drive straight to work.

Now, this is not proof that Wayne and Diana face different doxastic obligations. Perhaps Diana, like Wayne, is not rationally obligated to believe that Wayne objectively ought to drive straight to work. But, on the face of it, this would require some explaining. Diana is, after all, in quite a strong epistemic position with respect to an everyday proposition: *that Wayne has hit a tree branch*. Why couldn't she be rationally required to respond to her strong perceptual evidence forming an outright

belief in this everyday proposition? And, if she is in a position to form that belief, why couldn't she be required to reason from it to the conclusion that Wayne objectively ought to drive to work?

Suppose that Diana is indeed obliged to form the belief that Wayne objectively ought to drive straight to work. If she is, then Wayne and Diana have different doxastic obligations with respect to the proposition *Wayne objectively ought to drive straight to work*. Diana rationally ought to believe that proposition, but it's not the case that Wayne should. But the only difference between Wayne and Diana is a difference in the choices that they face—a difference, more specifically, in the ethical weight of the epistemic possibilities they both share. Therefore, Weight-Sensitivity for Belief is true.

The most immediately attractive way to resist this case for Weight-Sensitivity for Belief is to hold that Diana (and anyone in a relevantly similar situation) simply could not be rationally required to form the belief for which she has such strong evidence. But this is a surprising and strong conclusion, and anyone who takes it on takes on a heavy explanatory burden. It's worth mentioning up front one tempting way to address that burden: one might say that the ethical weight *for Wayne* explains why Diana may suspend judgment. I grant that this explanation may be right. But note that offering this explanation is not a way of avoiding Weight-Sensitivity for Belief; this is simply a different picture of the particular role that ethical weight plays in fixing epistemic standards. On the alternate picture, it is not ethical weight for the *believer*



(Diana), but ethical weight for the *agent under evaluation* (Wayne) that makes a difference for rational belief. The question of which of these two proposals to accept lies beyond the aims of this chapter, and in what follows, I'll be neutral between them. We've now seen that the epistemologist who seeks to explain why suspense of judgment is appropriate in Wayne's (and possibly also Diana's) case, *without* appealing to facts about ethical weight, faces a difficult explanatory task. In our next section, we'll consider some strategies for taking up that explanatory task, and I'll argue that none of those strategies can be made to work.

The positive argument of this chapter is now on the table: I've explained why Weight-Sensitivity for Action and the Belief-Action Link support Weight-Sensitivity for Belief. Weight-Sensitivity for Action and the Belief-Action Link are intuitively attractive, but Weight-Sensitivity for Belief is not. So the argument I've offered puts us in an apparently unfortunate position; we must either reject one of its initially plausible premises or accept its initially implausible conclusion.

#### **Section 4: Objections to the Belief-Action Link**

I've now made a case for Weight-Sensitivity for Belief. I've argued, in other words, that what a person rationally ought to believe depends on facts about ethical weight. Many readers will find this conclusion odd. The cases of Wayne and Diana highlight the oddness: how could two people with just the same epistemic position face different requirements on belief?

I agree that the idea that ethical weight makes a difference in epistemology is counterintuitive, and I'm sure that many readers will want to find a way to reject it. To do so, the reader will have to find a way to reject one of my two premises. Since Weight-Sensitivity for Action is an extremely weak and plausible claim, the best hope for the objector is to reject the Belief-Action Link. This section considers several attempts to do just that. In the end, none of these strategies are attractive: they either fail to provide principled grounds for rejecting the Belief-Action Link, or they incur unacceptable costs elsewhere. The result is that, oddness or no, we should accept Weight-Sensitivity for Belief.

#### *4.1 Objective Obligations and Rational Obligations*

The argument above makes use of a distinction between *objective* obligations and *rational* obligations. Cases like Gift Bag, Game Show, and Jackson's case of Jill the physician show the appeal of this distinction; intuitively, these are cases in which the thing to do *objectively* speaking differs from the thing to do *rationally* speaking. But the precise nature of the distinction illustrated by these cases is a matter of controversy. Indeed, some approaches seem, at first, incompatible with the argument I've offered for Weight-Sensitivity for Belief.

Consider, for example, an 'objectivist' view according to which no all-things-considered ethical obligations are relativized to an agent's epistemic possibilities. On this view, there are no *rational* obligations; the only obligations worthy of the name are objective ones. If this picture is right, it might seem that no premise even close to

the Belief-Action Link will be true, because the mere fact that I rationally believe something about my objective obligations does not entail anything at all about my actual obligations. Does this theoretical commitment provide a way to neatly dodge the argument for Weight-Sensitivity for Belief?

It does not. Roughly, that's because talk of 'rational obligation' is not a vital part of the argument for Weight-Sensitivity for Belief. Suppose that a strong objectivist view wins out: the best view, at the end of the day, entails that no all-things-considered obligations are relativized to epistemic possibilities. In that case, my argument can be recast to capture, in different terminology, the insight that rational beliefs about objective obligations have important implications for action.

To see how this might work, consider a particular way of developing objectivism, inspired by Graham (2010).<sup>30</sup> Having denied that there are any rational obligations, this objectivist view faces a riddle: how can we explain why, in Jackson's physician case, Jill ought to give her patient Drug A? This 'ought'-claim cannot be a claim about objective obligation. To solve riddles of this sort, the objectivist must supplement her account of obligation with a theoretical account of the requirement—whatever it is—that Jill so clearly faces. One way to do this is to say that, though there is no sense in

---

<sup>30</sup> It's worth noting that Graham defends objectivism about *moral* obligation. But Graham's arguments in favor of objectivism can, in principle, be applied to all-things-considered obligations in just the same way that they apply to moral ones.

which Jill is obliged to give her patient Drug A, administering Drug A is nevertheless *the action that a conscientious agent in her position would take*.

The argument for Weight-Sensitivity for Belief can be recast using this language in place of talk of 'rational obligations':

**Weight-Sensitivity for Action\*** When a person's epistemic possibilities are divided, the action that she would perform if conscientious can depend on the degree to which each of those possibilities is ethically weighty for her.

**Belief-Action Link\*** If you rationally ought to believe of a particular action,  $\varphi$ , that ought to perform it, then acting conscientiously does not preclude you from  $\varphi$ ing.

**Weight-Sensitivity for Belief** When a person's epistemic possibilities are divided, the doxastic states that she rationally ought to have can depend on the degree to which each of those possibilities is ethically weighty for her.

The objectivist about obligation should accept the premises of this revised argument.

A conscientious agent's behavior is surely sensitive to weight. And Belief-Action Link\* is entirely plausible; it would be a failure of conscientious action for an agent to believe herself (objectively) obliged to perform some particular action,  $\varphi$ , while failing to  $\varphi$ .

I've chosen a particularly austere picture of the nature of all-things-obligation to show that my argument can be adapted to accommodate even extreme departures from talk of rational and objective obligation. The lesson of this section will also hold for

less extreme departures.<sup>31</sup> On any plausible view of obligation, premises suitably similar to Weight-Sensitivity for Action and the Belief-Action Link will be true.

#### *4.2 Infallibilism about Rational Belief*

In section 3, I charged the opponent of Weight-Sensitivity for Belief with an explanatory challenge: she must explain, without appealing to facts about ethical weight, why it's not the case that Wayne (and possibly also Diana) ought to believe that he objectively ought to drive straight to work. Some may want to take up this explanatory challenge by pointing to Wayne's (and possibly also Diana's) uncertainty. Perhaps, whenever a person is rationally required to form some belief, there is no epistemic possibility for her that the belief is false. Call this principle *infallibilism about rational belief*.

If infallibilism about rational belief is true, whenever epistemic possibilities about some ethical proposition are divided for an agent, she is not rationally required to form a belief about that proposition's truth or falsity. As a result, the question of whether an agent is required to believe that  $p$  is never affected by ethical weight. If  $p$  is certain for her, then there is no chance that, as in Wayne's case, a very ethically weighty epistemic possibility that  $\neg p$  will have implications for her actions. If, on the

---

<sup>31</sup> See Lord (2015, 31-2) for a more moderate stance toward the roles played by objective and rational obligations in deliberation.

other hand,  $p$  is *not* certain for her, then it is not the case that she rationally ought to believe that  $p$ .

I'll now distinguish between two ways of developing infallibilism about rational belief. The first is an unattractive one because it threatens to diminish our rational obligations in an implausible way. The second, when developed to avoid this problem, is entirely compatible with the spirit of the argument above.

First, consider an infallibilist view on which everyday, prosaic propositions are generally not certain for us. There is, for instance, a possibility for me that my wife has been replaced by a doppelgänger, so it is not certain for me that she is standing in front of me right now. On an infallibilist view, the fact that these propositions are not certain for me entails that I am not rationally required to believe them. And, similarly, this sort of infallibilism can easily explain why both Diana and Wayne are not rationally required to believe that he objectively ought to drive straight to work; the proposition that Wayne hit a tree branch is not certain for either of them.

This sort of infallibilism rejects the argument for Weight-Sensitivity for Belief at an unacceptable cost.<sup>32</sup> Surely, there are plenty of facts about what a rational person would have to believe about her surroundings. The version of infallibilism we are currently considering seems unable to account for those facts. The conclusion that the

---

<sup>32</sup> This is a cost borne most obviously by some who work in the Bayesian tradition; see, e.g., Jeffrey (1971, 171-2).

rational requirements on our beliefs are so sparse seems, in fact, far less plausible than the conclusion that those requirements are sensitive to the choices we face.

Second, consider an infallibilist view on which many everyday propositions *are* in fact certain for us.<sup>33</sup> On this view, it is usually *not* possible for me that my wife has been replaced by a doppelgänger; in ordinary cases, this proposition is certain for me, and I may well be required to believe it.

This variant on infallibilism does not threaten to vitiate our rational obligations in the same way that the previous one did. But it is also entirely compatible with the argument offered above for Weight-Sensitivity for Belief. This sort of infallibilism must allow that, in the face of great ethical weight, ordinary propositions often are *not* certain for us. And the notion that our epistemic possibilities *themselves* are sensitive to ethical weight provides a tidy way to explain this phenomenon. Indeed, it's entirely compatible with this sort of infallibilism that, while it is certain for Diana that Wayne ought to drive straight to work, that proposition is not certain for Wayne, and that the difference between the choices faced by the two explains this difference in their epistemic possibilities. This is an attractive way of using infallibilist machinery to explain just why, in virtue of their different choices, Diana and Wayne have different doxastic obligations.

---

<sup>33</sup> According to Williamson (2000), many everyday propositions are known, and knowledge requires probability 1.0. So, if we understand 'certain for an agent' to mean 'having probability 1.0 for an agent' (rather than, for instance, 'justifying accepting bets at any odds'), Williamson's view offers an infallibilism on which everyday propositions are certain for us.

Infallibilism, then, does not provide an attractive way to avoid Weight-Sensitivity for Belief. Some ways of developing the view are objectionably deflationist about rational obligation, and others do not suggest any reason to worry about the argument above.

#### *4.3 Level-Splitting Views of Rationality*

One key premise in the argument for the Belief-Action Link is the claim that, if rationality requires you to believe that you rationally ought to  $\phi$ , then you rationally ought to  $\phi$ . I motivated this claim briefly by noting that, if it were false, an ideally rational agent could sometimes be required to be akratic—that is, to believe that she rationally ought to  $\phi$  while simultaneously making no effort to  $\phi$ . And it's very plausible that rationality does not *require* us to be akratic in this way.

There is at least one way of approaching debates about rationality, however, on which this initially plausible claim is false. This is the “level-splitting” view of rationality, a view on which the thing one ought to do (or believe) can diverge from the thing that one ought to *believe* that one ought to do (or believe). In the form relevant to this chapter's argument, a level-splitting view of rationality would have to hold that what one rationally ought to do can come apart from what one rationally ought to believe about what one rationally ought to do.<sup>34</sup>

---

<sup>34</sup> An anonymous referee notes that, given a level-splitting approach to rationality, there may also be counterexamples to Weight-Sensitivity for Action. But, for the reasons I cite in the main text, it would be unacceptably *ad hoc* to treat a case like Wayne's as one of those counterexamples.



Adopting a level-splitting approach to rationality to explain cases like Wayne's, however, is an objectionably *ad hoc* move. Broadly speaking, there are two important reasons for which philosophers adopt level-splitting views, and neither offers a principled explanation for the sort of *akrasia* that concerns us.

First, several epistemologists accept a level-splitting view on the grounds that it is the most viable way of accounting for different roles played by first-order evidence and higher-order evidence.<sup>35</sup> But, even if we accept this reason for accepting ideally rational epistemic *akrasia*, we should not take it to justify *akrasia* in a case like Wayne's. Wayne's case is thinly sketched, and it need not be understood to involve the characteristic tension between first-order and higher-order evidence.

Second, some have motivated a level-splitting view of rationality by noting that, in some cases that involve purely normative uncertainty, our evaluation of *action* seems to proceed very differently from our evaluation of *beliefs about norms on action*.<sup>36</sup> We consider a morally monstrous person's action blameworthy, for instance, even when we consider her to have reasoned to her flawed moral beliefs impeccably. This approach offers a different sort of case in which the rational requirements on action and belief might diverge: a case in which a person deliberates under conditions of purely normative uncertainty. But it would be *ad hoc* to construe Wayne's case in the same way; the uncertainty in Wayne's case derives from purely *non-normative*

---

<sup>35</sup> See especially Horowitz (2012) and Lasonen-Aarnio (2014).

<sup>36</sup> See, e.g., Coates (2012).

uncertainty about which object he has hit. And many have noted that level-splitting is not as tempting in cases of non-normative uncertainty as it is in cases of normative uncertainty.<sup>37</sup> Consider, for instance, a case in which a person believes that she ought to take a road that, unbeknownst to her, is closed for construction. In this case, there is no appeal to the view that, while the person's belief is rationally impeccable, her action is blameworthy.

Since Wayne is uncertain solely in virtue of uncertainty about non-normative facts, and there are principled reasons for thinking that this sort of uncertainty is importantly different from normative uncertainty, we should not be tempted to use this line of thought to justify a level-splitting approach to rationality in the cases that concern us.

### **Section 5: Implications of Weight-Sensitivity for Belief**

This section discusses the implications of Weight-Sensitivity for Belief and situates the thesis with respect to some current trends in epistemology.

If Weight-Sensitivity for Belief is true, then rationally holding beliefs about our own obligations in certain choice situations will sometimes require us to have higher levels of rational confidence than would otherwise be required. Agents in positions

---

<sup>37</sup> See, e.g., Weatherson (2014) and Harman (2015).

like Wayne's will be harder-put to justify beliefs about objective obligations than will agents whose beliefs are not connected to any pressing decisions.

In fact, there are also good reasons to think that weight-sensitivity in epistemology extends beyond beliefs about objective obligations. If, for instance, the question of whether Wayne objectively ought to drive to work hinges entirely on the question of whether he has hit a tree branch, it would be very odd for him (or Diana) to form the outright belief that he has hit a tree branch while suspending judgment about whether he ought to drive to work. Since epistemic rationality requires that beliefs be coherent, weight-sensitivity in epistemology will likely spread from beliefs about our obligations to the other beliefs with which beliefs about obligations must cohere. This means that even the rationality of beliefs about prosaic matters of fact (e.g. *that my car has just hit a tree branch*) will depend on facts about ethical weight.

This conclusion echoes a trend that has some defenders in contemporary analytic epistemology. Recent years have seen an increase in views that posit connections between paradigmatically practical notions like *stakes* or *actionability* and paradigmatically epistemological notions like *knowledge*.<sup>38</sup> In a slogan, these views posit 'pragmatic encroachment' in epistemology. The defenders of pragmatic encroachment may be happy to accept that certain facts about our choices are connected to the epistemic standards for belief about any given subject matter. This

---

<sup>38</sup> For defenses of pragmatic encroachment, see Stanley (2005), Hawthorne and Stanley (2008), and Fantl and McGrath (2009). For a useful overview, see Kim (2017).

conclusion of this chapter, then, entails a result that already has some defenders: norms on action encroach on epistemic normativity. But the argument offered here differs from extant arguments for pragmatic encroachment in two crucial ways.

First, arguments in favor of pragmatic encroachment generally rely on premises that connect belief and action in ways that are far more systematic, and potentially more difficult to swallow, than the Belief-Action Link offered above.<sup>39</sup> The argument offered above, then, may attract even some who resist systematic defenses of pragmatic encroachment.

Second, my argument suggests that encroachment in epistemology may not be limited to merely *pragmatic* encroachment. The defenders of pragmatic encroachment generally restrict their interest to connections between a person's beliefs and her practical interests.<sup>40</sup> Very few suggest that rational belief has anything to do with

---

<sup>39</sup> Three examples: Fantl and McGrath (2009) argue for pragmatic encroachment on the grounds of entirely general principles about the way that reasons justify both beliefs and actions. (Consider, for example, their *Unity Thesis*: "If p is warranted enough to be a reason you have to believe that q, for any q, then p is warranted enough to be a reason you have to  $\phi$ , for any  $\phi$ " [73].) Hawthorne and Stanley (2008) draw on quite general considerations regarding the role that epistemic appraisals play in the evaluation of action (cf. Fantl and McGrath [2007]). Basu and Schroeder (2019) argue for moral encroachment by defending the controversial idea that we can morally wrong others through the beliefs they form. I have taken on no commitments about either of these topics.

The argument in the current literature that resembles mine most can be found in Fantl and McGrath (2007). My argument differs from Fantl and McGrath's in two crucial ways. First, it provides an explicit argument in favor of, and thereby sheds light on the nature of, my proposed connection between belief and action. Fantl and McGrath, by contrast, do not provide an argument in favor of their analogous "KB inference." Second, while Fantl and McGrath's argument shows merely *pragmatic* encroachment in epistemology, I argue for a broader connection between ethics and rational belief. I say more about this in the main text below.

<sup>40</sup> Since defenders of pragmatic encroachment do not usually draw distinctions between different norms on action, they often stop short of explicitly defining pragmatic encroachment as a connection between epistemology and practical interests. Nevertheless, their writing does consistently suggest a focus on matters of importance *to the subject* rather than matters of *all-things-considered* importance;

*other* norms on action: for instance, a person's moral obligations or all-things-considered obligations.

To take an example: there are several extant arguments that imply that whether I know that a particular train is leaving on time can depend on how important it is *for me* that I complete my journey promptly. But almost no one in the current literature argues that whether I know that a particular train is leaving on time can depend on how *all-things-considered* important it is that I complete my journey promptly. (Imagine, for instance, that I am carrying the only antidote to a deadly disease that is racing toward my train's destination.) *Pragmatic* encroachment in epistemology is frequently defended, but encroachment from *all-things-considered ethical obligations* in epistemology is not.

In the previous chapter ("Pragmatic Encroachment and Moral Encroachment"), I argued that this asymmetry is not a principled one; the most prominent arguments for pragmatic encroachment can be generalized to support the notably different thesis of *moral* encroachment. But that chapter remained neutral on the question of whether there are, in the final analysis, any convincing arguments for either thesis. The argument of this chapter represents an important further step: here, I have offered an original positive argument for the conclusion that encroachment in epistemology is not merely pragmatic. To see this, note that the argument from

---

see, e.g., Fantl and McGrath (2007, p. 55 n1), Hawthorne and Stanley (2008, p. 583), and Stanley (2005, pp. 92-3).

Weight-Sensitivity for Action and the Belief-Action Link is carried out in terms of *all-things-considered* obligations. It relies not on the claim that Wayne *prudentially* ought not drive straight to work, but that he rationally ought not drive straight to work *all-things-considered*. It's an argument, then, for the conclusion that the rationality of belief is importantly tied to our all-things-considered normative obligations.

Wayne's case illustrates the most striking upshot of this conclusion: since moral facts can be difference-makers for all-things-considered obligations, the moral importance of a state of affairs (in Wayne's case, the moral importance of stopping to see if he has hit a child) can be a difference-maker in epistemology. Encroachment in epistemology from all-things-considered norms on action, then, carries with it *moral* encroachment in epistemology.

This chapter has argued that our epistemic obligations depend on our all-things-considered ethical obligations. This final section has highlighted a few respects in which, even given the recent "practical turn" in epistemology, my conclusion is a powerful and distinctive one. But, as the example of Wayne suggests, its primary upshot is that we should be more careful to form outright beliefs when those beliefs suggest that our choices matter a great deal. Though surprising, this may be a result that we can live with—and, perhaps, even a result that can tell us something about how to live well.

## Chapter 3

### Moral Encroachment and Reasons of the Wrong Kind

The previous two chapters made a case for the conclusion that there is moral encroachment in epistemology. In this chapter and the next one, I take that conclusion for granted. My goal in those chapters is to show, in more detail, just how moral encroachment in epistemology works.

This chapter is concerned with a distinction that I have gestured at throughout the past two chapters: the distinction between *epistemic* and *non-epistemic* factors. It's now time to take a closer look at that distinction. By understanding the distinction as a subcase of a more general distinction (the distinction between *reasons of the right kind* and *reasons of the wrong kind*), I'll argue, we can make progress in determining the extension of moral encroachment—in determining, that is, precisely *which* moral features of the world are apt to make a difference to knowledge.

To frame our discussion of the epistemic/non-epistemic distinction, let's consider two cases:

**Bribe for Belief** A demon will donate \$1 million to a good charity only if you form the belief that the number of stars in the universe is even.

**Bribe for Withholding** A demon will donate \$1 million to a good charity only if you withhold belief regarding the proposition that  $2 + 2 = 4$ .

Suppose that you somehow earn the demon's bribe for withholding; you withhold belief as to whether  $2 + 2$  is 4. It is morally good that you do so.<sup>41</sup> Nevertheless, by withholding, you violate an important norm. Intuitively, that norm is more deeply concerned with *truth* and *knowledge* than with the other good-making features of doxastic states—including, for instance, the consequences of having those states. Call norms like this *epistemic norms*.

Moral encroachment is best understood as a claim about epistemic norms. If moral features of the world only made a difference to the *desirability*, or *moral goodness*, of having a belief, they would not play any notable role in epistemology. So defenders of moral encroachment must defend a distinction between epistemic and non-epistemic norms, and they must hold that epistemic norms are sensitive to certain moral facts. But they must also resist the extreme conclusion that epistemic norms are sensitive to *all* moral facts that, intuitively, have some bearing on belief. If they went that far, they would run the risk of losing their grip on the epistemic/non-epistemic distinction entirely; if even a demon's bribe can make the difference to whether a

---

<sup>41</sup> Throughout this chapter, I'll be neutral about the particular moral properties that apply directly to belief, using "good" and "bad" as placeholders for the terms licensed by the true first-order moral theory.



belief meets epistemic norms, epistemic norms seem to simply collapse into norms regarding the *desirability* or *goodness* of belief.

Defenders of moral encroachment, then, face a pressing question: if a moral bribe makes no difference to epistemic norms, which moral considerations do? The examples I've used in the previous chapters suggest one way to answer this question: certain morally-fraught choice situations make a difference to epistemic norms. Consider, for instance, a person (we'll call him César throughout the chapter, for ease of reference) who faces the maniacal-traffic-officer scenario from chapter 1:

**Parked Car High Stakes** César parked his car four hours ago, and he cannot currently see it. César's friend Maryam informs him that there is a maniacal traffic officer on the loose, and if the officer sees César's car parked illegally, he will fly into a homicidal rage and kill five innocents. César thinks back, and he seems to remember (although not too vividly) that he parked it legally. He forms the belief that his car is currently parked legally, and he remains sitting in his easy chair.

César's example illustrates how moral facts about one's actions and options might make a difference to the epistemic status of one's belief. On a moral-encroachment story of the sort I've told in the past two chapters, César's belief might lack epistemic rationality, and therefore fail to be knowledge, precisely because of the moral features of his choice situation.

César's case closely resembles prominent cases from the literature on pragmatic encroachment, as developed by Stanley (2005), Hawthorne and Stanley (2008), and Fantl and McGrath (2009). And some, including Sarah Moss (2018a, sec. 4), consider it a desideratum for moral-encroachment views that they be broadly continuous with

the broader, more familiar, phenomenon of pragmatic encroachment. But a different recent approach to moral encroachment radically departs from the pragmatic-encroachment literature. To see how, consider the following cases:

**Birdwatching Stereotype** Fatima’s friend tells her that a canary is in the next room. Fatima has strong, but not flawless, inductive evidence supporting the prediction that any given canary in her country will be yellow. She forms the belief that the canary in the next room is yellow.<sup>42</sup>

**Racial Stereotype** Aidan is a waiter at a restaurant. As he leaves work for the night, he crosses paths with a Black family entering the restaurant. He has strong, but not flawless, inductive evidence supporting the prediction that any given set of Black diners at his restaurant will give their waiters tips lower than 20%. On the basis of the family’s race, he forms the belief that they will leave one of his colleagues a tip lower than 20%.<sup>43</sup>

Fatima and Aidan base their beliefs on similar bodies of inductive evidence. But there seems to be an important moral difference between the two cases: while Aidan’s belief is a paradigm of racist reasoning, Fatima’s seems entirely morally unproblematic. Several philosophers have recently argued that the moral problems with Aidan’s reasoning can explain why his belief is also epistemically problematic.

The four cases we’ve just seen raise a key question for defenders of moral encroachment: *which* moral considerations make a difference for epistemic rationality? *Parked Car High Stakes* raises moral questions about *action*; César’s action is morally problematic, in a way that I’ve previously argued may explain his belief’s epistemic shortcomings. *Racial Stereotype*, on the other hand, does not

---

<sup>42</sup> I adapt this case from Moss (2018a, 220).

<sup>43</sup> I adapt this case from Basu (forthcoming).

obviously raise any questions about action. To the extent that we think that there is a moral problem with Aidan, it is not with his action, but with his character, his belief-forming practices, or his belief itself.

Some defenders of moral encroachment (including Renee Bolinger, Sarah Moss, and myself) claim that epistemic norms are sensitive to moral features of *actions* and *options*. I'll call this sort of sensitivity *moderate* moral encroachment. Others (including Rima Basu, Michael Pace, and Mark Schroeder) claim that epistemic norms are sensitive to moral features of *beliefs themselves*. I'll call this sort of sensitivity *radical* moral encroachment.

The goal of this chapter is to argue against radical moral encroachment while defending moderate moral encroachment. In section 1, I raise a challenge for all defenders of moral encroachment: they must explain why the moral considerations they cite are not *reasons of the wrong kind* within epistemology. In section 2, I show that defenders of moderate moral encroachment can meet this challenge. In section 3, I show that defenders of radical moral encroachment cannot. In section 4, I explain how we can approach cases like *Racial Stereotype* without taking on the unattractive commitments of radical moral encroachment.

### **Section 1: Reasons of the Wrong Kind**

This section introduces a distinction between reasons of the right kind (RKRs) and reasons of the wrong kind (WKR). I'll argue that we can use this distinction to make

headway in answering the core question of this chapter: which moral considerations, if any, make a difference to epistemic norms?

What does it mean to say that a reason is “of the right kind” or “of the wrong kind”? We first grasp this distinction through examples—usually, examples involving *incentives* for having a mental state. The fact that there is a poisonous snake next to me is a RKR to fear the snake. The fact that someone will pay me if I fear a teddy bear, by contrast, is a WKR to fear the teddy bear. The fact that a flight would bring me to an exciting destination is a RKR to desire to buy a plane ticket. The fact that Donna will punch someone in the face unless I desire to buy a plane ticket, by contrast, is a WKR to desire to buy a plane ticket.

Many have noted that there is a unified phenomenon here—a single distinction that applies to a host of mental states (including, for instance, fear and desire). And it’s striking that the cases with which I began this chapter, *Bribe for Belief* and *Bribe for Withholding*, seem to be paradigmatic instances of the phenomenon: more specifically, they seem to involve paradigmatic WKR. There are good grounds for thinking, then, that the difference between bribes for belief and paradigmatically epistemic reasons for belief is one instance of a general pattern: the difference between RKR and WKR.<sup>44</sup>

---

<sup>44</sup> See Way (2012, 491-2) and Schroeder (2012b, 458-9).

By focusing on RKR and WKR, we can reframe the debate about moral encroachment in epistemology.<sup>45</sup> The defender of moral encroachment claims that certain moral features bear not only on the desirability but also on the epistemic rationality of belief. She must explain why the moral features she cites, unlike moral bribes, are reasons of the *right* kind.

How can a theorist justify claims of this sort? How, in other words, can we determine whether a consideration is a RKR or a WKR? Broadly speaking, there are two methods. The first is the *method of analogy*. In order to determine whether some consideration is a WKR for belief, we can ask whether a consideration of that sort would be a WKR for a different mental state—including, for instance, emotion, desire, or intention. Of course, we should not erase important differences between types of mental states. Nevertheless, I'll show in section 3 that certain analogies provide powerful evidence about the scope of epistemic rationality.

The second method for answering questions about WKRs and RKRs involves appealing to a *theory* of the RKR/WKR distinction. We can gain evidence that a moral consideration is a WKR by showing that a promising theory classifies it as a WKR. Now, there are many existing theories of the RKR/WKR distinction, and I do not want to base my conclusions on any particular one: after all, perhaps the best theory has yet to be discovered! So, in what follows, I will not rely on any particular theory; instead, I'll appeal to the two most promising general approaches to the RKR/WKR

---

<sup>45</sup> Schroeder (2012a, 284-5) also frames the debate in this way.

distinction.<sup>46</sup> My arguments will show that, on *either* of these general approaches, we should reject radical moral encroachment.

The first promising approach to the RKR/WKR distinction is a *constitutivist* one. On a constitutivist approach, we can explain the difference between RKRs and WKRs for a given mental state by appealing to facts about *what it is* to be in that mental state. Take an example: fear seems connected, by its very nature, to the question of whether something is threatening or dangerous. And RKRs for fear seem, in a systematic way, to be considerations regarding danger. WKRs for fear, like bribes, are not connected in the same way to considerations regarding danger. Constitutivist approaches to the RKR/WKR distinction can be found in D'Arms and Jacobson (2000, 2006), Hieronymi (2005), Schroeder (2010), and Sharadin (2016).

The second promising approach to the RKR/WKR distinction emphasizes a putative asymmetry in *efficacy*. Generally speaking, it seems easier to form a mental state (or, perhaps, to *directly* form it) on the basis of an RKR than on the basis of a WKR. For example, it is easier to fear a snake on the grounds that it is poisonous than it is to fear a teddy bear on the grounds that one has been bribed to do so. Perhaps this asymmetry in efficacy points the way toward the correct general explanation of the

---

<sup>46</sup> Another theory identifies RKRs with “object-given reasons” and WKRs with “state-given reasons.” See Parfit (2001; 2011, App. A); for criticism, see Rabinowicz and Rønnow-Rasmussen (2006), Hieronymi (2005, 441–43), and Schroeder (2012b, 2013). I follow Nye (2017) in supposing that this is not the most promising approach to the RKR/WKR distinction.

RKR/WKR distinction. Proponents of this efficacy-based approach include Persson (2007), Raz (2009), and Rowland (2015).<sup>47</sup>

This second approach is often paired with a commitment to *WKR skepticism*: the view that there are, strictly speaking, no *reasons* of the wrong kind at all.<sup>48</sup> On this view, apparent wrong-kind reasons against a mental state are, at most, reasons for *wanting* to be in the mental state, or for *bringing the mental state about*. In what follows, I'll refer to certain considerations as 'reasons of the wrong kind,' but WKR skeptics should feel free to interpret these as references to, e.g., reasons for bringing a mental state about.

I do not aim, in this chapter, to settle the question of how we should theorize the RKR/WKR distinction. I aim, instead, to reach conclusions that are compatible with either of the most plausible approaches to that distinction. So, in what follows, I'll treat facts about *what it is* to believe (and to withhold belief) as potential evidence about the shape of the RKR/WKR distinction, and I'll also treat facts about *efficacy* as evidence. Section 2 shows that both of these approaches are nicely compatible with moderate moral encroachment. Section 3, however, shows that both approaches raise serious problems for radical moral encroachment.

---

<sup>47</sup> The best approach might be both constitutivist *and* concerned with efficacy; see, e.g., Hieronymi (2005).

<sup>48</sup> For defenses, see Kelly (2002), Parfit (2011, App. A), Skorupski (2007), Way (2012), and Rowland (2015).

## Section 2: Moderate Moral Encroachment and Reasons of the Wrong Kind

Defenders of moral encroachment hold that some moral considerations, like bribes for belief, are WKR within epistemology, but that some *other* moral considerations are RKR within epistemology. But should we believe that any moral reasons really are RKR within epistemology? And if so, which ones? In this section, I'll show that defenders of *moderate* moral encroachment are in a position to answer these questions successfully.

Recall the contrast between moderate and radical moral encroachment: defenders of moderate moral encroachment hold that norms of epistemic rationality are sensitive to facts about the moral status of one's *actions* and *options*. Defenders of *radical* moral encroachment go farther: they argue that norms of epistemic rationality are sensitive to facts about the moral status of one's *beliefs themselves*. Some defenders of radical moral encroachment also defend moderate moral encroachment.<sup>49</sup> But, for now, let's consider moderate encroachment alone.

Defenders of moderate moral encroachment are interested in choice scenarios like the ones illustrated by *Parked Car Low Stakes* and *Parked Car High Stakes*. They hold that, while being offered a bribe to believe (or withhold) does not make a difference to epistemic rationality, facing certain choice scenarios (like the one César faces in *Parked Car High Stakes*) can. To make this claim plausible, she must argue that a case

---

<sup>49</sup> See Schroeder (2012a, 2018).



like César's involves a RKR for withholding belief (or, put differently, for adopting higher evidential standards). I'll now argue that, on either of the most plausible approaches to the RKR/WKR distinction, the defender of moral encroachment is in a good position to make this argument.

On a constitutivist approach, RKRs for a mental state bear some important connection to facts about *what it is* to be in that mental state. Certain mental states, on this view, simply "bring with them" an evaluative standard or presentation.<sup>50</sup> Fear, for example, is constitutively concerned with danger, so RKRs for fear are considerations that have to do with danger. WKRs, like bribes to be afraid or amused, are notably disconnected from the core evaluative concerns of the mental states they favor.

At first, the constitutivist approach may seem to present a problem for moral encroachment. It's tempting to think that belief is constitutively concerned solely with *truth*.<sup>51</sup> This suggests a simple picture, on which evidence of truth, and nothing else, is an RKR in epistemology. If this simple picture is right, it's bad news for moderate moral encroachment: the fact that I face a certain choice is not (generally) evidence for the truth or falsehood of my beliefs.

An idea familiar from the pragmatic encroachment literature defuses this point. Though it's very plausible that *belief* is constitutively concerned with truth (or,

---

<sup>50</sup> See Sharadin (2015), or D'Arms and Jacobson (2000) on the core "evaluative presentation" or "concerns" of mental states.

<sup>51</sup> See, for instance, Wedgwood (2002); for a response, see Smithies (2012, sec. 6).

perhaps, evidence), it's much less plausible that the mental state of *withholding* belief is constitutively connected to truth in such a straightforward way.<sup>52</sup> Just what would it mean for a state of withheld belief to meet its *constitutive* standard for correctness?

At a first pass, withheld belief as to *p* seems to “bring with it” a concern for whether one has *enough* epistemic support for *p*.<sup>53</sup> But this first pass does not seem to rule out practical or moral considerations; in fact, some have suggested that practical and moral considerations are the only ones that could possibly give an informative answer to the question of how much epistemic support is *enough*.<sup>54</sup> So there is room in epistemology for constitutive standards that are sensitive to practical and moral considerations. I'll now sketch a positive story about the constitutive concerns of withheld belief—one that vindicates the presence of some, but not all, moral considerations in epistemology.

Many have observed that coarse-grained doxastic states (like belief, disbelief, and withheld belief) seem fit to play a role that finer-grained doxastic states (like credences or “degrees of belief”) cannot.<sup>55</sup> When I believe that *p*, I settle the matter as to whether *p*—at least provisionally, I commit myself to treating it as true. When I

---

<sup>52</sup> See Schroeder (2012a, 2013).

<sup>53</sup> NB: I am neutral as to whether withholding is a distinctive doxastic state. The key idea in the main text can be made without reference to withholding: the question of whether to have a belief about *p* is not merely constitutively concerned with *evidence*; it is constitutively concerned with the *sufficiency* of one's evidence *p*. Moral and practical concerns seem apt to make a difference to the question of whether one's evidence is *enough*. Thanks to Justin D'Arms for useful discussion.

<sup>54</sup> See Owens (2000, 25-6), Pace (2011).

<sup>55</sup> See, for instance, Ross and Schroeder (2014), Smithies (2012, sec. 4).

withhold belief that  $p$ , by contrast, I actively leave my view of  $p$  unsettled. By adopting coarse-grained doxastic states, in other words, I adopt a policy about how to treat a proposition in future reasoning.

This can teach us something about the constitutive standard for correctness in withheld belief. On this story, we evaluate withheld belief *qua* withheld belief, at least in part, by assessing whether it is apt to play its distinctive role in future episodes of theoretical or practical reasoning. In other words, the question of whether it's correct to withhold belief is intimately connected to the question of whether, by doing so, one takes up a mental state that will facilitate the projects of representing and navigating the world.

This story explains why it's correct to withhold belief in *Parked Car High Stakes*, but incorrect to withhold belief in *Bribe for Withholding*. In the latter case, withholding belief will have attractive downstream effects, but they have nothing to do with future episodes of practical or theoretical reasoning. In the former, by contrast, withholding belief is correct precisely because it's part of a mental scheme that is apt to play a particular role in helping César to reason well—specifically, it ensures that he will not inappropriately assume that his car is parked legally.

I've now sketched, in broad outline, a story on which coarse-grained doxastic states are constitutively concerned with practical and moral matters. The outline could be filled out in a number of ways; the crucial point is that moderate moral encroachment

seems entirely compatible with a constitutivist approach to the RKR/WKR distinction. In section 3, we'll see that the same cannot be said for radical moral encroachment.

Let's move on to the second promising general approach to the RKR/WKR distinction. This general approach emphasizes the asymmetry in *efficacy* between RKR and WKR; it distinguishes between RKR and WKR by noting the difficulty of adopting (or, perhaps, *directly* adopting) a mental state on the basis of a WKR. If, as moral encroachment suggests, some moral considerations are WKR and others are RKR in epistemology, then this approach suggests that we should see a noteworthy gap in the difficulty of responding to those considerations by forming new doxastic states.

Interestingly, we find just such an asymmetry between *Parked Car High Stakes* and *Bribe for Withholding*. To see this, imagine yourself in the former case. It would be very natural for you to respond to the news of the maniacal traffic officer by thinking, "Probably, my car is parked legally. But what if it's not? What if I'm misremembering, and because of my illegal parking, innocent people will be murdered?" This reasoning seems apt to naturally, and directly, facilitate withheld belief.<sup>56</sup>

Contrast this with a modified version of the case. In the modified version, you do not learn about the maniacal traffic officer; instead, you learn that a benefactor will give

---

<sup>56</sup> One might argue that this sort of choice scenario does not allow for a sufficiently *direct* or *straightforward* way of withholding to count as an RKR. But this claim takes up a heavy burden of proof; on the face of it, the withholding I've sketched is entirely straightforward.

money to charity if you withhold belief about whether your car is parked legally. In this modified version, it would not be nearly as natural to focus on the possibility that your belief is false. It would be more natural to focus on your belief itself, and on possible ways to change it. You might think, for instance, “Wow, it sure would be good if I stopped believing that my car is parked outside!” This reasoning seems less likely to directly facilitate withholding belief.

In short, being in a situation like *Parked Car High Stakes* tends to bring one to focus on the possibility that one’s belief is false. Being in a situation like *Bribe for Withholding*, by contrast, only makes salient the benefits of withholding. It’s very plausible that the former psychological state tends to facilitate withholding belief in a different way—a more natural way, and perhaps a more direct way—than the latter does.<sup>57</sup> Now, perhaps this difference in salience does not suffice to explain the asymmetry between the cases. But, regardless of the precise nature of that asymmetry, these two cases seem to involve an asymmetry of just the sort that many theorists take to be the core difference separating WKR from RKR. If an efficacy-based theory of the RKR/WKR distinction is on the right track, then, the defender of moderate moral encroachment will be in a strong dialectical position. She has

---

<sup>57</sup> If the salience of the possibility that  $\neg p$  generally brings with it a RKR to withhold, does the salience of the possibility that  $p$  generally bring with it a RKR in favor of belief? In short, no. Attending to one’s credence in the possibility that  $\neg p$  facilitates withholding belief regarding  $p$ , but merely attending to one’s credence that  $p$  does not directly facilitate believing that  $p$ . Thanks to Tristram McPherson for discussion.

evidence that, while a moral bribe for withholding is a WKR, certain choice situations (like César's) provide RKR in favor of withholding.

As the next section will show, the same cannot be said for defenders of radical moral encroachment.

### **Section 3: Radical Moral Encroachment**

#### *3.1 Against Radical Moral Encroachment*

In this section, I'll turn from moderate moral encroachment to radical moral encroachment. There is radical moral encroachment in epistemology just in case norms of epistemic rationality are sensitive to moral features of *belief itself*. My discussion will focus on a recently popular proposal, one that has been defended by both Rima Basu and Mark Schroeder.<sup>58</sup> Basu and Schroeder both claim that the moral badness of a belief itself can make a difference to the epistemic rationality of that belief. I'll argue against this approach, on the grounds that it cannot adequately distinguish between RKR and WKRs.

Why think that belief itself can be morally bad? Defenders of radical moral encroachment use a variety of examples to make this notion plausible. Some have to do with beliefs that undermine personal relationships; Basu and Schroeder (2019), for instance, describe a person who believes on inconclusive evidence that her spouse

---

<sup>58</sup> Michael Pace (2011) also defends radical moral encroachment. His proposal faces a particularly intense version of the WKR-related challenge that I pose in the main text.

has started drinking again. But the examples that are most frequently used to motivate radical moral encroachment involve beliefs based on inferences from statistics about demographic groups. In the *Racial Stereotype* case from the introduction, Aidan forms such a belief; he judges that the people entering his restaurant will leave a tip below 20%, solely on the basis of their race. Gendler (2011) offers a similar case involving racial profiling, and Schroeder (2018) offers a similar case involving sexist profiling.

Defenders of radical moral encroachment make two distinctive claims about their cases. First, these cases involve beliefs that are morally bad in a non-derivative way; the beliefs' moral badness does not depend, for instance, on the beliefs' downstream consequences, or on the believer's character. Second, epistemic norms are sensitive to the non-derivative badness of such beliefs. Armed with these claims, the defender of radical moral encroachment can use the morally problematic nature of a belief to explain its epistemic irrationality.

The first of these two claims is quite controversial, but I'll grant it for the sake of argument.<sup>59</sup> I'll argue that, *even if* some beliefs are non-derivatively morally bad, their moral badness does not make a difference to norms of epistemic rationality.

The easiest way to see this point is to consider an analogy with mental states other than belief. Generally, the fact that having a mental state would be non-derivatively

---

<sup>59</sup> For a defense, see Basu and Schroeder (2019).

morally wrong is a paradigmatic WKR. Consider two examples. First: some jokes are morally bad jokes, in the sense that there are moral reasons that count against anyone's being amused by them. Second: it's very tempting to think that there are often powerful moral reasons against envy. Further, these moral reasons need not arise solely in cases where there's nothing at all funny about a joke, or when the envied party has nothing worth desiring; in at least some cases, it's morally bad to be amused or envious even though, in some sense, amusement or envy is clearly appropriate. On the grounds of cases like these, it's widely believed that the mere fact that amusement would be morally bad is a WKR against amusement, and the mere fact that envy would be morally bad is a WKR against envy.<sup>60</sup>

Why? Recall the cases that inspire the RKR/WKR distinction in the first place: cases like *Bribe for Belief*. These cases cry out for a distinction between two ways of evaluating a mental state: we can evaluate a mental state for whether it is *all-things-considered good to have*, but we can also evaluate a mental state for whether it is *fitting* (or *correct*, or *rational*) in a narrower sense. Cases in which moral reasons count against emotions also cry out to be evaluated along two distinct lines. Even if we agree that it would be best if no one were amused by a joke, there is a second evaluative question that we have not addressed: is the joke funny?

---

<sup>60</sup> See D'Arms and Jacobson (2000) and Nye (2017). Within the dialectical context of this chapter, it's particularly noteworthy that Schroeder (2010) grants this point.



In short, a mental state's moral badness is typically a WKR. This provides evidence that the moral badness of a *belief* is, likewise, a WKR against having that belief. In other words, the moral badness of a belief does not bear on its epistemic rationality. So radical moral encroachment goes too far.<sup>61</sup>

We don't have to rely on analogy alone to see this point. On either of the most promising approaches to theorizing the RKR/WKR distinction, the moral badness of belief is a strong candidate to be a WKR. Consider, first, the constitutivist approach. As we saw in section 2, there is a promising way to explain why high-stakes choice scenarios are relevant to the constitutive standard of correctness for withhold belief. The defender of radical moral encroachment cannot tell a story of this sort. It's just not plausible that the core standard of correctness for withholding places any particular emphasis on the avoidance of morally bad psychological states. To the contrary, other mental states seem just as apt—and perhaps more apt—to be evaluated for their moral badness. This provides excellent evidence that, on a constitutivist approach to the RKR/WKR distinction, the moral badness of belief is a WKR.

Move on, now, to the efficacy-based approach to the RKR/WKR distinction. Here, again, the defender of radical moral encroachment is on shaky ground; noting that a

---

<sup>61</sup> Keen readers may wonder why I have not tested *moderate* moral encroachment via analogy. The answer is straightforward; there is no analogue in the realm of emotion or desire for the distinction between *all-in* belief and *degrees* of belief. And that distinctive role for coarse-grained doxastic states, as we saw in section 2, is a crucial part of the explanation of how moral factors impact epistemic rationality.

belief is morally bad is a comparatively *ineffective* route to withholding. This becomes particularly vivid when we compare a situation like *Racial Stereotype* with a situation like *Parked Car High Stakes*. As section 2 noted, being placed in the latter sort of situation naturally calls attention to the high-risk possibility that one's belief is false. It would be highly natural for César to wonder, "but what my car *isn't* parked legally? Then five innocent lives would be in danger!" Reactions of this sort, I've argued, naturally facilitate withholding belief. Attending to the possibility that one's belief is morally wrong, on the other hand, does not—perhaps, in part, because it does not tend to bring to mind the possibility that the belief is false. When I note that my belief is morally bad, I am apt to react in just the same way I would react if faced with a bribe for withholding: by thinking something like, "wow, it sure is important that I get rid of this belief!" The defender of radical moral encroachment, then, cannot lay claim to even a *prima facie* asymmetry in efficacy between cases of morally bad belief and cases like *Bribe for Withholding*. This is powerful evidence that, if an efficacy-based treatment of the RKR/WKR distinction is on the right track, the moral badness of belief is a WKR against it.

Taking stock: the method of analogy suggests that the moral badness of a belief is a WKR. And the evidence regarding *what it is* to withhold belief, along with the evidence regarding *efficacy* in withholding, also suggests that the moral badness of a belief is a WKR. This amounts to a powerful case against the notion that a belief's moral badness makes a difference to epistemic rationality.

### *3.2 Interlude: Why Go Radical?*

The debate over radical moral encroachment is not over; in subsection 3.3, I'll consider a way in which radical moral encroachers can avoid the problems I've raised so far. But, before we move on to consider that revision, it's worth pausing to ask about what motivates it. Why bother sticking with the radical moral encroachment hypothesis?

As we've already seen, defenders of radical moral encroachment are interested in cases where a belief seems both well-supported by evidence and also morally bad. They aim to make room for the claim that such beliefs are epistemically irrational. In the relevant set of cases, the thought is, it would be unacceptable for a person's belief to be both morally bad and also epistemically rational.<sup>62</sup> Radical moral encroachment, then, is primarily motivated by an interest in precluding the possibility of tension between a doxastic state's epistemic status and its moral status.

But this is a bad motivation. The defenders of moral encroachment have excellent reason to think that tension between a doxastic state's epistemic status and its moral status is not merely possible, but actual. To see this, consider a revised version of Aidan's case:

---

<sup>62</sup> Basu and Schroeder (2019) place a great deal of weight on the claim that this tension is problematic.

**Racial Stereotype 2** Aidan is a waiter at a restaurant. As he leaves work for the night, he crosses paths with a Black family entering the restaurant. He has evidence that suggests, to degree 0.8, that any given Black diner at his restaurant will give her waiter a tip lower than 20%. On the basis of the family's race, he adopts credence 0.8 that they will leave one of his colleagues a tip lower than 20%.

*Racial Stereotype 2* is morally worrisome in just the same way that the original *Racial Stereotype* case is. Aidan's updated credence constitutes a racist judgment, and a problematic one; if a Black diner became aware of Aidan's high credence, she could rightly complain, and she could rightly demand an apology. These points about blame and apology are just the considerations that defenders of radical moral encroachment tend to cite as evidence that beliefs can be non-derivatively morally bad. To the extent that we have reason to think that beliefs can be non-derivatively morally bad, then, we also have reason to think that credal states alone can be non-derivatively morally bad.<sup>63</sup>

Importantly, however, all parties should agree that Aidan's updated credence, in *Racial Stereotype 2*, is epistemically rational. The case simply stipulates that his evidence makes it likely to degree 0.8 that any given Black diner at his restaurant will leave a tip lower than 20%. If he refuses to bring his credences about individual Black diners in line with his evidence, he will be epistemically irrational. The defenders of radical moral encroachment, rightly, tend to grant this point: they suggest that cases

---

<sup>63</sup> Moss (2018a, sec. 2) makes a related point.

like *Racial Stereotype* make increased confidence (albeit not all-in belief) epistemically rational.

If this is right, however, the defenders of radical moral encroachment are committed to acknowledging a tension regarding *Racial Stereotype 2*: in that case, Aidan's credence could be both epistemically rational and morally problematic. And, as we've seen, there are good reasons for them to take on this commitment. But once we acknowledge that an epistemically rational credence can be morally problematic, we should be much less worried about the prospect that a belief might display just the same sort of tension.<sup>64</sup>

There are also independent reasons for thinking that beliefs can be both morally bad and epistemically rational: the tension between RKR's in favor of a mental state and moral reasons against it is an entirely general one. Sometimes, it's morally bad to envy someone else's possession, but the possession is nevertheless enviable. Sometimes, it's morally bad to have a positive aesthetic reaction to a work of art, but the artwork is nevertheless aesthetically impressive. Mature moral agents have to learn to navigate difficult situations like this: situations in which the moral reasons against an attitude are both powerful and reasons of the wrong kind.

---

<sup>64</sup> Buchak (2014, sec. 4) suggests that holding someone responsible involves forming beliefs (not merely credences) about her. But this point does nothing to motivate the idea that we cannot be held responsible *for* mere credences; at most, it suggests that we cannot hold others responsible *with* mere credences.

We'll now move on to consider a way of revising radical moral encroachment to address the WKR problem. I'll argue that this revision is unsuccessful on its own merits. But we should also worry about whether it is well-motivated. The primary motivation for refining a theory of radical moral encroachment is to avoid tension between the epistemic status and the moral status of a doxastic state. But, since defenders of radical moral encroachment are already committed to accepting that tension regarding credences, this is weak motivation indeed.

### *3.3 Radical Moral Encroachment Redux*

There is a way to develop radical moral encroachment that avoids the problems raised in section 3.1. The development involves two key moves. First, the defender of radical moral encroachment accepts that, when a moral reason against belief has nothing to do with that belief's truth or falsehood, it is a WKR. Second, she posits a class of moral reasons against belief that *are* intimately connected to the belief's truth or falsehood. Within some range of cases, she must argue, it would be morally bad to believe that  $p$  only if  $p$  were false.

Mark Schroeder (2018) defends a radical view of moral encroachment with just this shape. On Schroeder's view, the fact that a belief would wrong someone is a moral reason against holding it—but only a false moral belief can wrong someone.

Schroeder reaches his conclusion by appealing to three other commitments:

- (1) Only considerations that provide evidence for or against  $p$ , or that bear on the *cost of error* about  $p$ , can make a difference to the epistemic rationality of belief about  $p$ .
- (2) There is a set of cases,  $S$ , in which belief would be irrational, and the only viable explanation for the irrationality of belief appeals to the fact that the belief might morally wrong someone.
- (3) The fact that forming a belief that  $p$  might morally wrong someone does not provide evidence for or against  $p$ .

On the grounds of these commitments, Schroeder infers that the fact that a belief might wrong someone is (at least in the cases in  $S$ ) closely associated with the costs of error—in other words, the costs of believing falsely. He then suggests a general explanation for the required connection between morally wronging belief and false belief: only a false moral belief can wrong someone.

The claim that a belief's moral badness depends on its falsehood is counterintuitive. Insofar as we are tempted by the thesis that beliefs can wrong others, we generally do not think that the question of whether they do so hinges on their truth or falsehood. We can think that Aidan wrongs the family entering his restaurant by forming his racist belief about their tipping practices, for instance, without our judgment being sensitive in any way to the question of whether his belief is false.

Since Schroeder motivates his counterintuitive conclusion through several controversial assumptions about the ethics of belief, it's tempting to apply a Moorean shift here, using the implausibility of Schroeder's conclusion to reject one of the commitments with which he supports it. Schroeder is sensitive to this, and he therefore attempts to debunk the intuition that his conclusion is false. He does so by

drawing a distinction between two ways in which we can morally evaluate a person's belief: we can ask whether the belief is *objectively* bad, or on the other hand, whether it is *subjectively* bad. People whose beliefs are true, Schroeder suggests, have not wronged anyone, and their beliefs are therefore guaranteed not to be morally bad in an *objective* sense. But this does not mean that every true belief is morally acceptable in a *subjective* sense. Perhaps, just as it is subjectively morally bad to poison a random meal in a crowded cafeteria, even if (by good fortune) no one ends up eating it, it is subjectively morally bad to form certain beliefs on the basis of racial stereotypes, even if those beliefs (by good fortune) end up being true. By leaning on this distinction, Schroeder makes room for the claim that there is something morally bad about a belief like Aidan's, even though only false beliefs wrong.

At first, the distinction between subjective and objective moral evaluation might seem to give Schroeder all the argumentative fuel he needs to push back against the Moorean shift. If his conclusion follows from an otherwise attractive picture moral encroachment, and there's a viable approach to the ethics of belief on which his conclusion is not so counterintuitive, then perhaps his argument should persuade us to endorse that approach to the ethics of belief.

But there are reasons to worry about the way that Schroeder applies the distinction between subjective and objective moral evaluation. To bring this out, I'll note a general feature of objective moral evaluation: even when she knows that one of her



past actions, *A*, was subjectively bad, a virtuous person will have a disposition to feel relief upon learning that *A* was not also objectively morally bad.

Consider an example:

**Deathbed Promise** As a benighted youth, Duane was inadequately attentive to his grandmother. After she passed away, he was not sure of whether he had made her a deathbed promise: to put flowers on her grave on October 1<sup>st</sup>, 1992. But when October 1<sup>st</sup>, 1992 arrived, rather than trying to determine whether he really did make the promise, Duane decided to stay home and play video games rather than putting flowers on her grave.

Duane has now grown up, and he has become a virtuous person. He learns that he did not actually make his grandmother this deathbed promise.

The moral badness of young Duane's action has more to do with the way he acted *given his evidence* than with the way he acted *given all the facts*. In other words, his action is easier to criticize as *subjectively* morally bad than as *objectively* morally bad. Had he actually made the deathbed promise, his action would have been morally bad in an objective sense as well. In this case, I suggest, Duane might well be disposed to feel relief when he learns that he never *actually* broke a deathbed promise. Perhaps that disposition would not be activated; perhaps, for instance, it would be overwhelmed by his sense that his action was subjectively morally bad. Nevertheless, it would surely be sensible if Duane had the sense of having escaped doing something that was morally bad in an importantly different way.

The problem is this: in the range of cases that motivate radical moral encroachment in the first place, a virtuous person would *not* be disposed to feel relief if her belief

turned out to be true. Return to Aidan's case: suppose that, after forming his racist belief about the diners entering his restaurant, he becomes a virtuous person, and he also learns that his racist belief was true. In this case, I suggest, Aidan would not have any disposition to be relieved. He would regard the diners' actual tipping as irrelevant to his moral self-assessment.<sup>65</sup>

This provides evidence that Schroeder's debunking maneuver falls flat. If his application of the subjective/objective distinction were apt, we would regard true racist beliefs, roughly, like we regard actions that narrowly avoid breaking promises. But, morally speaking, forming a true racist belief is more like *actually* breaking a promise than like narrowly avoiding breaking a promise. So, even in the face of Schroeder's debunking story, there are good reasons to be suspicious of the claim that the moral badness of racist beliefs like Aidan's has something to do with the possibility that they are false.<sup>66</sup>

---

<sup>65</sup> Some readers might feel that there is something preferable about Aidan's racial profiling if the diners actually leave a tip below 20%. I think this is misguided; generally, inaccurate racial profiling is just as bad as accurate racial profiling. But even if the badness of some racial profiling can be mitigated by accuracy, some surely is not. Take, for instance, a racist belief that someone has a genetic predisposition toward low intelligence. Surely, having such a predisposition does not make one fair game for racial stereotyping. So the defender of radical moral encroachment should not rest easy with the notion that the moral wrongness of racial profiling always comes from the risk of inaccurate profiling.

<sup>66</sup> A defender of Schroeder's view might argue: "it's a striking fact that no morally bad beliefs are *guaranteed* to be true by the believer's evidence. The best explanation of this striking fact is that the moral badness of belief is rooted in the risk of falsehood." But this striking fact is equally well-explained by the hypothesis that it's morally important to avoid certain *inadequately supported* beliefs. (Note, too, that if the badness of racist belief does *not* hinge on its falsehood, it is plausibly a WKR. Compare to the moral reason that arises if a demon threatens to murder five innocents unless you withhold belief about *p*, and he does so on the grounds that your belief is not guaranteed to be true by your evidence.) Thanks to Tristram McPherson for discussion.

To sum up: by positing a connection between a belief's moral badness and its falsehood, the radical moral encroacher makes it more plausible that a belief's moral badness is a RKR. But she also signs up to implausible claims about the source of moral badness in beliefs. Of course, if radical moral encroachment were well-motivated on independent grounds, this cost might be bearable. But in section 3.2, we saw that the primary motivation for radical moral encroachment is no motivation at all. So it makes sense to respond to the many challenges that face radical moral encroachment not by further refining the theory, but instead to look for the best available alternate theory. In this chapter's final section, I'll do just that.

#### **Section 4: Bad Beliefs Without Radical Moral Encroachment**

This chapter aims to show that, although we can safely accept moderate moral encroachment, we should not accept radical moral encroachment. So far, I've been making the latter point by showing that radical moral encroachment commits us to an unattractive normative theory: either it draws the RKR/WKR distinction poorly, or it locates the moral problem with bad beliefs in the wrong place. In this final section, I'll take a different approach: I'll note some alternate treatments of the cases that motivate radical moral encroachment. If these cases do not require us to adopt radical moral encroachment, and radical moral encroachment is also both ill-motivated and beset with problems, we can comfortably reject it.

As I mentioned in section 3, the cases that are most frequently cited by defenders of radical moral encroachment are structurally similar to *Racial Stereotype*. They

involve beliefs about particular individuals that are based on information about statistical regularities. Many such cases seem morally problematic, and many also seem to involve epistemic irrationality. Can we explain the irrationality of beliefs like these without appealing to radical moral encroachment?

In the vast majority of cases, I think that we can. Most regularities that hold within demographic groups in modern societies, especially the ones that are most likely to be cited by bigoted thinkers, are remarkably weak. What's more, most people have plenty of evidence to this effect. When a person sincerely avows the belief that some enormous percentage of a demographic group shares a trait of any importance, we should suspect that she's approaching her evidence in a flawed way. So, in most real-life cases of beliefs based on putative statistical regularities, there's no puzzle as to why the beliefs are epistemically irrational; they are based on assumptions that are ill-founded, irrational, or wildly inaccurate.<sup>67</sup>

What should we say, though, about the rare cases in which there really is strong evidence of a demographic regularity? Plausibly, in some cases of this sort, there is moral pressure against forming judgments about particular individuals based on these regularities. I'll now survey two ways in which we could interpret this pressure without taking on the worrisome costs of radical moral encroachment. On the first approach, the pressure is both moral *and* epistemic. On the second, the pressure is moral alone. Throughout, I'll illustrate the views at hand by discussing *Racial*

---

<sup>67</sup> Gardiner (2018) makes related points.

*Stereotype*, and simply stipulating that Aidan's evidence genuinely does make it very likely that any given Black diner will leave a tip below 20%.

First, perhaps appeals to *moderate* moral encroachment are sufficient to explain why Aidan's belief is epistemically irrational. Recall that, on moderate approaches to moral encroachment, a belief's rationality depends on certain moral facts having to do with *actions* or *options*. Moss (2018a, sec. 4) and Bolinger (2018, sec. 4) have both applied this view to cases like *Racial Stereotype*. Both suggest that, when we adopt certain beliefs based on statistical generalizations about demographic groups, we immorally risk relying on those beliefs in action, and thereby contributing to pernicious shared social practices.

This approach faces two initial problems. One has been noted by the proponents of radical moral encroachment: in some cases like Aidan's, there does not seem to be any risk that the relevant belief will inform any future action.<sup>68</sup> Aidan forms his belief while leaving work, and even if he bumps into the family of diners again, he will surely not remember them. Why think that, by forming his belief, he imposes on them a risk of any kind?

The second problem for this approach is similar to the problem that I posed for radical moral encroachment in section 3. Even if we grant that a belief like Aidan's may dispose him to act badly, this possibility doesn't seem closely connected to the *truth*

---

<sup>68</sup> See Schroeder (2018, sec. 3).

or *falsehood* of that belief. To see this, suppose that Aidan reasons as follows: “It’s very likely that this family will leave a tip below 20%. But what if I act on the expectation that they are low tippers, but they turn out to be high tippers? *Then* my action would be morally problematic!” Here, Aidan seems to be assuming that it is morally acceptable for him to act in a certain way toward the family, *unless* they actually leave a tip below 20%. But this is a bad assumption: the moral status of his action does not depend on the family’s actual tipping practices.

The point generalizes: in general, the most serious moral problems with racial profiling do not depend on whether the profiling in question is accurate. It’s morally important that we put a stop to patterns of behaviors based on expectations about members of oppressed groups. But it’s no less important to do so when our expectations turn out to be accurate than when they turn out to be inaccurate. For instance, people who have never spent time in jail deserve not to be treated as felons solely on the basis of their race. But felons also deserve not to be treated as felons solely on the basis of their race.

This is a problem for the claim that Aidan has a RKR for withholding belief. As we saw in sections 2 and 3, we should prefer a view on which the moral reasons that bear on epistemic rationality are intimately tied to the risk of falsehood. This is the most promising way to distinguish between cases like *Parked Car High Stakes* and *Bribe for Withholding*. But the most noteworthy problems associated with cases like *Racial Stereotype* are not associated with the risk of acting on the basis of stereotypes *when*

*they do not hold*; instead, they're associated with the risk of acting on those stereotypes *at all*.

So there are reasons to think that, even if there is moderate moral encroachment in epistemology, it does not extend to *Racial Stereotype*. Now, perhaps this initial challenge can be handled. Moss (2018a) briefly suggests that the moral badness of acting as if someone has a statistically prevalent trait is indeed distinctively serious when she lacks that trait. Perhaps this is right. But note that, for this proposal to be made good, this distinctive badness would have to be *comparably weighty* to the moral badness that comes from acting on the basis of racial profiling in the first place. If the latter moral badness simply settles all questions of how to act, after all, the risk of error makes no difference to the policies it's best to adopt for future episodes of practical reasoning.

There are reasons to worry, then, that moderate moral encroachment cannot establish that *all* cases like Aidan's involve epistemically irrationality. I therefore want to offer an alternative approach—a second position that does not require us to take on the unattractive commitments of radical moral encroachment. On this second approach, the vast majority of beliefs like Aidan's are epistemically irrational for banal reasons: they are based on spurious evidence, bad theory, projection errors, or irresponsible motivated reasoning. This approach also grants that, in some cases, questions about how to treat a person might hang on whether she actually fits a

particular demographic trend; in those cases, *moderate* moral encroachment can explain why all-in belief is epistemically irrational.

In the rare cases where neither of these explanations is available, however, this second approach simply grants that the belief could be epistemically rational. Importantly, this is not to say that the belief is morally kosher. To the contrary, this second approach explicitly embraces the possibility of a tension between the moral status of a belief and its epistemic rationality. As I argued in section 3.2, this is no cost to the theory; there can be tension between epistemic rationality and moral norms when it comes to credences, and there can be tension between the RKR that favor an emotion and the moral reasons against it. It should be no surprise that this tension afflicts belief as well.

I'll close by considering an objection: doesn't this view let believers like Aidan off the hook?<sup>69</sup> One way to make this objection more precise is to lean on the notion that WKR are comparatively inefficacious. When we accept that a belief's moral badness is a WKR, we may thereby imply that the belief is difficult to abandon. And the fact that meeting a moral demand would be very difficult sometimes mitigates blame. So

---

<sup>69</sup> Basu and Schroeder (2019) offer another criticism of views that allow this tension: they note that it would not be much of an apology to say "I'm sorry for believing... even though my belief was epistemically impeccable, short of being true." But Basu and Schroeder's view (on which there are very few positive epistemic duties) makes room for a much better sort of apology: "I'm sorry for believing. My belief was one reasonable response to the evidence, but there was another equally reasonable response available to me, and it would've been much more decent to you." Perhaps believers like Aidan, when their beliefs are epistemically rational, are called upon to offer apologies of this sort.



it may seem that, by allowing that beliefs like Aidan's might be epistemically rational, I wrongly imply that Aidan might deserve little blame.

I'll make two points in response to this worry. First, some—especially those who worry about the sort of control we have over our beliefs—may want to say that some sorts of moral badness (say, perhaps, *viciousness*) do not presuppose agential control, or aptness for blame. My discussion leaves open the possibility that, though Aidan cannot be blamed for his belief, his belief is still very seriously morally bad in some other sense.

Second, those who are inclined to make room for blaming Aidan can certainly do so. Though withholding belief *on the basis of* a moral consideration is indeed distinctively psychologically difficult, *getting oneself* to withhold belief regarding an uncertain proposition is often not difficult at all. Getting oneself to withhold is, generally, nowhere near as difficult as getting oneself to believe against the evidence. If Aidan claimed, "I'm trying to abandon the belief that this diner will leave a tip below 20%, but I'm just having such a hard time keeping an open mind," we would generally not accept his claim as an excuse.

Throughout this chapter, I've argued that the moral badness of a belief does not make a difference to its epistemic rationality. Some have taken cases like *Racial Stereotype* to provide evidence to the contrary. In this final section, I've cast doubt on the evidential force of those cases by noting other available ways to interpret them.

In conclusion, we need not embrace radical moral encroachment; what's more, by rejecting it, we can avoid a host of problems. The problems I've raised for radical moral encroachment, however, are not shared by moderate moral encroachment. Certain moral facts, then, may indeed play a surprising and important role in setting epistemic standards. But the fact that a belief would be morally bad to hold is not among them.

## Chapter 4

### Knowledge and the Many Norms on Action

In this chapter, I add detail to my picture of an impurist epistemology. I add that detail through the framing device of posing a challenge to impurism, and sketching the form for the best answer to that challenge. The challenge, very roughly, is to explain just how some norms on action, but not others, are relevant for epistemic evaluation. The answer that I favor, very roughly, is that each of several norms independently brings its own force to bear on epistemic evaluation, and each one does so in virtue of the role that it occupies in our social lives.

In section 1, I lay out the challenge in further depth and explain why any acceptable version of impurism must address it. In section 2, I defend two desiderata for a response to the challenge. In section 3, I show that contemporary versions of impurism do not have the tools to offer a response that meets those desiderata. To the extent that contemporary impurists address my challenge, they tend to do so by citing standards of ‘rationality’ (or ‘practical rationality’). No matter how we understand these appeals, they lead to unacceptable results for impurist views.

Having shown that contemporary impurists are in trouble, I spend section 4 discussing ways for them to get out of trouble. This section sketches the approach that I favor to impurism, and rules out some other approaches (including the one found in Hawthorne and Stanley 2008).

### **Section 1: The Challenge, and Why It Needs Answering**

Throughout the dissertation, I've illustrated impurism by offering examples of characters who face fraught choice situations. Think back, for instance, to the example of César, who stays in his easy chair rather than going to check on his car even though five innocent lives hang in the balance. Impurism makes room for the following claim about César: the fact that it's unacceptable for him to rely on the proposition that *his car is parked legally* explains why he is not in a position to know, or to rationally believe, that proposition.

I think a view in this vein is on the right track. But so far, this is just a suggestive sketch, and a fully satisfactory impurism would have to offer a more precise formulation of the connection between knowledge and norms on action. It'll be useful to distinguish between two burning questions for impurists, like myself, who aim to do so.

First: I've talked about whether César is in a position to *rely on* the proposition that the bank is open tomorrow. But what does it mean to *rely on* a proposition? Suppose, as many impurists claim, that César can know a proposition only when it's

appropriate for that proposition to play some role in her action. Well, then: what is that role? There has been a great deal of writing on this topic, and impurists suggest many different candidate roles. Fantl and McGrath (2009, 66), for instance, call attention to the question of whether a proposition can play a *justifying* role. Hawthorne and Stanley, by contrast, focus on whether “it is appropriate to treat the proposition... as a reason for acting” (2008, 578). Other precisifications abound.<sup>70</sup> Since the challenge that I offer afflicts all versions of impurism, I won’t say more about this first question. In what follows, I’ll suppose that the impurist can successfully answer it, and I’ll use talk of *relying on a proposition* as a theory-neutral placeholder for the best answer.

My challenge concerns a second, more frequently neglected, question about how to flesh out impurism: what does it mean for a person’s relying on a proposition to be *acceptable*? There are many normative standards against which behavior can be measured, and therefore many ways to answer this question. The impurist might say that César can know that his car is parked legally only if it would be acceptable, *morally* speaking, for to rely on that proposition. Or only if it would be acceptable, *prudentially* speaking. Or only if it would be acceptably *polite*. The list goes on. Which of these assessments of action make a difference for knowledge, and why?

---

<sup>70</sup> Fantl and McGrath (2002, 77) focus on whether an agent is “rational to prefer as if *p*”; in their (2007, 559), they instead foreground the condition of being “rational to act as if *p*.” Two other examples: Ross and Schroeder (2014, 272) ask whether it is acceptable to “treat *p* as true in [one’s] reasoning”, and Locke (2014, 43; cf. Hawthorne 2004, 30) draws attention to whether it is acceptable to “premise” a proposition in the situations where one is disposed to do so.

I've now posed my challenge to impurists: they must say which norms on action are distinctively connected to knowledge, and they must explain why. But why must the impurist address my challenge? In short, because without the right answer, the impurist's theory will have unacceptable results. I'll close this section by surveying two such results: outright contradiction and unacceptably skeptical verdicts.

First, consider the threat of outright contradiction. This threat is particularly pressing for impurists, like John Hawthorne and Jason Stanley, who posit a *biconditional* link between knowledge and action. According to Hawthorne and Stanley,

Where one's choice is  $p$ -dependent, it is appropriate to treat the proposition that  $p$  as a reason for acting iff you know that  $p$ . (2008, 578)

The trouble for this biconditional claim is that, in some cases, it is appropriate according to one norm, but inappropriate according to another, to treat the very same proposition as a reason for acting. To see this, consider a variant on the high-moral-stakes cases I've offered throughout the dissertation:

**Naomi's Medical Supplies** It is Friday afternoon, and Naomi is heading home from work. Her boss has asked her to put a bag of medical supplies in the mail by Saturday. She knows that, unless the medical supplies are in the mail by Saturday, five innocents who need treatment will die. But Naomi is radically morally apathetic: it does not matter to her whether the innocents live or die. Naomi sees that there is a long line at the only post office where she can put the supplies in the mail. If she waits in line, it will make her late to a dinner party that does matter a great deal to her. Remembering that the post office was open on Saturday a few weeks ago, Naomi decides to pass the post office by and to come back tomorrow to put the supplies in the mail.

Suppose that, when she passes the post office by, Naomi treats the proposition *that it will be open tomorrow* as a reason for action. Is it acceptable for her to do so?

Well, some norms surely call her behavior unacceptable. For instance, Naomi's action seems to violate *moral* norms, and perhaps also *all-things-considered* norms, on action. Given the weakness of her evidence that the post office will be open tomorrow, it's unacceptable for Naomi to rely on that proposition. The claim that it's (morally) unacceptable for Naomi to rely on the proposition that the post office will be open tomorrow, in conjunction with Hawthorne and Stanley's biconditional knowledge-action link, entails that she does not know that the post office will be open tomorrow.

On the other hand, some norms call Naomi's behavior acceptable. Though it's quite clear that Naomi's behavior violates moral standards, it's far less clear that there's anything *imprudent* or *irrational* about her action. Another way of putting the point: although the *moral* stakes in this case are high, the *prudential* stakes (and perhaps also the *rational* stakes) are low. The claim that Naomi's pattern of practical reasoning is (prudentially) acceptable, in conjunction with Hawthorne and Stanley's biconditional, entails that she knows that the post office will be open tomorrow.

We've reached a contradictory result: Naomi both knows and doesn't know that the post office will be open tomorrow. And it's clear what must be done to avoid this result: impurists like Hawthorne and Stanley must get more specific about the norm on action operative in their biconditional knowledge-action links.

Now, not all impurists defend biconditional knowledge-action links. Some defend the weaker claim that knowledge is a sufficient, but not a necessary, condition on acceptable action. Take, for example, a principle defended by Jeremy Fantl and Matthew McGrath:

If you know that  $p$ , then  $p$  is warranted enough to justify you in  $\phi$ -ing, for any  $\phi$ . (2009, 66)

This principle does not run the risk of licensing contradiction in the same way that a biconditional one does. Although one can apply *modus tollens* to Fantl and McGrath's principle to determine that an agent lacks knowledge, it cannot be similarly used to show that an agent has knowledge.

Defenders of principles like Fantl and McGrath's, however, are still obliged to answer my challenge about the many norms on action. Unless they do so, they run the risk of unacceptably skeptical results. Why? Well, simply put, there are a great many norms on action, including merely stipulated ones. If one must be able to act acceptably according to *every* norm in order to have knowledge, knowledge will be rare indeed.

To see this point, consider a further case:

*Kayla's Shirt* Kayla is deciding whether to do laundry for tomorrow. She decides not to, on the grounds that she (somewhat vaguely) remembers seeing a clean white shirt in her dresser two days ago. Unbeknownst to her, there is a club on the other side of the globe called the Clean White Shirt Group. According to the bylaws of the club, everyone must always have a clean white shirt in his or her dresser. The group's rules explicitly state that, whenever someone is not rationally certain that there is a clean white shirt in her dresser, that



proposition is insufficiently warranted to justify her in acting. As a matter of fact, Kayla does have a clean white shirt in her dresser.

Kayla's evidence that there is a clean white shirt in her dresser is fairly strong, but it is not strong enough, according to the bylaws of the Clean White Shirt Group, for that proposition to justify her actions.<sup>71</sup> So, if knowledge that  $p$  entails that relying on  $p$  will be acceptable even according to merely stipulated norms like the bylaws of the Clean White Shirt Group, Kayla cannot know that there is a clean white shirt in her dresser. This is absurd.

Moreover, this threat seems to generalize; for almost any item of ordinary knowledge held on the basis of a non-maximally-strong epistemic position, we can cook up a set of norms that threaten that knowledge. Unless the defender of impurism ensures that the ability to meet standards like these is not required for knowledge, her view will have inappropriately skeptical results.<sup>72</sup>

We've now seen why impurists must be specific about the norms on action that are distinctively connected to knowledge. In the next section, I'll defend two desiderata for any answer to this challenge.

---

<sup>71</sup> Suppose, for any given way of precisifying 'rely on,' the bylaws of the Clean White Shirt Group forbid Kayla to rely on the proposition that there is a clean white shirt in her dresser, and that they do because of the weakness of her epistemic position.

<sup>72</sup> You might think that the norms of the Clean White Shirt Group cannot threaten knowledge simply because the group does not exist. But even if such a group arose tomorrow, I could still know about the clean white shirt in my dresser. Indeed, even if the group is passing out fliers on my street, it's plausible that they do not threaten my knowledge of the shirts in my dresser.

## Section 2: Desiderata

First, impurists should offer an account on which knowledge is sensitive to norms of practical coherence.

Before defending this desideratum, I'll say a bit about the norms I mean to pick out with the term 'practical coherence.' These are, roughly, norms that evaluate whether one pursues one's final ends sensibly, regardless of which final ends one has. To the extent that an agent adheres to these norms, the mental states involved in her action cohere with one another. Her plans and her actions make sense given her goals, her desires, and her credences. She takes (what, on her evidence, are) the necessary means to her most important goals. She takes the actions that she intends. But, importantly, an agent cannot violate these norms *merely* in virtue of having any particular intention or desire.

The following case will help to show why impurists should connect practical coherence to knowledge.

*Tracy's Cigarettes* Tracy has a standing intention not to run out of Wolverines, a particularly dangerous and addictive brand of cigarettes. This is her most firmly held, most devout intention, and she fully endorses it. In terms of the overall success of her current plans and goals, running out of Wolverines constitutes the maximally disastrous scenario.

On her way home from work on Saturday, Tracy passes the only drugstore that sells Wolverines. She knows that, if she does not buy more Wolverines by the end of the day tomorrow, she will run out. She sees that there is a long line at the drugstore. Remembering that the

drugstore was open on Sunday a few weeks ago, she decides to pass the store by and to come back tomorrow to buy Wolverines.

We can suppose that, in this case, replenishing Tracy's stock of Wolverines would be bad for her well-being. We can even suppose that she knows this. This alone does not mean that Tracy has fallen short of the norms of practical coherence; she only falls short of the norms of practical coherence insofar as her plan to buy Wolverines fails to cohere with other plans or desires, or she fails to go about buying her cigarettes in a sensible way.<sup>73</sup>

Nevertheless, the case as described seems to involve a failure of practical coherence. To see this, note that a friend might challenge Tracy's behavior by asking, "why are you passing the drugstore by? Given how important it is to you to get those cigarettes, shouldn't you just wait in line?" And Tracy might respond by attributing knowledge to herself. She might answer, for instance, "what's all the fuss? I know that the store will be open tomorrow, so my decision to come back then is totally sensible."

Importantly, this is just the sort of exchange commonly cited by defenders of impurism to illustrate and to defend their view.<sup>74</sup> Knowledge, according to many impurists, puts a person beyond the reach of some (although, importantly, not all) sorts of criticism. And it's very plausible that, if Tracy were to use a knowledge-attribution to portray her action as beyond the reach of her friend's criticism, she

---

<sup>73</sup> You might worry that any plan that is knowingly imprudent must involve practical incoherence, and so that the case as presented is impossible. For more on this worry, see section 3.2.

<sup>74</sup> See, for instance, Hawthorne and Stanley (2008, 571-2); Fantl and McGrath (2007, 561-4; 2009, 60).

would be aiming to avoid criticism *in terms of practical coherence*. After all, Tracy acknowledges that her plan is imprudent. So the question of whether she has enough evidence to act *prudently* is not at issue between her and her friend. The question at issue, rather, is whether she has enough evidence for her actions and intentions to cohere with her evidence. Insofar as impurists take our practices of knowledge-attribution at face value, then, they have reason to suppose that knowledge is sensitive to the norms of practical coherence.

The foregoing remarks give some support to the notion that impurists should accept our first desideratum. And we don't have to look far for further support. Impurists frequently offer arguments that are easiest to read as attempts to establish connections between knowledge and practical coherence. Many of the most influential arguments for impurism, for instance, explicitly claim that they are evaluating subjects' actions within a decision-theoretic model.<sup>75</sup> It's easiest to read this language as aimed at picking out norms that do not evaluate a subject's final ends. Moreover, these arguments often draw on examples that are entirely neutral regarding the value of subjects' final ends. Take two examples: Fantl and McGrath (2009, 48-9) discuss an agent who aims to make it to Foxboro on time, and Hawthorne and Stanley (2008, 571) discuss agents who aim to find their way to a restaurant. The plausibility of these examples does not hinge on whether the agents' aims are

---

<sup>75</sup> Hawthorne and Stanley (2008, 583); Fantl and McGrath (2002, 75); Fantl and McGrath (2009, 76); Locke (2014, 47-50).

worthwhile; they serve, rather, to bring out the question of whether the agents are pursuing their aims in a sensible way.

Arguments in favor of impurism, then, are frequently executed in terms that suggest a focus on practical coherence. Moreover, they are no less plausible when understood as picking out norms of practical coherence. Insofar as the current literature gives us reason to think that impurism is true, we have reason to think there are connections between knowledge and practical coherence.

Second, impurists should offer an account on which knowledge is sensitive to the norms of morality; in other words, they should provide for moral encroachment. This is precisely the conclusion defended with the first two chapters of this dissertation, and I'll take it for granted in what follows. The question that we'll face in the remainder of this chapter, then, is how impurism can best provide for both moral encroachment and encroachment from practical coherence norms.

### **3. Appeals to Practical Rationality**

This section shows that, insofar as actual impurists address the challenge I've raised, they fail to meet our desiderata. The problem is this: to the extent that impurists refer to particular bodies of norms on action, they tend to simply cite 'rationality' or 'practical rationality.'<sup>76</sup> But, on any plausible theory of practical rationality,

---

<sup>76</sup> The language of 'rationality' or 'practical rationality' predominates, for instance, in Stanley (2005), Weatherston (2005), Fantl and McGrath (2002; 2007), Hawthorne and Stanley (2008), Ross and Schroeder (2014), and Locke (2015).

connections between knowledge and practical rationality do not suffice to meet both of our desiderata in a satisfying way. In this section, I'll survey three approaches to practical rationality, and I'll show why none of these approaches allows the impurist to avoid trouble.

### *3.1 Rationality as Practical Coherence*

Some use the term 'practical rationality' to pick out the notion that I've called 'practical coherence.'<sup>77</sup> On this picture, an agent is practically rational just to the extent that the mental states involved in her action cohere with one another.

As I noted in section 2, it's plausible that many impurists do have something like practical coherence in mind when they appeal to practical rationality. Section 2 also mentioned the problem with this approach: it fails to establish the right sort of connections between knowledge and *moral* norms. Morally-apathetic Naomi illustrates the issue: if Naomi's epistemic position is sufficiently weak, there is a serious normative problem with her action. But the norms of practical coherence cannot capture this problem. If Naomi is coherently apathetic about whether the medical supplies are mailed tomorrow, her evidence about the post office's hours is simply irrelevant to questions about how to pursue her final ends in a sensible way.

---

<sup>77</sup> Scanlon (1998, 25-30) and Broome (2007) argue in favor of using 'rationality' to pick out a property closely associated with coherence. Some simply stipulate that they will use 'rationality' in this way; see, for instance, Kolodny (2005, 509-10) and Southwood (2008, 9-10)

It's worth considering a complication: there are metaethical views on which being practically coherent simply entails following moral norms.<sup>78</sup> If a view of this sort is correct, then I've misdescribed Naomi's case, because anyone who violates a moral norm thereby violates a (perhaps more fundamental) norm of practical coherence. This opens the door to the possibility that, when a friend criticizes Naomi for lacking sufficient evidence, she really means that Naomi lacks sufficient evidence to pass the post office by *while meeting standards of practical coherence*. So if impurists' appeals to practical rationality are to be understood as appeals to norms of practical coherence against the background of this metaethical theory, they may indeed meet both of our desiderata.

I'll come back to this strategy in section 4. There, I'll sketch three available paths forward for the impurist. One of the three involves taking on a metaethical commitment of just this sort: a commitment to the underlying unity of norms of practical coherence and moral norms. I'll suggest that this is not the most promising way forward for the impurist—loosely speaking, because it would be an unfortunate surprise for impurists if the prospects for their view depended on the truth of a highly controversial metaethical view. For now, let's move on to see why two other approaches to practical rationality cannot give the impurist all she needs.

### *3.2 Rationality as Prudence*

---

<sup>78</sup> Korsgaard (1996, ch. 3) defends such a view.

Perhaps we should interpret impurists' references to norms of rationality as references to norms of *prudence*—that is, norms that prioritize pursuit of self-interest or well-being.

On the plausible assumption that prudential norms sometimes differ from moral norms, this proposal faces the same problem as the previous one. Just as the distinctive normative problem with Naomi's decision seems not to be grounded in practical incoherence, it also seems not to be grounded in imprudence. Just as there is no reason to suppose that the post office's hours matter for the pursuit of Naomi's final ends, there is no reason to suppose that they matter for the pursuit of her well-being.

Again, the impurist might retreat to a highly controversial position to solve this problem: she might claim that the norms of morality fall out of the norms of prudence. The impurist who takes this approach gives up a substantial hostage to metaethical fortune. But the impurist who interprets 'practical rationality' as prudence also faces a distinctive second problem—one that will be instructive for our discussion moving forward.

The second problem is this: an impurism that connects knowledge solely to prudence cannot thereby establish the right *sort* of connections between knowledge and practical coherence. Interestingly, this problem arises even if we grant that prudence and practical coherence are tightly related. Suppose, for instance, that maximal practical coherence and maximal prudence necessarily coincide. In other words, any



departure from maximal prudence will be accompanied by some incoherence in the mental states involved in action. You might think that, on this sort of view, connections between knowledge and prudence would automatically establish just the sort of connection between knowledge and coherence that the impurist needs.

But things are more complicated. The case of Tracy's cigarettes will help to illustrate the problem. Although characters like Tracy are imprudent, and they also do lack knowledge, their imprudence is not related to their action in the right way to explain *why* they lack knowledge. To see why, we'll need to take a closer look at the structure of defensible impurist views.

As we saw in sections 1 and 2, impurists generally defend some precisification of the following conditional:

If  $S$  knows that  $p$ , then it is acceptable for  $S$  to rely on  $p$ .

As it stands, the principle faces serious problems. For one, paradigmatic *reasons of the wrong kind* seem to present counterexamples. Say, for instance, that an evil demon threatens to destroy the world if I rely on the proposition *that I exist* (fill in the details about "relying on" however you like). Surely, in such a case, it would be unacceptable for me to rely on the premise that I exist. It would be *morally* unacceptable for me to do so, it would be *prudentially* unacceptable, and for anyone who has the goal of existing, it would be forbidden by norms of practical coherence. As it stands, then, the principle above suggests that I cannot know that I exist in the face of the demon's

threat. But defenders of impurism generally do not want to collapse the distinction between right and wrong kinds of reasons for belief in this way.

Impurists are aware of this issue. To address it, they distinguish between sources of normative problems with a person's action. Fantl and McGrath, for instance, offer the following precisification of the principle above:

When you know a proposition  $p$ , no weaknesses in your epistemic position with respect to  $p$ ... stand in the way of  $p$  justifying you. (2009, 64)

This more sophisticated formulation helps to address the case of the demon's threat. Even if the demon's threat makes it the case that, *ultima facie*, the proposition that I exist does not justify me (in believing some further proposition, acting in a certain way, etc.), it is not my *epistemic position* that stands in the way of my being justified. Now, Fantl and McGrath's approach is only one of many available strategies for the impurist. The key point is that any viable impurism must rule out knowledge only in response to certain normative problems—problems that must, at least, be associated with the agent's epistemic position.

Now, return to Tracy's case. There is a sense in which Tracy is unable to pass the drugstore while meeting the norms of prudence: her plan to get more Wolverines is an imprudent one, and as long as she retains that plan, she cannot fully meet norms of prudence. But, as we've just seen, impurists must take care *only* to rule out knowledge that  $p$  when a problem with action is associated with the agent's epistemic position with respect to  $p$ . And the imprudence of Tracy's plan to get more Wolverines

has nothing to do with her epistemic position regarding the drugstore. To make this clear, we can stipulate that Tracy formed the plan to get more Wolverines before she had any idea about where she might get them. So, even if we grant that Tracy's imprudent plan necessarily carries incoherence with it, prudence is the wrong norm to explain the particular incoherence that comes from the way that she treats her evidence about the drugstore. To explain why her action falls short *in the right way to make a difference for knowledge*, then, the impurist must appeal directly to practical coherence.

So mere appeals to prudence do not allow impurists to satisfy our desiderata. Let's turn to a third way of understanding appeals to 'practical rationality.'

### *3.3 Rationality as Reasons-Responsiveness*

On a third approach to practical rationality, to be practically rational is to respond to one's reasons when planning and acting.<sup>79</sup> This approach also fails to give the impurist a way to meet our desiderata.

Just what does it take for a person to respond to her reasons? Say that my prudential reasons speak in favor of staying home, but my moral reasons speak in favor of going out. In this case, there is a sense in which I will inevitably fail to respond to some of my reasons. But it seems wrong to say that, in a case like this, I will inevitably fall short of practical rationality. To avoid this result, a defender of the reasons-

---

<sup>79</sup> For two recent defenses, see Kiesewetter (2017) and Lord (2018).

responsiveness approach can say that practical rationality is a matter of responding to the *balance*, or *correct weighting*, of one's reasons.<sup>80</sup>

Suppose that the impurist connects knowledge to this sort of reasons-responsiveness; she claims that knowledge that *p* requires the ability to rely on *p* in a way that respects the balance of one's reasons. At first, this approach may seem promising. After all, the primary problem I've cited for other understandings of 'practical rationality' has been the inability to account for connections between knowledge and morality. But it's very plausible that, in many (perhaps all) cases, morality provides us with decisive reasons for action. So practical rationality, understood as reasons-responsiveness, may be violated in all the cases in which morality makes a difference for knowledge.

Unfortunately, this approach does not make room for our other desideratum: connections between knowledge and practical coherence. The problem is clearest in cases like Tracy's. When Tracy passes the drugstore on too little evidence, she is open to a particular sort of criticism—a criticism that, I've argued, the impurist should consider relevant for knowledge. But this criticism cannot be straightforwardly put in terms of her failure to respond to the balance of her reasons. The relevant problem with Tracy's action, loosely speaking, is that she is not cautious enough. But her reasons for action do not (on balance) speak in favor of cautiously getting more

---

<sup>80</sup> Or, better yet, responding to the balance of one's *possessed* or *available* reasons (see Lord 2018 and Kiesewetter 2017 respectively). I'll set this complication aside.

cigarettes; they speak in favor of kicking the habit altogether. So the reasons-responsiveness account of practical rationality seems unable to provide the right sort of normative criticism for agents like Tracy.

But perhaps this is too quick. Although, on the face of it, norms of practical coherence are remarkably different from the norm of reasons-responsiveness, perhaps there are subtle connections between the two. If so, the impurist may seem to have the key to defusing my challenge: she will have an account on which a single sort of normativity—reasons-responsiveness—absorbs norms of practical coherence and norms of morality alike.

On closer inspection, however, things are not so simple. Even if norms of coherence are connected to reasons, they are not connected to reasons for action in the way that the impurist needs. It's worth taking a closer look to see why.

For clarity's sake, I'll limit my discussion to one norm of coherence: the norm that one intend (what one believes to be) the necessary means to one's intended ends. What is the connection between this norm and reasons? Perhaps the connection is straightforward: whenever one intends an end, one thereby obtains a reason to intend the necessary means. Many are suspicious of this approach, on the grounds that it licenses objectionable bootstrapping. Say, for instance, that I have no reason to go to City Hall. Then I take a pill that brings me to (coherently) intend to lie face-down in the middle of City Hall. Going to City Hall is a necessary means to lying face-down

there. So, on the current proposal, taking the pill generates a reason to go to City Hall. Many find this result unacceptable.<sup>81</sup>

On another proposal, one does not always have a reason to take the necessary means simply because one intends the ends. Rather, one always has a reason to be such that, if one intends the ends, one intends the necessary means.<sup>82</sup> In other words, one always has a reason to *avoid means-ends incoherence*. This proposal avoids the bootstrapping problem, but it faces problems of its own.<sup>83</sup> One of the most notable, powerfully posed by Niko Kolodny (2005, 542-7; cf. 2008, 374-82), begins from the observation that mere coherence between mental states does not seem terribly important in its own right. But, on the current proposal, there is *always* a reason that counts in favor of avoiding incoherence.<sup>84</sup> What could that reason be? This question is challenging enough that many reject the proposal that there is a standing reason that counts in favor of coherence.

Does this mean that there is no relationship between practical coherence and reasons? Not necessarily. On a currently popular proposal, even though our reasons

---

<sup>81</sup> See Bratman (1987) and Broome (1999). For a response, see Schroeder (2005).

<sup>82</sup> Note a related proposal: one ought to be such that (when one intends the end, one also intends the necessary means. This is sometimes called a “wide-scope” approach, since it suggests that ‘ought’ is a propositional operator that takes scope over a conditional. Defenders include Broome (1999), Hill (1973), Gensler (1985), and Wallace (2001).

<sup>83</sup> Among these problems are problems with asymmetry, which I set aside; for more, see Schroeder (2004, 339-40) and Kolodny (2005, 528-42).

<sup>84</sup> In fact, the problem is worse; on any plausible wide-scope approach, there is an exceptionless *requirement* to avoid incoherence. In other words, rules of practical coherence are *strict* (Broome 1999) or *stringent* (Schroeder 2009, 233). But it’s difficult to say what could count against incoherence so powerfully that it generates an exceptionless requirement.

do not directly favor coherence, the attitudes best supported by our reasons always (or almost always) cohere with one another. As a result, an agent who responds to his reasons in every way will inevitably be fully coherent. Contraposing: any agent who is not fully coherent must not be responding to his reasons in every way.<sup>85</sup>

To see this proposal at work, return to Tracy's case. Suppose Tracy's reasons favor her giving up the intention to get more Wolverines, and they also favor her passing the drugstore by. As we've already seen, Tracy does *not* adopt both of these mental states. Instead, against the balance of her reasons, she maintains the intention to get more Wolverines. This failure to respond to reasons is accompanied by an incoherence: given her evidence, Tracy's plan to get more Wolverines does not sit well with her risky decision to pass the drugstore by. On the current proposal, this is an entirely general trend. Any time an agent is incoherent, she will have failed, in one way or another, to respond to her reasons.

Unfortunately, even this approach to rationality cannot help the impurist. Why? Well, recall the impurist proposal that we're currently considering: when an agent knows that  $p$ , she is in a position to rely on  $p$  in a way that responds to reasons. It's clear how this formula rules out knowledge in cases like Naomi's; Naomi has decisive reasons, of a moral sort, to be more cautious. It's much less clear, on the other hand, that the formula establishes the right sort of connection between knowledge and practical coherence. Is Tracy in a position to rely, in a way that responds to reasons, on the

---

<sup>85</sup> This proposal can be found in Kiesewetter (2017), Kolodny (2005, 2008), and Lord (2018).

proposition that the drugstore will be open tomorrow? Well, in one sense, she is not. On the view we're currently considering, whenever Tracy acts in a way that betrays practical incoherence, there will be a failure to respond to reasons somewhere in her psychology.

As we saw in section 3.2, however, this is not the right sort of failure to rule out the possibility that Tracy has knowledge. The mere fact that *something* about an agent puts her in violation of a norm on action cannot be enough to rule out knowledge. Rather, a viable impurism must rule out knowledge of *p* *only* when the normative problem with an agent's action arises from the weakness of her epistemic position with respect to *p*. But it is the fact that Tracy intends to get more Wolverines *at all* that stands in the way of her responding to reasons. And the problem with this intention, presumably, does not arise in virtue of the weakness of her epistemic position about the drugstore. So Tracy does not fail to respond to reasons *in the right way* to preclude knowledge.

Taking stock: we've now seen that impurists cannot account for the connection between knowledge and practical coherence simply by alluding to norms of reasons-responsiveness. This approach to practical rationality, like the other two we've considered, did not provide an attractive way for the impurist to meet our two desiderata.



Nevertheless, the discussion in this section might seem to point the way toward a more promising strategy for the impurist. That discussion suggested that moral norms and norms of practical coherence both have the following trait:

**Trait T** any violation of the norm by any agent is accompanied, somewhere in the agent's psychology, by a failure to respond to reasons

Here is a candidate attempt to meet my challenge on behalf of the impurist: when one knows that *p*, one is in a position to rely on *p* in a way that meets *each of several* normative standards—specifically, all and only those standards that have trait T.

Note two things about this proposal. First, it departs dramatically from contemporary impurists' actual proposals. Contemporary impurists write as if there is a single master norm on action—the norm of *practical rationality*—and, when one knows that *p*, one's epistemic position with respect to *p* does not stand in the way of one's meeting that norm. The proposal on the table makes progress toward meeting my desiderata precisely by rejecting that assumption.

Second, even if this proposal is extensionally adequate, it seems inadequately illuminating. Why think that norm-violations are particularly worthy of theoretical attention—or particularly apt to make a difference within epistemology—when the norm in question has trait T? On the face of it, the fact that I have violated a norm that has trait T is not interesting in itself; at most, it provides evidence that something with genuine theoretical importance has happened somewhere in my mind. In the next

section, I'll sharpen this charge, and I'll say a bit about some more promising ways forward for the impurist.

#### **Section 4: Ways Forward for Impurism**

In section 1, I posed a challenge to impurists: they must specify just which norms on action are distinctively connected to knowledge. To the extent that contemporary impurists address this challenge, they do so by naming norms of practical rationality alone. We've now seen why this strategy is inadequate: on any theory of practical rationality, it fails to appropriately cover both moral norms and norms of practical coherence.

Of course, the fact that existing impurist views fail to meet my challenge doesn't mean that no impurist view possibly could. And indeed, I hope to show that there is a way forward for the impurist. I begin this final section by making a proposal about the best format for a view that successfully meets my challenge. But adopting this format, I'll note, is only half the battle; the impurist must also explain *why* some norms are connected to knowledge while others are not. I close the chapter by surveying three strategies the impurist might adopt to meet this explanatory challenge.

##### *4.1 The Best Format for an Impurist Response*

In order to meet the challenge I've set, impurists should abandon biconditional knowledge-action connections of the following form:

Where one's choice is  $p$ -dependent, it is appropriate to treat the proposition that  $p$  as a reason for acting iff you know that  $p$ . (Hawthorne and Stanley 2008, 578)

The problem for this principle is simple. It suggests that there is a single sense of “appropriate” on which my ability to appropriately treat  $p$  as a reason for acting covaries perfectly with my knowledge that  $p$ . But, as we’ve seen, the impurist should hold that some people are precluded from knowing by *moral* problems with relying on a proposition, and others are precluded from knowing by problems with the *coherence* of relying on a proposition. Moreover, it’s very plausible that these problems are not coextensive: sometimes, one’s epistemic position makes it problematic by moral standards, but not problematic by standards of practical coherence, to rely on a proposition. Whatever reading we give to “appropriate,” then, it will not pick out all the normative violations that it should.<sup>86</sup>

To avoid this problem, the impurist should instead embrace a view on which knowledge entails, but is not entailed by, avoidance of certain normative criticisms of action. She should, in other words, embrace several principles of the following form:

If  $S$  knows that  $p$ , then  $S$ ’s epistemic position with respect to  $p$  does not stand in the way of her relying on  $p$  in a way that meets the standards of morality.

If  $S$  knows that  $p$ , then  $S$ ’s epistemic position with respect to  $p$  does not stand in the way of her relying on  $p$  in a way that meets the standards of practical coherence.

---

<sup>86</sup> To avoid this result, the impurist could replace “appropriate” with “appropriate according to each of a variety of distinct standards.” But this loses the sense of the original proposal and turns it into shorthand for the alternative that I go on to discuss.

This list of necessary conditions might continue; perhaps, for instance, norms of prudence also place a separate constraint on knowledge.

A caveat: *if* following the norms of practical coherence entails following all moral norms (a possibility I noted in section 3.1), then there is no possibility of tension between these two norms. On this picture of the relationship between morality and coherence, Hawthorne and Stanley's biconditional principle can be vindicated.

Nevertheless, since it can only be made good against the background of a particular metaethical view, impurists would do best to leave Hawthorne and Stanley's principle behind. Although the idea that any immoral agent must be incoherent certainly has defenders, it is highly controversial. Indeed, the burden of proof surely lies with its defenders; on the face of it, it's quite easy to imagine an entirely coherent Caligula.<sup>87</sup> If she adopts a biconditional knowledge-action link, the impurist rests the prospects for her view on the impossibility of such a figure. She concedes that her epistemological program can only be made good if a second ambitious program, in metaethics, is also successful. If possible, the impurist should not make this concession.

Suppose, then, that the impurist embraces multiple distinct necessary conditions on knowledge, each of which is associated with a different flavor of normativity. This is the most appealing format for a response to my challenge, but it is only half the battle;

---

<sup>87</sup> This case is from Gibbard (1999, 145).

the impurist should also explain *why* these necessary conditions exist while others do not. Why does knowledge rule certain *moral* criticisms of action (namely, ones that have to do with the weakness of one's epistemic position) but not the possibility of analogous criticisms *in terms of the norms of the Clean White Shirt Group*?

#### *4.2 Three Strategies*

In this final subsection, I'll sketch three strategies that the impurist could use to explain why some norms on action, but not others, are distinctively connected to knowledge.

**Brute List:** On this first approach, the impurist offers a list of norms that are distinctively connected to knowledge, and calls it a brute fact—one that stands in no need of further explanation—that each norm has that distinctive connection.

It may be that this is the best that the impurist can do. And, of course, all explanations have to stop at some point. But it would be doubly unfortunate if the impurist could not say anything informative to distinguish the norms that matter for knowledge from the norms that don't. For one, this leaves a troubling lack of unification in her picture of epistemology. But, even more troublingly, the brute list approach may undercut the motivation for adopting impurism in the first place.

To see this, recall one of the major sources of evidence for impurism: many impurists are impressed by the fact that we can use knowledge-attributions to make claims about the normative status of actions. Naomi, for instance, can deflect certain moral

criticisms of her choice to pass the post office by saying, “I know that it’ll be open tomorrow,” and Tracy can portray her choice to pass the drugstore as coherent by making the same claim.

It’s striking that this point about knowledge-discourse seems to generalize beyond norms of practical coherence and moral norms. To see this, return to the example of the Clean White Shirt Group. Suppose that Kayla is confronted by an adherent of that group, who accuses her of failing to check her dresser on insufficient evidence. Kayla, like Naomi and Tracy, can defend her action with a knowledge-claim: she can say “I know that I have a clean white shirt in my dresser.” We appear to have just the same sort of evidence for a link between knowledge and merely stipulated rules that we have for a link between knowledge and practical coherence.

The problem here is a serious one, and it has not gone unnoticed by opponents of impurism.<sup>88</sup> If our knowledge-discourse is flexible in a way that can successfully convey information about any salient standard at all, evidence from everyday usage may fit best not with impurism but with purist contextualism. So, unless she can provide a compelling, illuminating account of the difference between Kayla’s situation and César’s, the impurist stands open to the charge that some of the apparent evidence in favor of her view is actually no evidence at all.<sup>89</sup>

---

<sup>88</sup> DeRose (2009, ch. 7) presses a related worry.

<sup>89</sup> Tristram McPherson, in conversation, has suggested that there is a principled way of taking up this challenge by discarding my second desideratum. On this proposal, knowledge-discourse that aims to evaluate *coherence* differs from other knowledge-discourse that aims at evaluation because, in the

So providing a principled, attractive story about the difference between norms on action is not merely a way for the impurist to flesh out the details her view. It is also a vital step in explaining why her view is genuinely well-motivated. As a result, impurists would do best to avoid the brute-list approach.

**Metaethical Unity:** On this second approach, the impurist claims that norms of morality and norms of practical coherence both make a difference for knowledge precisely because both are *genuinely normative* in just the same sense. Merely stipulated norms, by contrast, do not make a difference for knowledge precisely because they are not genuinely normative. Note that this proposal can be made with or without the claim that being practically coherent requires being moral.

The appeal of this approach is clear—intuitively, norms of practical coherence and norms of morality both *matter* in a way that merely stipulated norms do not. But there is also a problem with this approach: namely, there are good reasons to worry that moral norms and norms of coherence are not, in fact, genuinely normative in just the same sense.

---

former case alone, the subject of the knowledge-attributions is *committed* to caring about whether she meets epistemic standards. But note that the following, parallel proposal seems like an equally good impurist response: knowledge-discourse that aims to evaluate *responsiveness to genuine reasons* is different from all other sorts of knowledge-discourse that aims at evaluation because, in the former case alone, it's *genuinely important* for the subject to meet certain epistemic standards. As I showed above, impurists' references to 'practical rationality' can be understood in a way that privileges either of these proposals. So there is no more reason to think that the impurist should give up on my second desideratum than there is to think that she should give up on my first.

Section 3.3 showcased some reasons for doubting that the connection between practical coherence and *genuine normativity* is straightforward. Many doubt that, merely by adopting a mental state, one obtains a genuinely normative reason to make one's other mental states cohere with it. Many also doubt that there is any standing genuinely normative reason to keep one's mental states coherent. And weaker connections between genuinely normative reasons and coherence—recall trait T—do not seem to capture the distinctive importance of practical coherence norms. In the end, it seems likely that morality and practical coherence both make special and important claims on us, but that their claims on us are important in different ways. Perhaps, for instance, morality is special because it provides us with *robustly normative* reasons,<sup>90</sup> or *authoritative* reasons,<sup>91</sup> whereas practical coherence is special because it is *constitutive of agency*,<sup>92</sup> or because we have *committed* ourselves to it.<sup>93</sup>

Now, maybe these reasons for worry about the metaethical unity picture can be assuaged. Just as there is a noteworthy contingent of philosophers who believe that being coherent simply entails being moral, there is also a noteworthy contingent of philosophers who believe that morality and practical coherence exert just the same sort of pull on us. But, just as it would be unfortunate for the impurist to rest the prospects for her view on the controversial position that coherence requires morality,

---

<sup>90</sup> McPherson (2011, 233)

<sup>91</sup> Hampton (1998, 85ff)

<sup>92</sup> See, for instance, Davidson (2004) and Korsgaard (1996).

<sup>93</sup> See, for instance, Southwood (2008).



it would also be unfortunate for the impurist to rest the prospects for her view on the controversial position that morality and practical coherence are both genuinely normative in just the same sense.

**Social Role:** On this third strategy, the impurist appeals not to the idea that practical coherence and morality have similar standing *in the normative universe itself*, but instead to the idea that practical coherence and morality play similar roles in some real or idealized social practice.

This strategy seems more promising than the other two. It's much less ambitious to suggest that morality and practical coherence share a social role than it is to suggest that, for instance, being practically coherent requires being fully moral. And, by taking this approach, the impurist stands to gain a genuinely informative story about the connection between knowledge and action.

For an example of the sort of story that might emerge, consider Edward Craig's account of social role played by the concept of knowledge. Very roughly, Craig (1990) suggests that, by deploying the concept of knowledge, we fill an important social role: we flag (sufficiently) reliable sources of information. This proposal seems to provide an illuminating explanation of the close connections between knowledge and action. Perhaps knowledge is sensitive to moral norms precisely because our shared epistemic discourse stably aims to communicate about whether an informant is reliable enough that acting on the basis of her testimony will allow one to meet moral standards. Perhaps, by the same token, knowledge is sensitive to norms of practical

coherence precisely because our shared epistemic practice stably aims to communicate about whether an informant is reliable enough to help one pursue one's immediate plans (regardless of whether those goals have genuine merit).

Of course, the devil is in the details. Craig's particular story about the role of knowledge-claims is, most likely, not a perfect fit for the impurist who wants to meet my desiderata.<sup>94</sup> Most impurists are concerned to evaluate whether the knower's beliefs have enough support for *her* to rely on them; Craig's proposal, by contrast, most naturally calls attention to the question of whether a knower's beliefs are supported enough for *anyone in her community* to rely on them. But I've referred to Craig here solely by way of illustrative example; surely, there are other productive accounts to be given of our knowledge-discourse. In my view, appeals to the connections between knowledge and norms on action in society provide the most promising way forward for impurists.

A good deal of this chapter has been devoted to providing the defenders of existing impurist proposals with some reasons for pessimism. But I've concluded the chapter by providing reasons for optimism about impurism more broadly. It's true that existing impurist views are generally unacceptable, because they fail to approach norms on action with appropriate nuance. But, nevertheless, there is a way forward

---

<sup>94</sup> Stephen Grimm (2015) draws on Craig to develop an impurism that comes close to meeting my desiderata. But, on his view, a morally apathetic agent can sometimes know that *p* even though the weakness of her epistemic position makes it morally unacceptable for her to rely on *p*. Chapter 2 of this dissertation shows that this is a bad result.

for impurists. With some additional work, it seems likely that impurists can explain, in a genuinely illuminating way, why only some of the many norms on action make a difference to knowledge.

## Bibliography

- Baker, Derek. 2018. "Skepticism about Ought *Simpliciter*." In *Oxford Studies in Metaethics*, vol. 13, edited by Russ Shafer-Landau, 230-252. Oxford: Oxford University Press.
- Basu, Rima. Forthcoming. "The Wrongs of Racist Beliefs." *Philosophical Studies*.
- Basu, Rima and Mark Schroeder. 2019. "Doxastic Wrongs." In *Pragmatic Encroachment in Epistemology*, edited by Brian Kim and Matthew McGrath, 181-205. New York: Routledge.
- Bolinger, Renée. Forthcoming. "The Rational Impermissibility of Accepting (Some) Racial Generalizations." *Synthese*.
- Broome, John. 1999. "Normative Requirements." *Ratio* 12 (4): 398-419.
- . 2007. "Does Rationality Consist in Responding Correctly to Reasons?" *Journal of Moral Philosophy* 4 (3): 349-74.
- Brown, Jessica. 2008. "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning." *Noûs* 42 (2): 167-189.
- . 2013. "Experimental Philosophy, Contextualism, and SSI." *Philosophy and Phenomenological Research* 86 (2): 233-261.
- Buchak, Lara. 2014. "Belief, Credence, and Norms." *Philosophical Studies* 169 (2): 1-27.
- Craig, Edward. 1990. *Knowledge and the State of Nature*. Oxford: Oxford University Press.

- D'Arms, Justin and Daniel Jacobson. 2000. "The Moralistic Fallacy: On the 'Appropriateness' of Emotions." *Philosophy and Phenomenological Research* 61 (1): 65-90.
- DeRose, Keith. 1992. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research* 52 (4): 913-929.
- Fantl, Jeremy & McGrath, Matthew. 2002. "Evidence, Pragmatics, and Justification." *Philosophical Review* 111 (1): 67-94.
- . 2007. "On Pragmatic Encroachment in Epistemology." *Philosophy and Phenomenological Research* 75 (3): 558-589.
- . 2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Feldman, Richard. 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60 (3): 667-695.
- Foot, Phillipa. 1997. "Morality as a System of Hypothetical Imperatives." Reprinted in *Moral Discourse and Practice*, edited by Stephen Darwall, Allan Gibbard, and Peter Railton: 313-322. Oxford: Oxford University Press.
- Fritz, James. 2017. "Pragmatic Encroachment and Moral Encroachment." *Pacific Philosophical Quarterly* 98 (51): 643-661.
- Gardiner, Georgi. 2018. "Evidentialism and Moral Encroachment." In *Believing in Accordance with the Evidence: New Essays on Evidentialism*, edited by K. McCain, 169-195. Switzerland: Springer.
- Gendler, Tamar. 2011. "On the Epistemic Costs of Implicit Bias." *Philosophical Studies* 156 (1): 33-63.
- Gensler, H. 1985. "Ethical Consistency Principles." *Philosophical Quarterly* 35: 156-70.
- Gibbard, Allan. 1999. "Morality as Consistency in Living: Korsgaard's Kantian Lectures." *Ethics* 110 (1): 140-164.
- Graham, Peter A. 2010. "In Defense of Objectivism about Moral Obligation." *Ethics* 121 (1): 88-115.

- Grimm, Stephen. 2015. "Knowledge, Practical Interests, and Rising Tides." In *Epistemic Evaluation: Point and Purpose in Epistemology*, edited by John Greco and David Henderson. New York: Oxford University Press.
- Hampton, Jean. 1998. *The Authority of Reason*. New York: Cambridge University Press.
- Harman, Elizabeth. 2015. "The Irrelevance of Moral Uncertainty." In *Oxford Studies in Metaethics*, vol. 10, edited by Russ Shafer-Landau: 53-79. Oxford: Oxford University Press.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, John and Jason Stanley. 2008. "Knowledge and Action." *The Journal of Philosophy* 105 (10): 571-590.
- Hieronymi, Pamela. 2005. "The Wrong Kind of Reason." *The Journal of Philosophy* 102(9): 437-457.
- Hill, T. Jr. 1973. "The Hypothetical Imperative." *Philosophical Review* 82, 429-50.
- Ichikawa, Jonathan Jenkins and Steup, Matthias. 2016. "The Analysis of Knowledge." *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.).  
<https://plato.stanford.edu/archives/win2016/entries/knowledge-analysis/>
- Jackson, Frank. 1991. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101 (3): 461-482.
- Jeffrey, Richard C. 1971. "Dracula Meets Wolfman: Acceptance vs. Partial Belief." In *Induction, Acceptance, and Rational Belief*, edited by Marshall Swain. Dordrecht: Reidel.
- Kelly, Thomas. 2002. "The Rationality of Belief and Some Other Propositional Attitudes." *Philosophical Studies* 110 (2): 163-196.
- Kiesewetter, Benjamin. 2017. *The Normativity of Rationality*. Oxford: Oxford University Press.
- Kim, Brian. 2017. "Pragmatic Encroachment in Epistemology." *Philosophy Compass* 12 (5): 1-14.

- Kolodny, Niko. 2005. "Why Be Rational?" *Mind* 114 (455): 509-563.
- . 2007. "How Does Coherence Matter?" *Proceedings of the Aristotelian Society* 107 (1pt3): 229-263.
- Korsgaard, Christine. 1996. *The Sources of Normativity*, edited by Onora O'Neill. Cambridge: Cambridge University Press.
- Lasonen-Aarnio, Maria. 2014. "Higher-Order Evidence and the Limits of Defeat." *Philosophy and Phenomenological Research* 88 (2): 314-345.
- Locke, Dustin. 2014. "The Decision-Theoretic Lockean Thesis." *Inquiry* 57 (1): 28-54.
- . 2015. "Practical Certainty." *Philosophy and Phenomenological Research* 90 (1): 72-95.
- . 2017. "Implicature and Non-Local Pragmatic Encroachment." *Synthese* 194 (2): 631-654.
- Lord, Errol. 2015. "Acting for the Right Reasons, Abilities, and Obligation." In *Oxford Studies in Metaethics*, vol. 10, edited by Russ Shafer-Landau: 26-51. Oxford: Oxford University Press.
- . 2018. *The Importance of Being Rational*. Oxford: Oxford University Press.
- McPherson, Tristram. 2011. "Against Quietist Normative Realism." *Philosophical Studies* 154 (2): 223-240.
- . 2018. "Authoritatively Normative Concepts." In *Oxford Studies in Metaethics*, vol. 13, edited by Russ Shafer-Landau, 253-277. Oxford: Oxford University Press.
- Moore, Eric. 2007. "Objective Consequentialism, Right Actions, and Good People." *Philosophical Studies* 133 (1): 83 - 94.
- Moss, Sarah. 2018a. "Moral Encroachment." *Proceedings of the Aristotelian Society* 118 (2): 177-205.
- . 2018b. *Probabilistic Knowledge*. Oxford: Oxford University Press.

- Nye, Howard. 2017. "The Wrong Kind of Reasons." In *The Routledge Handbook of Metaethics*, edited by Tristram McPherson and David Plunkett, 340-354. New York: Routledge.
- Owens, David. 2000. *Reason without Freedom: The Problem of Epistemic Normativity*. New York: Routledge.
- Pace, Michael. 2011. "The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will to Believe'." *Noûs* 45 (2): 239-268.
- Parfit, Derek. 2001. "Rationality and Reasons." In *Exploring Practical Philosophy: From Action to Values*, edited by D. Egonsson, J. Josefsson, B. Petersson, & T. Rønnow-Rasmussen, 17-39. Aldershot: Ashgate
- . 2011. *On What Matters, Volume 1*. Oxford: Oxford University Press.
- Persson, Ingmar. 2007. "Primary and Secondary Reasons." In *Homage á Wlodek*, edited by T. Rønnow-Rasmussen, B. Petersson, J. Josefsson and D. Egonsson, URL = <http://www.fil.lu.se/hommageawlodek>.
- Rabinowicz, Wlodek, and Toni Rønnow-Rasmussen. 2006. "Buck-Passing and the Right Kind of Reasons." *The Philosophical Quarterly* 56(222): 114-120.
- Raz, Joseph. 2009. "Reasons: Practical and Adaptive." In D. Sobel and S. Wall (eds.), *Reasons for Actions*, Cambridge: Cambridge University Press, 37-57.
- Rinard, Susanna. 2017. "No Exception for Belief." *Philosophy and Phenomenological Research* 94 (1): 121-143.
- Ross, Jacob. 2006. "Rejecting Ethical Deflationism." *Ethics* 116 (4): 742-768.
- Ross, Jacob and Mark Schroeder. 2014. "Belief, Credence, and Pragmatic Encroachment." *Philosophy and Phenomenological Research*. 88 (2): 259-288.
- Rowland, Richard. 2015. "Dissolving the Wrong Kind of Reason Problem." *Philosophical Studies* 172(6): 1455-1474.
- Scanlon, T.M. 1998. *What We Owe to Each Other*. Cambridge: Belknap Press of Harvard University Press.



- Schroeder, Mark. 2004. "The Scope of Instrumental Reason." *Philosophical Perspectives* 18: 337-364.
- . 2005. "Instrumental Mythology." *Journal of Ethics and Social Philosophy* 1 (2): 1-12.
- . 2009. "Means-End Coherence, Stringency, and Subjective Reasons." *Philosophical Studies* 143 (2): 223-248.
- . 2010. "Value and the Right Kind of Reason." In *Oxford Studies in Metaethics*, volume 5, edited by Russ Shafer-Landau, 25-55. Oxford: Oxford University Press.
- . 2012a. "Stakes, Withholding, and Pragmatic Encroachment on Knowledge." *Philosophical Studies* 160(2): 265-285.
- . 2012b. "The Ubiquity of State-Given Reasons." *Ethics* 122(3): 457-488.
- . 2013. "State-Given Reasons: Prevalent, If Not Ubiquitous." *Ethics* 124: 128-140.
- . 2018. "When Beliefs Wrong." *Philosophical Topics* 46(1): 115-127.
- Sepielli, Andrew. 2009. "What to Do When You Don't Know What to Do." In *Oxford Studies in Metaethics*, vol. 4, edited by Russ Shafer-Landau: 5-28. Oxford: Oxford University Press.
- Sharadin, Nathaniel. 2016. "Reasons Wrong and Right." *Pacific Philosophical Quarterly*. 97 (3): 371-399.
- Skorupski, John. 2007. "Buck-Passing about Goodness". In *Hommage á Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*, edited by Toni Rønnow-Rasmussen et al. URL: <http://www.fil.lu.se/hommageawlodek>
- Smithies, Declan. 2012. "The Normative Role of Knowledge." *Noûs* 46 (2): 265-288.
- Southwood, Nicholas. 2008. "Vindicating the Normativity of Rationality." *Ethics* 119 (1): 9-30.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Thomson, J.J. 2001. *Goodness and Advice*. Princeton: Princeton University Press.

- Tiffany, Evan. 2007. "Deflationary Normative Pluralism." *The Canadian Journal of Philosophy* 37 (supplement), 231-262.
- Wallace, R. J. 2001. "Normativity, Commitment and Instrumental Reason." *Philosophers' Imprint* 1: 1-26.
- Way, Jonathan. 2012. "Transmission and the Wrong Kind of Reason." *Ethics* 122(3): 489-515.
- Way, Jonathan and Daniel Whiting. 2016. "If You Justifiably Believe that You Ought to  $\phi$ , You Ought to  $\phi$ ." *Philosophical Studies* 173, (7): 1873-1895.
- Weatherson, Brian. 2005. "Can We Do Without Pragmatic Encroachment?" *Philosophical Perspectives* 19 (1): 417-443.
- . 2014. "Running Risks Morally." *Philosophical Studies* 167, no. 1: 141-163.
- Wedgwood, Ralph. 2002. "The Aim of Belief." *Philosophical Perspectives* 16: 267-297.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. New York: Oxford University Press.
- Worsnip, Alex. 2015. "Two Kinds of Stakes." *Pacific Philosophical Quarterly* 96(3): 307-324.