

OBLIGATIONS AND BLAME AT THE MORAL EPISTEMOLOGICAL LIMIT

A Dissertation

Presented to the Faculty of the Graduate School
of Cornell University

in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

by

Adam True Bendorf

February 2016

© 2016 Adam True Bendorf
ALL RIGHTS RESERVED

OBLIGATIONS AND BLAME AT THE MORAL EPISTEMOLOGICAL LIMIT

Adam True Bendorf, Ph.D.

Cornell University 2016

The focus of this dissertation is a special type of moral ignorance: non-culpable ignorance of truths of the form *I am obligated to do/not do X*. Call ignorance of this type NCFPO ignorance (for **N**on-**C**ulpable **F**irst-**P**erson **O**bligation ignorance). Many philosophers believe that NCFPO ignorance exists. But whether NCFPO ignorance is possible hangs on the issue of whether a person can be obligated to perform an action that, for good reason, she does not believe that she is obligated to perform—an issue about which philosophers disagree. Is it the case, then, that NCFPO ignorance is possible? And if it is, just how extreme can a person's NCFPO ignorance be? Is it possible for a person to be non-culpably ignorant of *all* of her moral obligations? If so, under what conditions is this possible?

I argue here that NCFPO ignorance is not only possible, but it is possible for a person to be non-culpably ignorant of all of her moral obligations. This conclusion sheds light on several previously unexplored connections between epistemologists' views of external world skepticism and ethicists' views of moral skepticism. It also has implications concerning moral practice, since, as I argue in the dissertation, NCFPO ignorance constitutes an excuse for wrongdoing.

BIOGRAPHICAL SKETCH

Adam True Bendorf received a B.S. in Mathematics from the University of Michigan and an M.A. in Philosophy from Western Michigan University.

To my parents.

ACKNOWLEDGMENTS

I am indebted to many people for their help in bringing this project to fruition.

First, I am extremely grateful to all of the members of Cornell's Sage School of Philosophy—faculty, staff, and graduate students. Given the myriad ways in which the Sage School has influenced me, a proper acknowledgement would be at least as long as this dissertation and would take even more energy to write. I hope that I will be forgiven for simply singling out some members of the Sage School for their help. For penetrating criticism, I owe thanks to Richard Boyd, Brandon Conley, Hun Chung, Austin Duggan, Augie Faller, Gail Fine, Carl Ginet, Harold Hodes, Theo Korzukhin, Flora Lee, Kate Manne, Katie Mathie, Andrew McGonigal, Colin McLear, Dan Murphy, Scott O'Connor, Vlad Pintro, Sydney Penner, Will Starr, Erin Taylor, and Ru Ye. Brent Kyle also deserves to be acknowledged for his valuable contributions as a philosopher, roommate, and philanthropist. Stephen Humphreys-Mahaffey, Clifford Roberts, and Jona Vance are to be commended for providing me with crucial support in times of emotional and philosophical crisis; I have deeply treasured our friendship, as well as our philosophical interactions, the fruits of which are apparent on almost every page of this dissertation. Special thanks are also owed to the Sage School's office staff (past and present)—especially Paula Epps, Pam Hanna, and Dorothy Vanderbilt—who have served as models of professionalism, and who have saved me from myself on a number of occasions.

My greatest philosophical debts are to my advisor—Nick Sturgeon—and my other committee members: Derk Pereboom and Nico Silins. These three have served as ideal philosophical interlocutors. All three exhibit to the highest degree the cardinal virtues of the philosopher: clarity, precision, and insightfulness. They are also extremely patient, which has been a great help given the

numerous false starts and missteps that I have made in writing this dissertation.

Many people outside the Sage School also deserve thanks for their help over the years. Financial support from my parents, Bob Eagle, and Aunt Reet made possible my long, twisted educational journey. Support from the philosophy department at Western Michigan University was also indispensable; I owe thanks to Fritz Allhoff, Marc Alspector-Kelly, Arthur Falk, Joanne Lau, Tim McGrew, Ali Niedbalski, Chris Ragg, Jonah Schupbach, and Spencer Smith. Finally, I owe a special debt of gratitude to my girlfriend—Erin Taylor Wiley—and my parents. Erin has discussed with me nearly every argument that made it into this dissertation (as well as many that didn't); she also helped me through the dark times. As for my parents, this dissertation is dedicated to them. Their intelligence and pragmatism has shaped my philosophical development in ways I recognize but cannot properly articulate. I hope that this dissertation lives up to the examples that they have set.

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgments	v
Table of Contents	vii
1 Introduction	1
1.1 Preliminaries	3
1.2 Question 1: Is NCFPO ignorance possible?	4
1.3 Question 2: Assuming that NCFPO ignorance is possible, do the NCFPO ignorant deserve blame for acting wrongly?	6
1.4 Question 3: Assuming that NCFPO ignorance is possible, how extensive can a mature person’s NCFPO ignorance be?	9
1.5 The dissertation, summarized	12
1.5.1 NCFPO ignorance is possible	13
1.5.2 Moral skepticism can be a reasonable position for some people to adopt	15
1.5.3 NCFPO ignorance constitutes an excuse for wrongdoing .	18
1.6 Generalizing to other types of moral or ethical ignorance	20
2 Does a Person’s Epistemic State Constrain her Obligations?	25
2.1 Introduction	25
2.2 Clarifying the view that a person’s epistemic state constrains her obligations	27
2.2.1 Arguments that a person’s epistemic state constrains her obligations	30
2.2.2 Arguments that an agent’s epistemic state does not con- strain her obligations	39
2.2.3 Summary	43
2.3 A new argument against (Epistemic-Sensitivity)	43
3 Can Skepticism Be Reasonable?	50
3.1 Introduction	50
3.2 Autonomy claims	56
3.3 Counterexamples to autonomy	62
3.4 The (Fancy Inference Assumption)	66
3.5 Moral experiences	68
3.6 Moral intuitions	73
3.7 The Transition Objection	75
3.8 Epistemically required skepticism	78
3.9 Conclusion: psychological issues	83

4	Blame and Moral Ignorance	86
4.1	Introduction	86
4.2	Epistemological preliminaries	89
4.3	The state of the literature	90
4.4	Sharpening the Forgiving and Unforgiving Views	93
4.5	The abductive argument	96
	4.5.1 The first observation to be explained	96
	4.5.2 The second observation	100
	4.5.3 Formulating the abductive argument	105
4.6	The Quality of Will objection	106
	4.6.1 Talbert	108
	4.6.2 Harman	112
	4.6.3 An alternative	115
	4.6.4 Can defenders of the Quality of Will View explain the Comparison Judgments?	118
4.7	The second response on behalf of the Unforgiving View: epis- temic sensitivity, again	122
4.8	The Partially Forgiving View	125
4.9	Conclusion	128
	Bibliography	129

CHAPTER 1

INTRODUCTION

This dissertation is about a special type of moral ignorance: non-culpable ignorance of truths of the form *I am obligated to do/not do X*. I will refer to this type of ignorance as **NCFPO ignorance**, which stands for **Non-Culpable First-Person Obligation** ignorance, and I will say that a person who is non-culpably ignorant of a truth of the form *I am obligated to do/not do X* is **NCFPO ignorant**.

There are two ways in which an agent can be NCFPO ignorant. First, it might be that, for some action *X*, the agent is obligated to do *X*, but she justifiably believes that she is *not* obligated to do *X*; in other words, she might have a justified false belief about the moral status of *X*ing. Second, an agent might be obligated to do *X* while justifiably lacking any sort of belief about the moral status of *X*ing; in other words, she might reasonably be undecided about whether *X*ing is obligatory.

The term “NCFPO ignorance,” while inelegant, is useful because it draws attention to two features that differentiate this type of ignorance from other types of moral ignorance. First, not all moral ignorance is first-personal, as NCFPO ignorance is. A person could reasonably be said to be morally ignorant if, for example, she holds a false view of someone else’s moral situation. But her moral ignorance would not count as NCFPO ignorance, since it does not concern *her own* moral situation. Second, not all moral ignorance is ignorance of obligations. If a person is ignorant of whether it would be morally good or bad for her to act in a certain way, then she could be said to be morally ignorant. Again, however, this person’s moral ignorance would not count as NCFPO ignorance, since her ignorance would not concern what she is *obligated* to do.

In defining NCFPO ignorance as non-culpable ignorance of truths of the

form *I am obligated to do/not do X*, I am appealing to the commonplace that ignorance comes in two flavors: the culpable, and the non-culpable. The distinction between culpable and non-culpable ignorance is a familiar one. Someone who has no evidence that her car has been stolen during the night does not fall short epistemically if she fails to believe the truth that it actually has been. In such circumstances, the car owner's ignorance is non-culpable. But if the owner did have very strong evidence that her car had been stolen, and she nonetheless persisted in the belief that it had not been, she *would* be falling short epistemically. Her ignorance would be culpable. Arguably, the same distinction—between culpable ignorance on the one hand, and non-culpable ignorance on the other—applies to first-person ignorance of obligations.¹ It seems that some first-person ignorance of obligations is culpable, and some is not. The freshman egoist, who believes for bad reasons that he is never obligated to help others, *should* know better. His ignorance of what he is obligated to do is culpable. However, sometimes people's evidence about what they are obligated to do is simply inconclusive. If a person with inconclusive evidence about whether she is obligated to perform a certain action chooses to suspend judgment about whether she is so obligated, then although she would count as ignorant of her obligation, her ignorance would be different from that of the freshman egoist—her ignorance would be non-culpable.

The main goal of this introduction is to introduce three questions concerning the existence, scope, and importance of NCFPO ignorance, which will frame this dissertation. Much hangs on how these questions are to be answered. After explaining what does, I will preview the answers to these questions that I will defend as well as the arguments that I will use to do so. I will also, in the last

¹It will soon become clear why I say that the same distinction "arguably" applies to first-person ignorance of obligations. See the section below entitled "Question 1: Is NCFPO ignorance possible?"

section of this introduction, discuss the extent to which similar questions do or do not arise for other types of moral ignorance, such as ignorance of what would be good or bad for a person.

1.1 Preliminaries

There is an obvious reason why one might believe that people do sometimes suffer from NCFPO ignorance. If there are objective truths concerning what we are obligated to do, then our cognitive shortcomings would seem to ensure that we will be ignorant of many of them. Our knowledge is limited, and our faculties are imperfect. It should come as no surprise, then, if, through no fault of our own, we sometimes go wrong about—or simply fail to recognize—what we morally ought to do.

But even if there are no objective moral truths concerning what we are obligated to do, there is still room to recognize the existence of NCFPO ignorance. Moral relativists, who deny the existence of moral objectivity, can admit that, through no fault of our own, we sometimes fail to recognize what we are obligated to do. Suppose, for example, that cultural moral relativism is correct. A person might make a mistake about what she is obligated to do because she is mistaken about what is demanded by her culture's moral code; if she is mistaken through no fault of her own, then her ignorance would count as NCFPO ignorance. Quasi-realists can also make room for NCFPO ignorance. The quasi-realist's project is to find a way that an anti-realist can say almost everything that the realist does, without lapsing into inconsistency. Since the view that we can be mistaken about or simply ignorant of what we are obligated to do is an important commitment of realism, the quasi-realist will surely want to capture it.² It is a small step to the claim that NCFPO ignorance exists; all the quasi-

²As Blackburn does. See, for example, his discussion of the Taliban in Blackburn [3], 215.

realist needs to add is that such ignorance is sometimes non-culpable.

A wide range of philosophers, then, can admit the existence of NCFPO ignorance. Indeed, it might seem as if everyone *should* admit the existence of NCFPO ignorance—it might seem as if it is simply obvious that NCFPO ignorance exists. In fact, however, this is not at all obvious. There is reason to think that perhaps NCFPO ignorance isn't possible.

1.2 Question 1: Is NCFPO ignorance possible?

To see why one might think that NCFPO ignorance is not possible, consider the following case. Suppose that Dana is a doctor treating a patient, Pat, who has a life-threatening illness. Dana must choose one of two medicines to administer; I will refer to these as LOOKS-GOOD and REALLY-IS-GOOD. Unfortunately for Pat, Dana has misleading evidence about both. Dana has strong reason to believe that LOOKS-GOOD would cure Pat; in fact, however, administering LOOKS-GOOD would kill him. Dana also has strong reason to believe that REALLY-IS-GOOD would kill Pat; again, however, her evidence is misleading, and administering REALLY-IS-GOOD would cure him.

Ask yourself which medicine, if any, Dana is obligated to administer. Typically, doctors are obligated to administer the medicine that will cure their patients. But many people find it hard to believe that Dana is obligated to administer REALLY-IS-GOOD.³ Some will even go so far as to say that if Dana is obligated to administer either of the medicines (and perhaps she isn't), what she morally ought to do is to follow her evidence by administering LOOKS-GOOD.⁴

³One person who finds this hard to believe is Julia Markovits, who claims in Markovits [21] that “doctors are morally obligated to prescribe the course of treatment their evidence tells them is most likely to cure their patients, not the treatment that (against all evidence) happens to be best” (219). I will discuss Markovits’s view at length in chapter two.

⁴As Markovits does (see the previous footnote).

There is a lesson to be drawn here. If it really is true that Dana is not obligated to administer REALLY-IS-GOOD, a natural explanation of that fact would be: Dana is not obligated to administer REALLY-IS-GOOD because she does not believe that she is obligated to administer that drug, and she is reasonable not to believe that. This explanation suggests the following general principle, which I call (Epistemic-Sensitivity): one cannot be obligated to perform a certain action if one does not believe that one is under that obligation, and one is epistemically permitted not to believe that. In other words, obligations are “epistemically-sensitive.”

It follows from (Epistemic-Sensitivity) that NCFPO ignorance is not possible. To be non-culpably ignorant of the fact that one is obligated to do X, three conditions must be satisfied: (a) one must be obligated to do X, (b) one must not believe that one is obligated to do X, and (c) one must be epistemically permitted not to believe that one is obligated to do X.⁵ Condition (b) must be satisfied because any person who believes that she is obligated to do X will not count as being ignorant of the fact that she is obligated to do X.⁶ And condition (c) must be satisfied because any person who satisfies (a) and (b) but fails to sat-

⁵As Nicholas Sturgeon has pointed out to me, the view of ignorance that I assume here implies that a person cannot be ignorant of a truth that she believes. But this might strike one as an implausible consequence. Why should true belief *alone* be incompatible with ignorance? Suppose, for example, that one forms a true belief in *p*, even though one lacks good reason to believe *p*. Does the fact that one’s belief is true entail that one does not count as ignorant of *p*? It may seem that it does not.

This is a fair objection, which does put pressure on the idea that the account of ignorance that I presuppose above (and elsewhere in this dissertation) completely captures the folk concept of ignorance. That said, whether my account of ignorance captures that concept completely is not important for my purposes. The word “ignorance” could be entirely eliminated from my discussion. In its place, I could speak instead in terms of truths that one fails to believe and truths that one is epistemically permitted to fail to believe. The issues that I discuss will be just as pressing when reformulated in these terms, and my arguments will still work. For the sake of simplicity, however, I will speak as if the (arguably oversimplified) account of ignorance I assume here is correct.

⁶Of course, someone who believes for bad reasons that she is obligated to do X could be ignorant of the reasons why she is obligated to do X. But she couldn’t count as ignorant of the fact that she is obligated to do X.

isfy (c) would *culpably* ignorant of her obligation to do X. Note, however, that if (Epistemic-Sensitivity) is true, then any person who satisfies conditions (b) and (c) cannot also satisfy condition (a). If (Epistemic-Sensitivity) is true, then any person who suspends belief in the claim that she is obligated to do X, and who is epistemically permitted to do so is, in fact, *not* obligated to do X. It follows, then, that NCFPO ignorance is impossible.

Because (Epistemic-Sensitivity) poses a threat to the claim that NCFPO ignorance is possible, we face a problem. The problem is that both (Epistemic-Sensitivity) and the view that NCFPO ignorance is possible are quite plausible. As we have just seen, however, both cannot be true. But which is correct? Is (Epistemic-Sensitivity) true? Or is NCFPO ignorance possible?

1.3 Question 2: Assuming that NCFPO ignorance is possible, do the NCFPO ignorant deserve blame for acting wrongly?

As I have just noted, there is a philosophical debate that needs to be settled in order to establish that NCFPO ignorance is possible. But why should we care whether NCFPO ignorance is possible?

Setting aside the intrinsic philosophical interest of the concept of NCFPO ignorance, there is a down-to-earth reason to care whether NCFPO ignorance is possible. Many philosophers have argued that if NCFPO ignorance is possible, it constitutes an excuse for wrongdoing. If these philosophers are correct, then the possibility of NCFPO ignorance would have implications for moral practice.

Everyone treats some varieties of non-culpable ignorance as excuses for acting wrongly. Suppose, for example, that I greet a friend with a warm slap on the back. Unbeknownst to me, however, my friend is suffering from a sunburn,

and my slap causes him excruciating pain. If I had no reason to believe that my friend was suffering from a sunburn, it would be quite unfair for him to blame me for hurting him. That I hurt him was unfortunate, but I do not deserve to be the target of indignation or anger for doing so. My non-culpable ignorance excuses me, even if what I did was wrong.⁷

Sometimes, then, non-culpable ignorance constitutes an excuse for acting wrongly. But which types of non-culpable ignorance constitute an excuse for acting wrongly? In the sunburn case, I slapped my friend on the back because I was ignorant of a certain *non-moral* fact—namely, that my friend had a sunburn. But sometimes people are non-culpably ignorant of obligations that they are under, but not because they lack knowledge of some relevant non-moral facts. What should we say about these people? Are they blameworthy for acting wrongly?

Consider, for example, the ancient slaveholder discussed in Rosen [33]:

In the ancient Near East in the Biblical period . . . no one denied that it was bad to be a slave, just as it is bad to be sick or deformed. The evidence suggests, however, that until quite late in antiquity it never

⁷One might object that there is no need to appeal to the notion of non-culpable ignorance in order to explain why I am not blameworthy for hurting my sunburned friend. Blame would be appropriate only if it was wrong for me to slap my friend on the back. But one might argue that it was *not* wrong for me to slap my friend on the back if I had no reason to believe that his back was sunburned.

There are two reasons why one might raise this objection. First, one might raise the objection because one believes that no action can be wrong for one to perform—that is, no action can be such that one is obligated not to perform it—if, for good reason, one does not believe that it would be wrong for one to perform. But this would be to defend the claim that obligations are epistemically-sensitive, which implies that NCFPO ignorance is impossible; recall, however, that we are working now under the assumption that NCFPO ignorance *is* possible, so this objection is dialectically inappropriate. Second, even if one believes that NCFPO ignorance is possible, one might maintain that the sunburn case is not an example of it. But if this is the source of the objection, then the general point I am making here—that everyone treats some varieties of non-culpable ignorance as excuses for acting wrongly—can still be made, so long as the reader admits that there are some cases in which people who are non-culpably ignorant of obligations are excused for acting wrongly because of their ignorance. The reader should feel free, then, to substitute another example of NCFPO ignorance for the sunburn case.

occurred to anyone to object to slavery on grounds of moral or religious principle. So consider an ordinary Hittite lord. He buys and sells human beings, forces labor without compensation, and separates families to suit his purposes. Needless to say, what he does is wrong. The landlord is not entitled to do these things. But of course he thinks he is ... Let us suppose he asked himself, 'Is this wrong (shameful, base)?'—and coming up with nothing, went on as before. There may be cogent arguments against slavery from premises our Hittite accepts, or might come to accept. But if there are, they are not obvious, and his failure to discover them need not signal recklessness or negligence on his part. To the contrary, it is compatible with his having been unusually reflective. Given the intellectual and cultural resources available to a second millennium Hittite lord, it would have taken a moral genius to see through to the wrongness of chattel slavery.⁸

Arguably, ancient slaveholders like the one Rosen describes were non-culpably ignorant of the fact that they were obligated to free their slaves. But their ignorance of this obligation did not derive from non-culpable ignorance of any of the non-moral features of slaveholding. Would such slaveholders deserve blame, then? Or would the slaveholders' non-culpable ignorance excuse them?

The question here generalizes. One might wonder not only about ancient slaveholders, but anyone who suffers from NCFPO ignorance. Does NCFPO ignorance constitute an excuse for wrongdoing? Some philosophers, such as Rosen himself, answer that it does.⁹ However, other philosophers, such as Elizabeth Harman and Matthew Talbert, disagree; Harman and Talbert maintain

⁸Rosen [33], 64-66.

⁹See Rosen [33], Rosen [34], and Rosen [35].

that NCFPO ignorance does not constitute an excuse for wrongdoing.¹⁰ Given the central place that blame occupies in our moral lives, it is quite important for this debate to be settled.¹¹ For if NCFPO ignorance is possible, then the view that NCFPO ignorance is an excuse for wrongdoing will tell us something important about how we ought to treat (or at least feel towards) many people.

1.4 Question 3: Assuming that NCFPO ignorance is possible, how extensive can a mature person's NCFPO ignorance be?

If NCFPO ignorance not only exists but constitutes an excuse for wrongdoing, then that should have an impact on our moral practice. But how much of an impact should it have? To determine that, we need a sense of the *distribution* of NCFPO ignorance. There are, however, at least three different questions we may ask about how NCFPO ignorance is and can be distributed. We may ask: how many people are NCFPO ignorant? We may ask: how common are situations in which a person is NCFPO ignorant? Or we may ask: how extensive can a person's NCFPO ignorance be?

The first two of these questions are difficult to address philosophically, as their answers depend on contingent facts about the particular people that exist and what their epistemic situations are like. The third question, however, is more amenable to philosophical analysis, as it deals only with the *possible scope* of NCFPO ignorance.

¹⁰Harman [14], Talbert [42].

¹¹In this dissertation, I will proceed under the assumption, accepted by most philosophers, that people are sometimes blameworthy for acting wrongly. Some philosophers, such as free will skeptics, reject this view; according to these philosophers, nobody is ever blameworthy for anything. I admit that there are powerful arguments for this position, but I am not alone in thinking that the jury is still out on whether the free will skeptics' arguments succeed. It is worth considering, then, as I do in this dissertation, how things stand if people *do* sometimes deserve blame.

As a way to introduce some of the issues that must be addressed in order to answer this third question, it is helpful to consider an analogy to a different, but structurally similar question: given that it is possible for people to be non-culpably ignorant of facts about the external world, how extensive can a person's non-culpable ignorance of the external world be? In other words, about which truths concerning the external world can a person be ignorant? Skeptics will, of course, answer that every truth concerning the external world (if there are any truths about the external world) is a potential object of non-culpable ignorance, since, according to the skeptics, *every* agent is entirely ignorant of the external world, and non-culpably so. For those who reject skepticism, however, the answer is less clear. Rejecting skepticism requires one to say only that some people are sometimes justified in believing some truths about the external world. But one can maintain that even if one holds that some other people are justified in suspending belief in all external world propositions. Is it the case, then, that simple propositions about the external world (such as claims about one's immediate environment) are such that no person could reasonably suspend belief in them? Or is it possible that a person could reasonably suspend belief in all claims about the external world?

If NCFPO ignorance is possible, then similar questions arise about the possible scope of NCFPO ignorance. For example, we might wonder whether, for every mature person, there is some claim of the form *I am obligated to do/not do X* that she is epistemically required to believe. (I will not ask a similar question about the immature, since they might be epistemically permitted to suspend belief in all first-person obligation claims simply because their cognitive faculties are underdeveloped.) In other words, is it the case that every mature person is epistemically required to believe some first-person obligation claims? Or is

it the case that some mature people can reasonably suspend belief in all first-person obligation claims?

One might argue that there is a privileged class of moral truths that are so obvious that every mature person is epistemically required to believe them. Perhaps truths such as *torturing animals for fun is wrong* are among these privileged moral truths. To deny such truths, or even simply to fail to believe them, would be a colossal epistemic blunder. It is also conceivable, however, that there is no privileged class of moral truths at all. Maybe *some* mature people are permitted to suspend belief in all moral truths. To motivate taking this possibility seriously, it is helpful to think again about external world skepticism. Most of the time, it seems that there is a privileged class of truths about the external world that are so obvious that nobody could reasonably suspend belief in them—a class that perhaps includes truths about one's immediate environment, or truths such as "the sun comes up every day." But many philosophers find it not too difficult to get into a state of mind in which suspending belief in even these truths seems reasonable. That is what makes external world skepticism so tricky; even after one has convinced oneself that skepticism is not the *only* epistemically permissible stance to take towards claims about the external world, there is a temptation to regard it still as an epistemic *option*. Moral skepticism can seem tricky for a similar reason; even if one thinks that it is epistemically permissible for one to maintain one's moral views, there is a temptation to say that some people might nonetheless be epistemically permitted to suspend belief in all moral claims whatsoever.

But which is it: are some moral truths so obvious that no person can reasonably suspend belief in them, or is it possible for some people to rationally adopt a skeptical stance towards morality? The answer to this question will help us

not only to determine how widespread NCFPO ignorance *can* be, but it also has the potential to help us to determine how widespread NCFPO ignorance actually *is*. Learning about the conditions under which an agent can rationally adopt a skeptical stance towards morality (if it is possible to do so) might give us information that can be combined with empirical facts about the type of conditions that agents actually find themselves in to tell us something about how NCFPO ignorance is *actually* distributed.

1.5 The dissertation, summarized

We have now seen three questions about NCFPO ignorance. The first asks whether NCFPO ignorance is possible. The second asks whether NCFPO ignorance, if possible, can constitute an excuse for wrongdoing. The third asks how extensive a mature person's non-culpable ignorance of her obligations can be, assuming that NCFPO ignorance is possible.

This dissertation defends three main conclusions, which provide answers to these questions. The conclusions are the following:

1. NCFPO ignorance is possible.
2. NCFPO ignorance constitutes an excuse for wrongdoing.
3. Any first-person obligation claim is a potential object of non-culpable ignorance, since moral skepticism can be a reasonable position for some mature people to adopt.

What follows is a high-level overview of my arguments for each conclusion. I will present these arguments in the order that they occur in the dissertation, which, because of relations of logical dependence that hold between the arguments, will be different than the order of the questions above.

1.5.1 NCFPO ignorance is possible

As I have already explained, if (Epistemic-Sensitivity) is true—that is, if one cannot be obligated to perform an action that one permissibly does not believe that one is obligated to perform—then NCFPO ignorance is impossible. What I argue in chapter two, however, is that (Epistemic-Sensitivity) is false. Showing that (Epistemic-Sensitivity) is false does not by itself show that NCFPO ignorance is possible. But the argument that I use to refute (Epistemic-Sensitivity) is easily extended to establish this stronger claim as well.

I will start by summarizing my argument against (Epistemic-Sensitivity). To understand the basic idea behind my argument against (Epistemic-Sensitivity), it is helpful to first understand why the most natural approach to arguing against (Epistemic-Sensitivity) does not work. That approach is to try to construct a counterexample to (Epistemic-Sensitivity) by doing the following:

First, take someone who is obligated to perform a certain action X , and who knows that she is obligated to perform X ,

Second, subtract her evidence that she is obligated to perform X ,

And, third, subtract her belief that she is obligated to perform X .

For example, we can imagine a professor who knows that she is obligated to be in her office at 9 AM, since she promised to meet a student at that time. We might then subtract the professor's evidence and belief that she is obligated to be in her office at 9 AM. If her obligation remains, despite the fact that she reasonably does not believe that she is under that obligation, we would have a counterexample to (Epistemic-Sensitivity).

One problem with this strategy is that it is not obvious that it is possible to subtract the professor's evidence and belief that she is obligated to be in

her office at 9 AM without also removing her obligation. Indeed, it is likely that people who feel moved by the cases that support (Epistemic-Sensitivity) (such as the case of Dana and Pat, discussed above in section 1.2) will also feel that performing this subtraction will remove the professor's obligation as well. Thus, the obvious strategy for arguing against the view that obligations are epistemically-sensitive is unlikely to move anyone who was not already inclined to deny that view. Of course, the fact that an argument is dialectically limited does not mean that it is unsound. But it does mean that we should search for an argument with broader appeal.

In chapter two, I sketch such an argument. My strategy is to show that (Epistemic-Sensitivity) implies that satisfying a fairly minimal set of epistemic conditions is sufficient for a person not to be subject to any obligations whatsoever. There is, therefore, a serious cost to accepting (Epistemic-Sensitivity), a cost which many philosophers will not want to pay. My argument goes as follows:

1. By satisfying a relatively minimal set of epistemic conditions, a mature person can reasonably suspend belief in all claims of the form *I am obligated to do/not do X*.
2. Any such person would be under no obligations whatsoever if (Epistemic-Sensitivity) were true.
3. But satisfying a relatively minimal set of epistemic conditions cannot ensure that a person is under no obligations whatsoever.
4. Therefore, (Epistemic-Sensitivity) is false.

This strategy improves on the first strategy discussed above in two ways. First, my premises can be motivated in such a way that even those initially

sympathetic to the idea that obligations are epistemically-sensitive can feel their force. Second, my strategy is not susceptible to the objection that a person's obligations disappear when her corresponding moral evidence is removed; it's not very plausible, once the "relatively minimal set of epistemic conditions" mentioned above are spelled out, that a person's obligations do disappear once she satisfies those conditions.

Not only does the argument above show that (Epistemic-Sensitivity) is false, but it turns out that it also provides an argument that NCFPO ignorance is possible. Someone who satisfies the relatively minimal epistemic conditions mentioned above would still be subject to certain obligations; however, she would be non-culpably ignorant of those obligations. Thus, she would count as NCFPO ignorant.

A crucial component of the argument above—premise 1.—is not defended in chapter two. That premise is defended in chapter three, where the same premise is used to provide an answer to the question: assuming that NCFPO ignorance is possible, of what obligations can a mature person be non-culpably ignorant? This is what I will discuss next.

1.5.2 Moral skepticism can be a reasonable position for some people to adopt

My basic argument for the view that moral skepticism—that is, suspending belief in all (non-trivial) obligation claims—can be a reasonable position for some people to adopt is inspired by a view that has been growing in popularity among epistemologists. Many epistemologists have recently suggested that a person who becomes a thoroughgoing external world skeptic cannot be epis-

temically compelled to abandon her skepticism. A natural argument for this view goes as follows: anyone who has become a thoroughgoing external world skeptic has abandoned all of her beliefs about the external world. However, it is impossible to support a belief about the external world without appeal to other beliefs about the external world, so no argument *against* external world skepticism can get a foothold against the skeptic. Thus, the thoroughgoing external world skeptic cannot be epistemically compelled to abandon her position, which means that her skepticism counts as *reasonable*.

As is, this argument faces a number of objections. Even so, there is something to the argument; indeed, I will argue in chapter three that the moral analogue of a more sophisticated version of this argument helps to show that moral skepticism can be a reasonable position for some people to adopt.

Establishing this requires several steps. The first step is to clean up the argument above and to translate it into a moral idiom. What results is the following argument. Someone who becomes a thoroughgoing moral skeptic—that is, someone who comes to suspend belief in all claims of the form *Xing is obligatory*, except those trivial claims that are logically entailed by her non-moral beliefs—no longer has any substantive moral beliefs at all. But that means that no argument for a claim of the form *Xing is obligatory* can gain a foothold against the skeptic. Thus, no moral skeptic can be epistemically compelled to abandon her skepticism *by any argument*.

Notice, however, that this does not yet imply that every thoroughgoing moral skeptic cannot be epistemically compelled to abandon her skepticism, for it might be that something *other* than an argument has the power to do that. Some philosophers believe, for example, that there are moral experiences. One might maintain that the presence of such an experience can epistemically re-

quire a person to believe a claim of the form *Xing is obligatory*. Or one might maintain that the presence of a moral intuition with content *Xing is obligatory* can epistemically require a person to believe a claim of that form.

What is necessary, then, is to show that moral experiences or intuitions are not sufficient to cure every case of moral skepticism. This I do by suggesting that even if it were the case that everybody has moral experiences or intuitions, there are scenarios in which a person might find herself that would prevent such experiences or intuitions from epistemically requiring her to abandon her skepticism. A person who really does have moral experiences, for example, might have good reason to believe that she does not have such experiences, and this, I claim, would be sufficient to undermine any justificatory force those experiences would have. Additionally, a person who has moral intuitions might also have “anti-moral” intuitions, and I argue that the presence of such intuitions would prevent her moral intuitions from epistemically compelling her to abandon her skepticism.

The arguments summarized above show that it is possible for a mature skeptic to find herself in a position in which no arguments, moral experiences, or intuitions generate an epistemic requirement that she abandon her skepticism. But this conclusion does not obviously entail that it is possible for skepticism to be *reasonable* or *epistemically permissible*, since it might be argued that in order for a person’s doxastic state to be reasonable or epistemically permissible, her transition to that state must have been epistemically respectable; this is what I refer to as the *Transition Principle*. If the Transition Principle is true, then it is necessary to know *how* a person became a skeptic in order to determine whether her skepticism is epistemically permissible; knowing that she is not epistemically compelled to abandon her skepticism by any argument, experience, or intuition

would not, by itself, be sufficient. One might worry, though, that there is no epistemically respectable way to transition to skepticism, since the existence of an epistemically respectable transition to skepticism would require there to be a compelling argument for adopting skepticism, which arguably does not exist.

In response, I concede that this objection might very well show that one's route to skepticism matters in determining whether one's skepticism is epistemically permissible. But I argue that there are, nonetheless, epistemically respectable ways to transition to skepticism.

The ultimate conclusion of chapter three, then, is that it is possible for moral skepticism to be reasonable for some people. This not only provides the missing premise in my argument that NCFPO ignorance is possible, but it also provides an answer to the third question discussed above—assuming that NCFPO ignorance is possible, how extensive can a mature person's non-culpable ignorance of her obligations be?—by showing that a person's non-culpable ignorance can be as broad as possible: it can extend to *all* of her moral obligations.

1.5.3 NCFPO ignorance constitutes an excuse for wrongdoing

Once the possible scope of NCFPO ignorance has been established, the next question to discuss is how one ought to treat (or at least feel towards) people who are NCFPO ignorant. If an NCFPO ignorant person acts wrongly, does she deserve blame? Or is she to be excused because she reasonably did not believe her action to be wrong?

Let the *Forgiving View* be the view that NCFPO ignorance constitutes an excuse for wrongdoing; let the *Unforgiving View* be the view that denies that NCFPO ignorance constitutes an excuse for wrongdoing. The debate between defenders of the Forgiven and Unforgiving Views seems to have reached a

stalemate, which recent contributions by Elizabeth Harman, Gideon Rosen, and Matthew Talbert have not been able to break. In the last chapter of this dissertation, however, I provide a new strategy for moving the debate forward.

Defenders of the Forgiving and Unforgiving Views often attempt to argue for their preferred view by formulating examples involving NCFPO ignorant agents who are supposed to be either clearly blameworthy or clearly blameless for acting wrongly. The problem with this strategy is that the examples provided thus far rarely move many philosophers on the other side of the debate; each side disagrees with the other side's judgments about these cases. What is needed to move the debate forward is to identify some more neutral assumptions that can be leveraged into an argument for either the Forgiving or the Unforgiving View.

This is precisely what I attempt to provide in the chapter four. The first phenomenon on which I focus is the fact that non-culpable non-moral ignorance excuses sometimes, but not always. Often, a person's non-culpable ignorance of relevant non-moral features of her action prevents her from deserving blame for acting wrongly. As I have already noted, if I give a friend a warm slap on the back, but, unbeknownst to me, my friend is suffering from a sunburn, then, while it is unfortunate that my slap will cause him pain, I do not deserve to be blamed for hurting him. The reason why, it seems, is that I was non-culpably ignorant of the fact that slapping my friend on the back would hurt him. But while non-culpable ignorance of non-moral facts sometimes excuses, it doesn't *always* excuse. A driver who intentionally hits a pedestrian deserves blame for doing so even if the driver is ignorant of the pedestrian's name. Both defenders of the Forgiving View and defenders of the Unforgiving View can accept these observations.

The second phenomenon on which I focus concerns a series of cases that I call the Comparison Cases. The Comparison Cases involve an agent with varying degrees of evidence that he is acting permissibly, though in reality he is acting wrongly. The Comparison Cases suggest that the greater the amount of evidence that an agent has that he is acting permissibly, the lower the *amount* of blame that he deserves is.

I will argue in chapter four that the Forgiving View, when supplemented with two closely related claims, provides a compelling explanation of each of these phenomena: it accounts for standard intuitions about when non-culpable non-moral ignorance excuses and when it does not, and it also provides a plausible account of the Comparison Cases. If it could be shown that the Forgiving View provides the *best* explanation of these phenomena, that would provide strong support for that view. To explore whether it does, I consider three alternative explanations that might be defended by opponents of the Forgiving View.

Ultimately, I argue that each of these alternative explanations fails. In the absence of a better explanation of the two phenomena mentioned above, I claim that the Forgiving View should be accepted on abductive grounds.

1.6 Generalizing to other types of moral or ethical ignorance

By now, I hope that it is clear why the notion of NCFPO ignorance is philosophically important. But even if this is clear, one might wonder whether questions similar to those discussed above arise for other types of moral or ethical ignorance. In this section, I will explain that while analogues of all of the questions discussed above can be raised concerning other types of moral or ethical ignorance, there is reason to think that NCFPO ignorance is a particularly interesting

type of ignorance.

One type of ignorance that raises some analogous questions is non-culpable first-person ignorance of what is good or bad for a person. Some things are good for us—like eating broccoli, or reading good philosophy—while others are bad for us—like binge-drinking, or feeling depressed. But is it always epistemically accessible to us what is good or bad for us? It might seem not. If there are objective facts about what is good or bad for us, then the fact that we are cognitively limited beings suggests that the facts about what is good or bad for us can outstrip our ability to recognize those facts.

On second thought, though, it can seem that the facts about what is good or bad for us *cannot* outstrip our ability to recognize what is good or bad for us. It is often said that what you don't know can't hurt you. A more plausible, but essentially similar, view says that something cannot be bad for a person if that person reasonably does not believe that it is bad for her—in other words, badness is epistemically-sensitive.¹² The view that badness is epistemically-sensitive provides a compelling explanation of many people's reactions to cases involving unknown betrayal. As Nagel [25] points out, it can seem that a person who has been betrayed behind her back is not harmed merely in virtue of being betrayed; she is harmed only if she finds out that she has been betrayed. Anyone who finds this view plausible will be attracted to the idea that badness for a person is epistemically-sensitive.¹³ If it really were true, however, that badness for a person is epistemically-sensitive, then it would follow that it is impossible for one to be non-culpably ignorant of the fact that something is bad for one. As-

¹²Why is this view more plausible? Because failing to believe that something is bad for you—which would prevent you from knowing that thing is bad for you—is not always sufficient to prevent that thing from being bad for you. The view that badness is epistemically-sensitive is compatible with admitting this; however, the view that what you don't know can't hurt you is not.

¹³See Nagel [25], 76-78.

suming that goodness for a person is also epistemically-sensitive, it would also be impossible for a person to be non-culpably ignorant of the fact that something is good for her.

The view that goodness and badness are epistemically-sensitive, then, raises a challenge to the existence of non-culpable first-person ignorance of what is good or bad for a person that is similar to the challenge, discussed above in section 1.2, to the existence of NCFPO ignorance. So far, the issues about NCFPO ignorance are simply special cases of more general issues concerning moral ignorance.

One can also show that issues analogous to those discussed in section 1.4 (“Question 3: Assuming that NCFPO ignorance is possible, how extensive can a mature person’s NCFPO ignorance be?”) arise for non-culpable first-person ignorance of what is good or bad for a person. In that section, I noted that the possibility of NCFPO ignorance raises the question of whether, for every person, there is a privileged class of truths about that person’s obligations that she is rationally required to believe. Similarly, one might wonder whether, for every person, there is a privileged class of truths about what would be good or bad for that person that she is rationally required to believe. It is natural to believe that there are some such truths—for example, the truth that every person is better off being healthy rather than sick. But one might also think that someone who was a thoroughgoing skeptic about goodness and badness could be rationally permitted to suspend belief in all claims of the form *X is good/bad for me*.

While questions analogous to those discussed in section 1.2 and 1.3 do arise for non-culpable first-person ignorance of what is good or bad for a person, what makes NCFPO ignorance special is that it is more deeply connected with

blame than other types of ignorance are. As I will argue later in this dissertation, there is good reason to believe that NCFPO ignorance always constitutes an excuse for wrongdoing. This is what differentiates NCFPO ignorance from other types of moral ignorance. Other types of moral ignorance—with one possible exception—constitute an excuse for wrongdoing at most some of the time.

Take, for example, non-culpable first-person ignorance of what is good or bad for a person. A case could be made that non-culpable first-person ignorance of what is good or bad for a person constitutes an excuse for wrongdoing some of the time. This would be clearest if everyone has moral obligations to themselves—such as an obligation to refrain from seriously unhealthy activities. Even if everyone does have such obligations, it could plausibly be argued that some people who have indulged in seriously unhealthy activities do not deserve blame for doing so. One could argue, for example, that someone who developed lung cancer in the 1950s because of a heavy smoking habit would, nonetheless, not deserve blame for harming herself by smoking, since people back then did not know that smoking causes cancer.

Still, even if one accepts that non-culpable first-person ignorance of what is good or bad for a person can excuse wrongdoing some of the time, it excuses wrongdoing in only a small range of cases—only when the agent's ignorance of her action's consequences for *herself* excuse her for wrongdoing. As far as I can see, something similar is true of every other type of moral ignorance: no type of moral ignorance other than NCFPO ignorance has the potential to excuse such a wide range of actions. The one possible exception to this generalization is non-culpable ignorance of the goodness or badness of the consequences of one's actions. One could argue that if an agent is non-culpably ignorant of whether her action's consequences are bad, then that would always constitute an excuse for

wrongdoing. This excuse has the potential to apply to a wide range of actions. It is possible, however, that this excuse is parasitic on the excuse provided by NCFPO ignorance; it may be that non-culpable ignorance of whether one's action is bad provides an excuse only because it prevents one from knowing that one's action is wrong. But I will not try to determine whether this is the case.

At any rate, the take-away of this section is that while many of the same issues that arise for NCFPO ignorance do arise for other types of moral ignorance, NCFPO ignorance does have a claim to being a particularly interesting type of moral ignorance. That said, it is likely that many of the arguments that I offer in this dissertation will generalize to answer some of the analogous questions that arise for other types of moral ignorance. I will not attempt to generalize my argument here, but this is a potential avenue for future research.

CHAPTER 2
DOES A PERSON'S EPISTEMIC STATE CONSTRAIN HER
OBLIGATIONS?

2.1 Introduction

Philosophers often maintain that NCFPO ignorance—that is, non-culpable ignorance of a true claim of the form *I am obligated to do/not do X*—is an excuse for wrongdoing.¹ But is NCFPO ignorance even possible? In this chapter, I will explore a plausible view of the connection between obligations and evidence that implies that it is not.

The view in question is what I will call (Epistemic-Sensitivity). (Epistemic-Sensitivity) states that for all actions *X*, a person *S* cannot be obligated to do *X* if *S* does not believe that she is obligated to do *X*, and *S* is epistemically permitted not to believe that. To see why (Epistemic-Sensitivity) implies that NCFPO ignorance is impossible, assume for reductio that there is a person *S* who is non-culpably ignorant of the fact that she is obligated to do *X*. If *S* is non-culpably ignorant of the fact that she is obligated to do *X*, she must not believe that she is obligated to do *X*; otherwise, she would not count as *ignorant* of the fact that she is obligated to do *X*.² To be non-culpably ignorant of the fact that she is obligated to do *X*, *S* must also be *epistemically permitted* not to believe that she is obligated to do *X*; otherwise, she could be only *culpably* ignorant of her obligation. But if *S* does not believe that she is obligated to do *X*, and she is epistemically permitted not to believe that, then (Epistemic-Sensitivity) implies that *S* is *not* obligated to do *X*—a contradiction, since we stipulated that *S* is obligated to do *X*. Thus, we

¹See chapter four for an extended discussion of this view.

²Note, however, that while a person who is obligated to do *X* and who believes that she is obligated to do *X* cannot count as ignorant of the fact that she is obligated to do *X*, she might still count as ignorant of the *reason* why she is obligated to do *X*.

must reject the assumption we started with: that NCFPO ignorance is possible.

(Epistemic-Sensitivity) is one way of capturing the idea, defended by a number of philosophers recently, that a person's epistemic state constrains her obligations. In order to evaluate (Epistemic-Sensitivity), then, we should begin by exploring the arguments that have been offered for and against the view that a person's epistemic state constrains her obligations. If there is a compelling argument that a person's obligations are *not* constrained by her epistemic state, then (Epistemic-Sensitivity) is dead in the water, since it is a view according to which a person's obligations are so constrained. On the other hand, if there is reason to believe that a person's obligations do constrain her epistemic state, then it is worth exploring whether those arguments will establish in what particular way they do—in particular, it's worth considering whether those arguments might be used to support (Epistemic-Sensitivity).

In the first part of this chapter, I will begin by discussing the best extant arguments for and against the view that a person's epistemic state constrains her obligations. I will show that none of these arguments—neither the arguments for the view, nor the arguments against the view—are convincing.

While we shouldn't be convinced by any of the extant arguments for or against the view that a person's epistemic state constrains her obligations, I do believe that there is good reason to think that a person's epistemic state does not constrain her obligations in the way that (Epistemic-Sensitivity) suggests. In the second part of this paper, I will outline an argument against (Epistemic-Sensitivity). As we shall see, my argument against (Epistemic-Sensitivity) not only undermines the challenge that claim poses to the possibility of NCFPO ignorance, but the argument can be extended to show that NCFPO ignorance is possible.

2.2 Clarifying the view that a person's epistemic state constrains her obligations

Philosophers have long wondered whether a person's epistemic state constrains the obligations to which she is subject. In the utilitarian tradition, for example, it is easy to find discussions of whether a person's action should be evaluated by exploring the *actual* consequences of her action or rather the *consequences she had reason to anticipate*.³ Suppose that someone flicks a switch, thinking—for good reason—that it would turn on the lights, but instead it sets off a bomb connected to the switch via a mechanism that she had no idea existed. Has she done something wrong? A simple form of utilitarianism that looks only at the *actual* consequences of a person's action would say that she has (assuming, of course, that not flicking the switch would have resulted in a better outcome). But a form of utilitarianism according to which the moral status of a person's action is determined only by the consequences *that she had reason to anticipate* could say that what she did was morally permissible, since she had no reason to think that flipping the switch would set off a bomb.

It is not only utilitarians, however, who are interested in the question of whether a person's obligations are constrained by her epistemic state. H. A. Prichard, for example, discussed this question in his well-known article "Duty and Ignorance of Fact."⁴ In the end, Prichard accepted what he called "the subjective view of the basis of an obligation," which states that "obligation depends on our being in a certain attitude of mind towards [our] situation in respect of knowledge, thought, or opinion."⁵ He is not alone among non-utilitarians in holding that some view like this is appropriate. Another non-utilitarian philoso-

³For a start, see Feldman [11] and the references therein.

⁴Prichard [30].

⁵Prichard [30], 89.

pher who holds a view like this is Julia Markovits, whom I will discuss shortly.⁶

Now, it is important to note at the outset that the claim that a person's epistemic state constrains her obligations stands in need of further clarification. As is, the claim does not say, for example, *which* aspects of a person's epistemic state constrain her obligations. Is it her evidence concerning what she is obligated to do? Is it her evidence concerning what it would be best for her to do?⁷ Is it her evidence concerning what she has reason to do? Or is it what she is in a position to know concerning her obligations, what it would be best for her to do, or what she has reason to do? It also does not clarify the nature of the constraint that a person's epistemic state imposes on her obligations. Is she required to have a certain kind of evidence in order to be under a certain obligation? Or is she required *not* to have a certain kind of evidence in order to be under a certain obligation? I will return to these important questions in the next section. But for now I want to note that even though the view that a person's epistemic state constrains her obligations stands in need of further clarification, as stated it is already importantly different from two views that are easily conflated with it. To avoid misunderstanding, it is important to point out the relevant differences.

A first view that might be conflated with the view that a person's epistemic state constrains her obligations says that a person's *doxastic* state constrains her obligations. To say that a person's epistemic state constrains her obligations is to say that (at least some of) the facts about what she does or doesn't know, what she has or doesn't have evidence for, and what she is epistemically obligated or permitted to believe or not to believe constrain her obligations. But to say that a person's doxastic state constrains her obligations is to say that *only* what she believes or doesn't believe constrains her obligations. These are very different

⁶Markovits [21].

⁷As Markovits [21] suggests (219).

claims. The view that a person's doxastic state constrains her obligations is open to the objection that a person's beliefs about what she is obligated to do cannot, by themselves, change her obligations. (It's not the case, for example, that an agent would cease to be obligated to fulfill her obligations if she were able to simply stop believing that she was so obligated.) But this objection does not face the view that a person's epistemic state constrains her obligations; it is important, then, to keep these views separated in one's mind.

A second view, which is more easily confused with the claim that a person's epistemic state constrains her obligations, is what Graham [13] calls "evidence-subjectivism."⁸ Evidence-subjectivism states that "necessarily, a person has the moral obligations that she has at a time solely in virtue of facts about her evidential situation at (or prior to) that time."⁹ This view is not the same as the claim that a person's epistemic state constrains her obligations. Evidence-subjectivism includes a claim about the *source* or *grounds* of a person's moral obligations; however, the claim that there exists a certain kind of *constraint* on a person's obligations does not. To say that a person's epistemic state constrains her obligations is to say only that being in a certain kind of epistemic state (or perhaps *not* being in a certain kind of epistemic state) is a *necessary condition* for having certain obligations. (Compare: to say that the amount of money in my bank account constrains my ability to travel to Paris just is to say that having more than a certain amount of money in my bank account is a necessary condition for me to travel to Paris.) As I have already noted, which epistemic state one must be in in order to have certain obligations is a matter that needs to be clarified. But my point here is simply that no matter how one precisifies the view that a person's epistemic state constrains her obligations, one need not think that a person's

⁸Graham [13], 89.

⁹Ibid.

epistemic state *grounds* her obligations. One need not say, for example, that a person's obligation not to torture animals is grounded in her evidence. Instead, one might say that a person's obligation not to torture animals is grounded in the fact that torturing animals causes them pain and distress. Saying that, however, is compatible with accepting that a person must satisfy a certain epistemic condition in order for her to be so obligated.

2.2.1 Arguments that a person's epistemic state constrains her obligations

As I have already indicated, (Epistemic-Sensitivity) is one way of capturing the more general idea that a person's epistemic state constrains her obligations. As we will see shortly, there are other ways of capturing this idea. But since (Epistemic-Sensitivity) is one of them, properly evaluating (Epistemic-Sensitivity) will require looking at the arguments for and against this general view. If there is a compelling argument against the idea that a person's epistemic state constrains her obligations, then that argument would be a compelling reason to reject (Epistemic-Sensitivity). If, on the other hand, there is a compelling argument in favor of the idea that a person's epistemic state constrains her obligations, then it is worth attending to that argument in order to see whether that argument could be used to support (Epistemic-Sensitivity).

I will begin my evaluation of the arguments that have been offered for and against the view that a person's epistemic state constrains her obligations by presenting, in this section, the two best arguments for this view.

The main argument in favor of the view that a person's epistemic state constrains her obligations is what I will refer to as the "appeal to intuitions." The

basic idea here is that the view that a person's epistemic state constrains her obligations best explains our intuitions about certain cases in which someone has misleading evidence about what it would be best for her to do.

Julia Markovits provides an argument along these lines. She imagines a doctor who has misleading evidence about two drugs: LOOKS-GOOD and REALLY-IS-GOOD.¹⁰ The doctor has good reason to believe that administering LOOKS-GOOD will save her patient, but in fact administering it would kill him. REALLY-IS-GOOD, on the other hand, really would save her patient. But the doctor has strong misleading evidence to the effect that REALLY-IS-GOOD would kill him.

Question: is the doctor morally obligated to administer REALLY-IS-GOOD? Markovits says that she is not. Assume that's right. What explains why the doctor isn't obligated to administer REALLY-IS-GOOD? The answer seems clear: the doctor isn't obligated to administer REALLY-IS-GOOD because her evidence suggests that doing so would not help her patient. That is, the best explanation of why the doctor is not obligated to administer REALLY-IS-GOOD seems to rely on the idea that her obligations are constrained by her epistemic state. Inference to the best explanation leads to the conclusion that the doctor's epistemic state really does constrain her obligations.

A related, though importantly different, argument for the claim that an agent's obligations are constrained by her epistemic state has been spelled out by Peter Graham.¹¹ The argument relies on the following case, modeled after a

¹⁰Markovits [21], 219. In the interest of clarity, I have changed the names of the drugs in Markovits's case.

¹¹The argument is found in Graham [13], 96-97. Graham actually formulates the argument as an argument for the claim referred to earlier as "evidence-subjectivism." But, as I noted earlier, evidence-subjectivism is a *stronger* claim than the claim that a person's obligations are constrained by her epistemic state. So the argument also provides an argument for the latter claim as well.

Note that, ultimately, Graham does not endorse the argument.

case due to Frank Jackson:

Jill's patient, John, has a very painful, nonfatal disease. Jill has three drugs: A, B, and C. A would completely cure him; B would only partially cure him; and C would kill him. But, though Jill knows that B would only partially cure John, her evidence leaves it open which of A and C would cure him and which would kill him.¹²

Assume for reductio that a person's epistemic state does *not* constrain her obligations. If that were the case, then it would seem that Jill would be obligated to administer the drug that would cure John—namely, drug A. Now, suppose that Jill knows that an agent's epistemic state does not constrain her obligations; however, since she isn't sure which of A and C would cure John and which would kill him, she knows only that she is obligated to administer either A or C. What, then, should Jill do? A morally conscientious agent, it seems, will never do something that she knows that she is obligated not to do. The problem, however, is that it also seems that the only option that Jill would consider, if she is morally conscientious, is administering B. But as we have already said, Jill *knows* that administering B would be wrong. Strangely, then, Jill cannot simultaneously be morally conscientious and know that an agent's epistemic state does not constrain her obligations. But a morally conscientious agent *should* be able to believe the true moral theory, whatever it is. So, it seems that we must reject the assumption we started with—namely, that a person's epistemic state *does* constrain her obligations.

¹²Graham [13], 96.

The many ways in which a person's epistemic state might constrain her obligations

Having seen two arguments for the view that a person's epistemic state constrains her obligations, I now want to consider whether either (or perhaps both) of these arguments can be used not only to support this general view, but also the more specific thesis of (Epistemic-Sensitivity). I'll begin by considering whether the appeal to intuitions can be used to support (Epistemic-Sensitivity).

To start, I'll provide a spoiler: even if we assume for the sake of argument that a person's epistemic state constrains her obligations—an assumption that I will evaluate in the next section—I don't think that the appeal to intuitions can be used to support (Epistemic-Sensitivity). Or, more exactly, I don't think that the appeal to intuitions can support (Epistemic-Sensitivity) if we mean by "support" something more than "slightly raise the probability of." Establishing that a person's epistemic state constrains her obligations should perhaps raise our credence in (Epistemic-Sensitivity) to some degree, since (Epistemic-Sensitivity) is one way in which a person's epistemic state could constrain her obligations. (Analogy: learning that I have beer in the fridge should lead you to raise your credence that I have three beers in the fridge, since that's one way in which I might have beer in the fridge.) But there's no reason to think that we should raise our credence in (Epistemic-Sensitivity) very much. (Just as learning that I have beer in the fridge shouldn't raise your credence that I have three beers in the fridge very much.) Why? Because, as I will now show, there are lots of possible views about the precise way in which an agent's epistemic state constrains her obligations that provide equally good explanations of the intuitions to which Markovits appeals. Thus, the appeal to intuitions doesn't provide much support for (Epistemic-Sensitivity).

To say that a person's epistemic state constrains her obligations invites two questions:

1. First, which aspects of a person's epistemic state constrain her obligations?
2. Second, what is the nature of the constraint that a person's epistemic state imposes on her obligations?

Answering these questions in different ways leads to different versions of the view that a person's epistemic state constrains her obligations. To give the reader a sense of the different options available in answering these questions, I'll consider the different accounts one might give of why the doctor in Markovits's example isn't obligated to administer REALLY-IS-GOOD (assuming, of course, that she isn't—an assumption which I will return to shortly).

The question of which aspects of the doctor's epistemic state constrain her obligations can be divided into two further questions:

- 1a. What *kind* of epistemic state of the doctor's constrains her obligations?
- 1b. And what is the *content* of the relevant epistemic state?

Start with 1a. Philosophers recognize a range of different *kinds* of epistemic states that a person might be in. A person might *know* that p. She might be *permitted* to believe that p. She might be permitted *not* to believe that p. She might *have strong reason* to believe that p, though her reason to believe that p might not be capable of grounding knowledge of p. And so on. It is important, then, to get clear on which kind of epistemic state is supposed to constrain the doctor's obligations. One might maintain, for example, that the reason why the doctor isn't obligated to administer REALLY-IS-GOOD is that in order for the doctor to be so obligated, the doctor must *know* that REALLY-IS-GOOD would cure her patient—which, of course, she doesn't. Alternatively, however, one

might maintain that the reason why the doctor isn't obligated to administer REALLY-IS-GOOD is that in order for the doctor to be so obligated, the doctor must simply have *some reason* to believe that REALLY-IS-GOOD would cure her patient—which, of course, she doesn't. These different explanations suggest different versions of the view that a person's epistemic state constrains her obligations.

Even once we have determined the *kind* of epistemic state of the doctor's that is supposed to constrain her obligations, a further question remains: what is the *content* of the relevant epistemic state? One possible explanation of why the doctor is not obligated to administer REALLY-IS-GOOD is that the doctor does not have good reason to believe *that REALLY-IS-GOOD would cure her patient*. But another possible explanation of why the doctor is not obligated to administer REALLY-IS-GOOD is that the doctor does not have good reason to believe *that she is obligated to administer REALLY-IS-GOOD*. Both explanations appeal to the same *kind* of epistemic state, but they differ in what they identify as the content of the necessary epistemic state.

Answering questions 1a. and 1b. gives us a better sense of precisely *which* aspects of the doctor's overall epistemic state constrain her obligations, but answering these questions does not clarify the precise nature of the *constraint* that the relevant aspects of the doctor's epistemic state imposes. We might wonder, however, whether the relevant constraint says that the *presence* of a certain kind of epistemic state is necessary in order for the doctor to be obligated to administer REALLY-IS-GOOD, or whether the *absence* of a certain kind of epistemic state is necessary instead. Consider, for example, the difference between the following two views:

1. For the doctor to be obligated to administer REALLY-IS-GOOD, she must

have evidence that she is obligated to administer REALLY-IS-GOOD.

2. For the doctor to be obligated to administer REALLY-IS-GOOD, she must lack evidence that she is not obligated to administer REALLY-IS-GOOD.

The first view is an instance of a more general principle according to which obligation requires possession of a certain kind of evidence; the second view, by contrast, is an instance of a more general principle according to which possession of a certain kind of evidence is *incompatible* with obligation. Both views agree, however, with respect to what the relevant *kind* of epistemic state is as well as what the content of that state is.

As I hope this discussion makes clear, questions 1. and 2. can be answered in many different ways, which provides a whole host of different accounts of the ways in which a person's epistemic state constrains her obligations. Consider, for example, the following selection of views:

(Epistemic-Sensitivity) For all actions X, a person S cannot be obligated to do X if S does not believe that she is obligated to do X, and S is epistemically permitted not to believe that.

(Incompatibility-Epistemic-Sensitivity) For all actions, a person S cannot be obligated to do X if S has reason to believe that she is not obligated to do X.

(Non-Moral) For all actions X, if a person S is obligated to do X, then S must be epistemically required to believe that X possesses the non-moral features that ground her obligation to do X, though she need not recognize that this grounding relation obtains.¹³

¹³I will discuss a view along these lines in chapter four. See my discussion there of the "weak epistemic-sensitivity view."

(Position-to-Know-Non-Moral) For all actions X , if a person S is obligated to do X , then S must be in a position to know that her action possesses the non-moral features that ground her obligation to do X , though she need not recognize that this grounding relation obtains.

(Best) For all actions X , if a person S is obligated to do X , then S must be epistemically required to believe that her action would be best for her to perform.

These principles differ from each other in many different ways. (Epistemic-Sensitivity) and (Incompatibility-Epistemic-Sensitivity) differ from each other on the question of what the *constraint* relation is that a person's epistemic state imposes on her obligations. (Non-Moral) and (Position-to-Know-Non-Moral) differ from each other on the question of what *kind* of epistemic state of a person constrains her obligations. And (Non-Moral) and (Best) differ from each other on the question of what the *content* of a person's epistemic state is that constrains her obligations.

Circling back to the appeal to intuitions, the point to take away from the above discussion is that there are a *number* of different explanations of why the doctor isn't obligated to administer REALLY-IS-GOOD. Thus, the appeal to intuitions does little to support (Epistemic-Sensitivity), even if that argument does establish that a person's epistemic state constrains her obligations.

A similar point can be made about Graham's argument. The only place in Graham's argument where he appeals to the idea that a person's epistemic state constrains her obligations is when he states that if a person's epistemic state does not constrain her obligations, then Jackson's doctor is obligated to do A . I think that there is reason to dispute this assumption (as I will explain in the next section), but even if it were accepted, Graham's argument does not rely on any

assumption about the particular way in which an agent's epistemic state constrains her obligations. Thus, even if Graham's argument succeeded, it would do little to support (Epistemic-Sensitivity), as opposed to any of the other views mentioned above.

Evaluating the arguments that a person's epistemic state constrains her obligations

Above, I have pointed out that even if the arguments for the general view that a person's epistemic state constrains her obligations succeed, then those arguments can't be used to support (Epistemic-Sensitivity). I now want to consider, however, whether those arguments *do* succeed in establishing this general view. I think that they do not. I will now explain why they don't.

The first argument that a person's epistemic state constrains her obligations—the "appeal to intuitions"—relies on the assumption that the doctor is not obligated to administer REALLY-IS-GOOD. But this can certainly be disputed. As Graham [13] has pointed out, what is obvious about this case is that the doctor would not be blameworthy for administering LOOKS-GOOD instead of REALLY-IS-GOOD.¹⁴ But that is compatible with saying that by so acting, the doctor would be doing something wrong.

The second argument discussed above suffers from a more serious flaw. The argument assumes that a morally conscientious agent will never choose to perform an action that she knows is morally wrong. But, as Graham [13] notes, a morally conscientious agent might in some cases choose to perform an action that she knows is morally wrong if she has reason to believe performing that action would minimize the chance of her doing something more seriously

¹⁴Graham [13], 93-94.

wrong.¹⁵ If Graham is correct about this, then there is reason to think that Jill could administer drug B *and* be morally conscientious. Even if, by administering drug B, Jill would be choosing to perform an action that she is obligated not to perform, she would also be minimizing the risk of doing something seriously wrong—namely, administering a drug that would kill her patient. This, it seems, is a reasonable compromise for a morally conscientious agent to make, which throws a wrench in the second argument's works.

A further criticism of the second argument is that the argument relies on the claim that if the doctor's epistemic state does not constrain her obligations, then she is obligated to administer drug A. But that does not follow. Even if it is true that the doctor is obligated to administer drug B if her epistemic state *does* constrain her obligations, learning that her epistemic state *does not* constrain her obligations tell us nothing about what she is obligated to do. Indeed, denying that the doctor's epistemic state constrains her obligations is perfectly compatible with saying that she is obligated to administer B; all that is ruled out is saying that the doctor is obligated to administer B *because* her epistemic state constrains her obligations. So it seems that Graham's argument goes off the rails right at the beginning.

2.2.2 Arguments that an agent's epistemic state does not constrain her obligations

We have now seen that the two best arguments for the view that an agent's obligations are constrained by her epistemic state fail. What I want to do now is to present two *prima facie* compelling arguments due to Peter Graham that are intended to show that an agent's epistemic state does *not* constrain her obliga-

¹⁵Graham [13], 99-103.

tions. Remember: if an argument against this more general view succeeds, then (Epistemic-Sensitivity) is dead in the water.

Graham's first argument for the claim that an agent's epistemic state does not constrain her obligations has to do with information acquisition. Graham claims that "when we change our view about our moral obligations as a result of acquiring information, we don't take it that our moral obligations have changed. Rather, it seems to us that we, at last, come to see what our moral obligations were all along."¹⁶ But, the idea goes, if an agent's epistemic state *did* constrain her obligations, we *would* sometimes think that our obligations changed as a result of acquiring information. So, an agent's obligations are not constrained by her epistemic state.

Graham's second argument for the claim that an agent's epistemic state does not constrain her obligations relies on considerations about advice-giving. Graham says, "The question I want answered when I ask myself what my moral obligations are is the same as that which I want answered when, in seeking your help, I ask you what they are; but to adequately answer me you don't need to consider my evidence concerning my situation; therefore, my moral obligations don't depend on my evidence concerning my situation; so, objectivism is true."¹⁷

Evaluating the arguments that an agent's epistemic state does not constrain her obligations

While both of Graham's arguments are compelling at first glance, I believe that they fail. In this section, I'll explain why they do.

Start with the first argument, which says that when we acquire new infor-

¹⁶Graham [13], 91.

¹⁷Ibid.

mation about what we are obligated to do, we don't think that our obligations change. To see this argument's limitations, think back to the case in which the doctor faces a choice between administering LOOKS-GOOD and REALLY-IS-GOOD. Suppose that *you* are the doctor, and that you have no idea that your evidence is misleading. (Recall that your evidence suggests that LOOKS-GOOD will cure your patient, when in fact it would kill him, and that REALLY-IS-GOOD would kill him, when in fact it would cure him.) Suppose, however, that you come to acquire more information and thus to realize the truth about LOOKS-GOOD and REALLY-IS-GOOD. Is it really as obvious as Graham assumes that you wouldn't believe that your moral obligations have changed? When imagining myself in the doctor's situation, I find it conceivable that I *was* obligated to administer LOOKS-GOOD (even though that would have disastrous consequences), but, upon acquiring more information, I am now obligated to administer REALLY-IS-GOOD. At the very least, I do not think that this possibility is self-evidently wrong. Anyone else who feels the same way will not be convinced by Graham's argument.

To rephrase the point slightly, the problem with Graham's first argument is the same sort of problem that faced Markovits's appeal to intuitions. As Graham points out, the problem with Markovits's argument is that it is not obvious that the doctor can permissibly refrain from administering REALLY-IS-GOOD. But, similarly, it is not obvious that the doctor is obligated to administer REALLY-IS-GOOD. It's simply not obvious what the doctor is permitted or obligated to do. Thus, both arguments suffer from the same defect: they rely on confident judgments that we have no reason to be confident of, in the absence of independent support.

The problem with Graham's second argument is more subtle. The flaw in

this argument is brought out best, I believe, by considering the following parody argument against the claim that an agent's *epistemic obligations* don't depend on her epistemic state:

The question I want answered when I ask myself what to believe is the same as that which I want answered when, in seeking your help, I ask you what I should believe; but to adequately answer me you don't need to consider my evidence concerning my situation; therefore, my epistemic obligations don't depend on my evidence concerning my situation.

This argument, I believe, sounds just as plausible, at first glance, as Graham's. But the argument *must* go wrong, since the conclusion it aims for is patently false: an agent's epistemic obligations obviously *do* depend on her epistemic state.

Seeing that the argument doesn't work, however, isn't the same thing as understanding why it goes wrong. It would be helpful to understand why Graham's argument seems so convincing at first glance.

I think that we can give a partial explanation of why Graham's argument seems convincing at first glance if we focus again on asking for advice about what to believe. If I ask a physicist what I have most reason to believe concerning quantum mechanics, I'm doing so because I want to believe true things about the physical universe, and the physicist is in a position to help me. My primary concern is *not* to form beliefs responsive to the evidence I had *before* speaking to the physicist. What I want is to gain *better* evidence and to believe what *that* evidence suggests I should believe. Similarly, when I ask someone what I am obligated to *do*, my goal is to gain better evidence about what I am obligated to do and then to act accordingly. My primary concern is *not* to act on

the basis of the evidence I had before.

Requests for advice about what we have most reason to believe or what we are obligated to do, then, are also requests for more evidence. Graham's argument seems compelling when we neglect this fact. Once it is recognized, however, we see that it is a simple matter to account for why moral advisors need not inquire into their advisees' moral evidence, and we see that Graham's argument does not threaten the view that an agent's epistemic state constrains her obligations.

2.2.3 Summary

We have now seen that the arguments in favor of the view that an agent's epistemic state constrains her obligations cannot be used to support (Epistemic-Sensitivity). We have also seen that the arguments *against* that view fail, and so those arguments pose no threat to (Epistemic-Sensitivity). We are left, then, to wonder: what else can be said either in favor of or against (Epistemic-Sensitivity)?

In the next section, I will present a novel argument against (Epistemic-Sensitivity). To motivate my argument, I will begin by considering a few simple-minded moves one might use to try to refute (Epistemic-Sensitivity). I will show that these moves need to be supplemented if they are to work. I will then argue that they can be supplemented in the right way.

2.3 A new argument against (Epistemic-Sensitivity)

(Epistemic-Sensitivity) says that *any* agent who permissibly does not believe that she is obligated to do X is not obligated to do X. To refute this claim, all one

needs is to identify an agent S and an action X such that S permissibly lacks a belief that she is obligated to do X, even though S really *is* obligated to do X. To try to come up with such an agent/action pair, the natural approach is to do the following:

First, take someone who is obligated to perform a certain action X, and who knows that she is obligated to do X,

Second, subtract her evidence that she is obligated to do X,

And, third, subtract her belief that she is obligated to do X.

If the agent's obligation to do X persists even after her evidence and belief that she is obligated to do X have been subtracted, then a counterexample to (Epistemic-Sensitivity) has been obtained.

The problem is how to identify an agent S and an obligation X such that the agent's obligation *does* persist through this subtraction.

Consider, for example, an agent who is obligated to go to her office hours because she promised a student that she would be there. Once we subtract the agent's evidence and belief that she is obligated to keep her office hours, why shouldn't one believe that her obligation also disappears? I, for one, don't find it obvious that the agent's obligation would persist.

However, it may seem that the problem in this case is that the agent's obligation doesn't concern a matter of *serious* moral importance. (Blowing off one's office hours is bad, but, in the grand scheme of things, it is really of minor significance.) Perhaps, however, if we consider an agent who is under a more serious obligation, it will be more plausible that the agent's obligation remains even after subtracting the agent's belief and evidence that she is under that obligation.

Imagine, for example, taking an agent who is obligated not to torture animals for fun, and who believes that she is so obligated, and then subtracting

both her belief and her evidence that she is so obligated. It is less plausible that such an agent's obligation disappears after performing this subtraction than it was to say that the agent in the previous case's obligation disappeared after subtracting her belief and evidence about her obligation.

The problem with this sort of case, however, is that there is reason to doubt whether one's judgments about it are reliable. The problem is that when we consider the obligation not to torture animals for fun, we not only increase the moral stakes, but we also increase dramatically the amount of evidence that the agent has that she is under that obligation. This introduces a confounding factor: it makes it hard to imagine how an agent really *could* lack this kind of evidence without suffering from a defect the presence of which would explain all by itself why she is not under that obligation.

To understand this point, consider what's involved in subtracting bit by bit an agent's evidence that torturing animals for fun is wrong.

One way to subtract an agent's evidence that torturing animals for fun is wrong would be to start by subtracting the agent's evidence of the non-moral effects of torturing animals for fun, which are surely *part* of her overall evidence that torturing animals for fun is wrong. So, we can imagine removing the agent's evidence that torturing animals causes those animals pain, etc. But it's hard to imagine how someone could lack this evidence without being so cognitively impoverished that she might not even count as a moral agent, and so wouldn't be subject to any obligations at all. At the very least, we should be extremely wary of relying on judgments about agents who have undergone such a subtraction.

There is, however, another way of trying to subtract the agent's evidence that torturing animals for fun is wrong. One could instead proceed by subtract-

ing *only* the agent's *moral* evidence that she is obligated not to torture animals for fun, while leaving fixed her evidence of the non-moral effects of torturing animals. The problem now, however, is that it is still not obvious how an agent could find herself in a position where she lacks this moral evidence while maintaining the corresponding non-moral evidence. One might imagine, for example, that the agent is a moral imbecile, despite having perfectly adequate *general* cognitive abilities. When she tries to reason about morality, however, she goes horribly wrong, or simply can't make heads or tails of much. The problem with this is that it is very hard to imagine how this could be possible. To ask one to imagine this is a bit like asking someone to imagine an intelligent person who can, for example, think in quite a sophisticated way about literature, or car repair, but simply cannot recognize even the simplest mathematical truths, such as $2 + 2 = 4$. It's difficult to wrap one's mind around such a scenario, which means that it's quite difficult to make confident judgments about what such an agent would or would not be obligated to do.

What one needs, then, if one is to pursue the strategy I have outlined for arguing against (Epistemic-Sensitivity) is reason to think that it is possible for an agent who does *not* suffer from a general cognitive deficit to permissibly lack a belief that she is obligated not to torture animals for fun. The hope, then, is that, given the moral seriousness of the obligation not to torture animals for fun, it would be plausible that the agent's obligation remains even once her belief and evidence that she is so obligated are subtracted. Of course, whether it will be plausible that the agent's obligation remains might depend on what she must be like in order for her to permissibly lack such a belief, so getting the details of the epistemological story right is crucial.

I believe that this strategy can be carried out. The following argument out-

lines how it can:

1. An agent can (epistemically) permissibly fail to believe that torturing animals for fun is wrong by satisfying a fairly minimal set of epistemic conditions, which do not require that the agent suffer from a general cognitive deficit.
2. If (Epistemic-Sensitivity) were true, then it would follow that any agent who (epistemically) permissibly fails to believe that torturing animals for fun is wrong would thereby not be subject to an obligation to refrain from torturing animals.
3. Thus, if (Epistemic Sensitivity) were true, then anyone satisfying a fairly minimal set of epistemic conditions would not be obligated to refrain from torturing animals.
4. But it's not true that someone who satisfies a relatively minimal set of epistemic conditions is thereby freed from their obligation to refrain from torturing animals.
5. Therefore, (Epistemic-Sensitivity) is false.

The next chapter is devoted to laying out a minimal set of conditions that are sufficient for an agent to be epistemically permitted to suspend belief in *all* claims of the form *Xing is obligatory*—which, of course, includes the claim that *refraining from torturing animals for fun is obligatory*.

The conditions that I claim are sufficient for an agent to be epistemically permitted to suspend belief in all claims of the form *Xing is obligatory* are the following: (a) the agent lacks any “substantive” moral beliefs—that is, she does not have any moral beliefs other than those that are logically entailed by her non-moral beliefs, (b) she either lacks moral experiences or justifiably believes

that she doesn't have any, (c) she either lacks moral intuitions or has anti-moral intuitions, and (d) if she originally became a moral skeptic because she was persuaded by a bad skeptical argument, she has since reassessed her reasons for being a moral skeptic and decided to remain a skeptic simply because there is no reason to abandon her skepticism. I will explain in the next chapter why these conditions are sufficient for an agent to be epistemically permitted to suspend belief in all claims of the form *Xing is obligatory*. How that argument goes is not important for our purposes at the moment, however. All that matters is whether an agent who satisfies these conditions would thereby come to lack any moral obligations at all. I believe that they would not. (Epistemic-Sensitivity), however, entails that they would. Thus, I conclude that (Epistemic-Sensitivity) is false.

It would be possible for a die-hard defender of (Epistemic-Sensitivity) to hold on to that view by biting the bullet and accepting that an agent who satisfies conditions (a)-(d) would thereby remove all of her obligations. However, given the choice between accepting that conclusion or rejecting (Epistemic-Sensitivity), it seems to me that the best option is to reject (Epistemic-Sensitivity).

If one accepts the argument against (Epistemic-Sensitivity) that I have just spelled out, then the threat that (Epistemic-Sensitivity) poses to the possibility of NCFPO ignorance dissolves. But refuting (Epistemic-Sensitivity) does not, by itself, show that NCFPO ignorance is possible. Perhaps there is another argument against the possibility of NCFPO ignorance that succeeds. So an additional argument needs to be offered in order to show that NCFPO ignorance is possible.

Luckily, it is quite easy to modify the argument above so as to show that

NCFPO ignorance is possible. With the appropriate modifications, the argument for the possibility of NCFPO ignorance goes as follows:

1. An agent can (epistemically) permissibly fail to believe that she is under any obligations by satisfying a fairly minimal set of epistemic conditions—conditions (a)-(d), outlined above.
2. An agent who satisfies conditions (a)-(d) and does not believe any claim of the form *Xing is obligatory* would, nonetheless, still be subject to obligations.
3. However, she would be non-culpably ignorant of those obligations.
4. Therefore, NCFPO ignorance is possible.

I have already explained that premise 1. will be defended in the next chapter, so no further remarks on that premise are necessary. Premise 2. is simply a summary of the main idea behind the previous argument—that satisfying conditions (a)-(d) is not enough to remove all of an agent's obligations—so that premise does not need further explanation either. Finally, premise 3. follows from the idea, which I take to be relatively obvious, that any agent who (epistemically) permissibly fails to believe a true claim P counts as non-culpably ignorant of P.

The same general idea, then, suffices to establish both that (Epistemic-Sensitivity) is false, and that NCFPO ignorance is possible. Of course, the key premise behind both arguments is that an agent can (epistemically) permissibly fail to believe all claims of the form *Xing is obligatory* by satisfying a fairly minimal set of epistemic conditions. Defending this claim will be the work of the next chapter.

CHAPTER 3
CAN SKEPTICISM BE REASONABLE?

3.1 Introduction

It is increasingly common for epistemologists to draw a distinction between two types of responses to external world skepticism: in James Pryor's terminology, the distinction is between *ambitious anti-skeptical responses* on the one hand and *modest anti-skeptical responses* on the other.¹ Ambitious anti-skeptical responses attempt to provide arguments against skepticism that appeal only to premises that even a skeptic would accept. These anti-skeptical responses try to answer the skeptic "on his own terms."² Modest anti-skeptical responses, by contrast, aim only to "diagnose and defuse" skeptical arguments, even if doing so requires us to appeal to premises that the skeptic will not accept.³

Many epistemologists believe that there is little hope for ambitious anti-skeptical responses; however, some of these same epistemologists maintain that the search for modest anti-skeptical responses has yielded results that should set the minds of anti-skeptics at ease.⁴ Thus, discussions of skepticism often take on a defensive character: they don't aim to provide arguments that show that skepticism is always unreasonable; instead, they aim only to show that we non-skeptics are permitted to hold on to common sense beliefs.

This is understandable. If it could be shown that our common sense beliefs survive the skeptic's attacks, then one might think that should be enough to let us sleep at night. None of us want to be skeptics, so it might seem that once we have established the permissibility of our point of view, the interesting work is

¹Pryor [32], 517-518.

²Pryor [32], 517.

³Ibid.

⁴Byrne [5], 300-302; Pryor [32], 517; Williamson [46], 26-27.

done.

But because discussions of skepticism typically aim to show only that our beliefs can survive the skeptical onslaught, they leave a very interesting question unanswered. The question is this: even if we friends of common sense are epistemically permitted to keep our common sense beliefs, might it not also be epistemically permissible—or even epistemically mandatory—for some people to be skeptics?

I don't know of any extended discussions of this question, though a number of epistemologists have at least gestured towards answers to it. To give just one example, here is Williamson:

If someone has already taken the route into scepticism offered by that fallacious argument, before it was blocked, and has become genuinely undecided, at least in principle, as to whether she is in a sceptical scenario, then the blocking of the route now comes too late to rescue her. Nothing said here should convince someone who has given up ordinary beliefs that they did in fact constitute knowledge, for nothing said here should convince her that they are true. The trick is never to give them up. This is the usual case with philosophical treatments of scepticism: they are better at prevention than at cure. If a refutation of scepticism is supposed to reason one out of the hole, then scepticism is irrefutable. The most to be hoped for is something which will prevent the sceptic (who may be oneself) from reasoning one into the hole in the first place.⁵

Williamson never argues for the views he states here. Nor do any of the epistemologists I have seen who have expressed similar sentiments. It is natural

⁵Williamson [46], 26-27.

to wonder, then, what sort of argument could be given for this view. This is one issue that I will be taking up in this chapter. But before I turn to that issue, I want to introduce a bit of Williamson-inspired terminology that will help to simplify the following discussion.

Let a *cure* (*in the narrow sense*) for a subject *S*'s skepticism be an argument with the following property: were *S* to be exposed to the argument, *S* would be epistemically required to abandon her skepticism. (Later in the chapter, I will introduce the notion of a cure in the broad sense for a subject *S*'s skepticism. But until I get there, whenever I use the word "cure," I mean a cure in the narrow sense.)

Let a *protective against skepticism*, on the other hand, be a demonstration that some argument for skepticism fails.

Cures, as the name suggests, are tools with the power to destroy the disease of skepticism whenever they come into contact with it. We might think of them as offensive weapons in the fight against skepticism. Protectives, by contrast, are defensive weapons. In order to count as a protective, an argument need only be able to arrest the spread of skepticism; it needn't be able to cure cases of skepticism once they are established.

One view that Williamson expresses in the above passage is that all cases of skepticism are incurable. (Or, to put his view another way, we might say that he thinks that all cases of skepticism are *terminal*.) This, I think, is the most commonly expressed view about the epistemic status of the skeptic's position.⁶ But it is important to note that this view stands at one extreme of a spectrum of possible views, which the notion of a cure helps us to see. At the other end of the spectrum is the view that no cases of skepticism are terminal—in other

⁶See, for example, Pryor [32], 517 and Byrne [5], 319 and f.n. 7, who suggest that the prospects of curing skepticism seem "somewhat dim".

words, all cases of skepticism are curable. According to this view, there is some argument (or maybe set of arguments) that could be wheeled out to deal with any case of skepticism whatsoever.⁷ But there are also intermediate views. One might hold that there are a few particularly troublesome cases of skepticism which are terminal, though most cases are curable. Or one might hold that most cases of skepticism are terminal, and the curable cases are only a small subset.

Though epistemologists do occasionally state one of these views, to the best of my knowledge, no one has ever made a serious effort to sort out which is correct. This is unfortunate, since it really is an interesting question which of these views is right. But, as I have already explained, epistemologists' neglect of this question is understandable, even if it is ill-motivated.

What is less understandable is why so little attention has been paid to the analogous question about moral skepticism: even if it is the case that it is reasonable for many people to believe that they are under obligations of various kinds, might it not be reasonable (or even epistemically mandatory) for other people to suspend belief in all moral claims? This question is important not only for its intrinsic theoretical interest, but also because it has implications for moral practice.

The implications for moral practice arise if one assumes, as many philosophers do, that NCFPO ignorance constitutes an excuse for wrongdoing.⁸ According to this view, an agent is not blameworthy for acting wrongly if, at the time she acted, she was non-culpably ignorant of the fact that it would be wrong for her to act in that way. To make things more concrete, consider the example of buying "happy meat"—where "happy meat" is meat from animals who lived

⁷Note that even if all cases of skepticism are curable, there might not be a *single* cure capable of treating every case—just as there is no single antibiotic that can clear up every treatable bacterial infection. Thus, the qualification "or set of arguments" really is necessary here.

⁸See chapter four for an extended argument that NCFPO ignorance constitutes at least a partial excuse for wrongdoing.

happy, healthy lives on a free-range farm, until they were slaughtered in a way that caused them as little pain as possible. Many people think that it is morally permissible to buy happy meat. But suppose that these people are mistaken. Even if it is wrong to buy happy meat, it is conceivable that some people who think that buying happy meat is permissible are nonetheless justified in believing that. (The murkier that you think this moral issue is, the more likely you'll be willing to admit that some people could be justified in holding a mistaken view about it.) Indeed, it is conceivable that some people might be justified in believing that buying happy meat is permissible even if they know as many of the non-moral facts about what is involved in production of happy meat as anyone does. Assuming that this is the case, what should we say about such people when they purchase happy meat? Do they deserve blame? Many philosophers would say that they do not, and the reason why is that these philosophers maintain that NCFPO ignorance constitutes an excuse for wrongdoing.

If NCFPO ignorance constitutes an excuse for wrongdoing, then the view that some people can permissibly suspend belief in all moral claims has striking implications for our moral practice of blaming. If some people can permissibly suspend belief in all moral claims, then such people will suffer from the most extreme kind of NCFPO ignorance possible—they will be non-culpably ignorant of all claims of the form *I am obligated to do/not do X*. Such people, therefore, will not deserve blame for *anything* they do—no matter how horrific. It is very important, then, to settle whether reasonable moral skeptics are possible. A good first step towards answering this question is to explore whether moral skepticism is terminal. (As we shall see later, showing that some cases of moral skepticism are terminal would not show that reasonable moral skeptics are possible, but it would get us most of the way there.)

With that said, here is my plan in this chapter. There is a simple argument that one might offer for the conclusion that every case of moral skepticism is terminal, which I will explain in section 2. This argument has a natural analogue that I suspect is what has convinced people such as Williamson to believe that every case of external world skepticism is terminal. As I will point out, however, both of these arguments run into trouble. In both cases, the arguments rely on a claim to the effect that a certain sort of “inference barrier” exists. For the argument that every case of moral skepticism is terminal, the inference barrier in question is the familiar “is-ought gap.” For the argument that every case of external world skepticism is terminal, the inference barrier is supposed to prevent us from reasonably inferring claims about the external world from claims that are not about the external world. The problem is that it is very difficult to spell out these kinds of inference barriers. But there are suggestions in the literature on the is-ought problem that look like they might help here. These suggestions seem as if they can help us to formulate a defensible “is-ought gap” claim. Furthermore, they seem as if they will help us to formulate a claim to the effect that there is an “external world gap,” which is precisely what people like Williamson need in order to support their view that every case of external world skepticism is terminal.

But in the next section (section 3), I’ll introduce an apparent counterexample to is-ought gap theses that looks like trouble for the argument that all cases of moral skepticism are terminal. (A similar example might also cause trouble for the claim that there is an “external world gap”, but I won’t pursue this point. For the rest of the paper, my focus will be restricted to the question of whether moral skepticism can be reasonable.) Despite the appearance of trouble here, I’ll suggest that there is reason to think that this counterexample can be dis-

solved. But even if it can't, I'll show that there is a way to repair the argument by weakening its conclusion; the weakened argument, instead of showing that all cases of moral skepticism are terminal, will show only that moral skepticism is "treatment resistant." (I'll explain later what that means.)

However, even this weakened argument faces trouble from the moral analogues of various views one finds in epistemology. In the next section of the paper, I'll suggest that these troubles can be avoided, and that we can show that some (or perhaps many) cases of moral skepticism are likely to be uncured, even if we grant a number of assumptions that appear at first glance to be hostile to the view that moral skepticism is never terminal.

Finally, I will discuss how my argument that some cases of moral skepticism are likely to be uncured can be used to show not only that some moral skeptics are reasonable, but also that moral skepticism can be epistemically mandatory for some people.

3.2 Autonomy claims

There is a tempting argument for the conclusion that all cases of moral skepticism are terminal. The argument goes as follows: A moral skeptic is someone who has given up all of her moral beliefs.⁹ But any argument that concludes with a moral claim must rely on some moral premises, so the skeptic isn't in a position to be epistemically compelled to accept any such argument. Thus, no moral skeptic can be cured; that is, all cases of moral skepticism are terminal.

⁹On some understandings of moral skepticism, a person must do more than give up all of her moral beliefs in order to count as a moral skeptic—perhaps she also needs to believe that moral knowledge is impossible or that some other equally bold metaethical thesis is true. But the understanding of moral skepticism at play in the argument above is quite different. On that understanding, moral skeptics are those who are maximally noncommittal about first-order moral claims. Moral skeptics of this kind, however, needn't possess any particular theoretical view about morality or the possibility of moral knowledge.

Laying out the premises of this argument a bit more precisely, the argument is the following:

(Is-Ought Gap) One can't reasonably infer any moral claim from a set of non-moral premises.

(Inference Assumption) If a moral skeptic is epistemically required to abandon her skepticism, that's because she is required to infer a moral claim from some other belief(s) that she has.

(The Strong Conclusion) Therefore, no moral skeptic is epistemically required to abandon her skepticism.

An analogous argument might be offered for the conclusion that external world skepticism is terminal. Again, the idea is simple. An external world skeptic is someone who has given up all of her external world beliefs. But any argument that concludes with an external world claim must rely on some external world claims, so the external world skeptic isn't in a position to be epistemically compelled to accept any such argument. Laying out the argument a bit more precisely, it goes as follows:

(External World Gap) No external world claim can reasonably be inferred from a set of premises that doesn't include some external world claims.

(Inference Assumption*) If an external world skeptic is epistemically required to abandon her skepticism, that's because she is required to infer an external world claim from some other belief(s) she has.

(The Strong Conclusion*) Therefore, no external world skeptic is epistemically required to abandon her skepticism.

I suspect that something like this argument is what has led epistemologists such as Williamson to accept (The Strong Conclusion*).

Both of these arguments are compelling at first glance. The logic is sound, and the premises are easily motivated.

Take (Is-Ought Gap), for example. Support for the claim that there is an is-ought gap is provided by the fact that arguments that move from non-moral claims to moral claims always seem to rely on “bridge principles” that spell out some sort of connection between the moral and the non-moral. As Nicholas Sturgeon has noted, the fact that such bridge principles are required in ethics isn’t surprising.¹⁰ There are many inferential gaps that can be bridged only via “bridge principles.” To give one example, claims about the future are inferentially isolated in just this way from claims about the past; in order to argue from some claims about the past to a claim about the future, you need to assume something about the connection between the past and the future—for example, that nature is suitably uniform. Similarly, in order to bridge the gap from claims about the “internal world” to claims about the “external world,” it seems that we must rely on an assumption about the connection between the two. This is just what (External World Gap) says.

What about the (Inference Assumption)? This premise is supported by a widely accepted epistemological view: coherentism. Coherentism says that beliefs can be supported only by other beliefs. But if beliefs can be supported only by other beliefs, then if you’re required to form a particular belief, that must be because you have some other beliefs that you’re required to infer that belief from. The (Inference Assumption) and the (Inference Assumption*) are simply special cases of this principle.

So, if we think that coherentism is right, and we think that there is an is-ought gap and an “external world gap,” then it seems that we have reason to believe that all cases of moral skepticism and external world skepticism are ter-

¹⁰Sturgeon [41], 200-202

minal.

In reality, however, things are not this simple. There are a number of problems with both of these arguments.

To see the first problem, notice that in presenting the above arguments, I made an assumption about what moral skeptics and external world skeptics are: I assumed they are people who have suspended belief in all moral claims and external world claims, respectively. So, in order to really understand what it is to be a moral skeptic or an external world skeptic, we need to know what counts as a moral claim and what counts as an external world claim.

How do we identify the moral and external world claims? Well, one might think this is simple. A moral claim is any claim that contains a moral predicate in it. And an external world claim is any claim that contains an external world predicate in it. So long as we can classify predicates into the moral/non-moral or the external world/non-external world, it will also be possible to sort claims into these categories.

The problem is that if we define our terms in this way, there are lots of ways to reasonably infer moral claims from non-moral claims and external world claims from non-external world claims, which means that (Is-Ought Gap) and (External World Gap) are false. To give just one set of examples (inspired by examples of Arthur Prior's): from *grass is green*, we can infer *grass is green or murder is wrong*.¹¹ And from *I am thinking*, we can infer *I am thinking or tables exist independently of the mind*. In the first case, we have an inference from a non-moral claim to a claim that counts as moral (because it contains a moral predicate). And in the second case, we have an inference from a non-external world claim to an external world claim. So we have counterexamples to both (Is-Ought Gap) and (External World Gap).

¹¹Prior [31].

That there is a problem here should not be surprising. We've known for a very long time now that it is not easy to say in what sense there is an Is-Ought Gap.¹² What is surprising is that this problem also causes trouble for the obvious way of defending the view that external world skepticism is terminal, yet epistemologists seem not to have noticed this.

Now, it might seem that this problem isn't so hard to fix. One might be tempted to say that the conclusion of the first argument isn't "purely moral," and the conclusion of the second argument isn't "purely external world." Maybe all we need to do, then, is to say that there is no way to reasonably infer a purely moral or purely external world claim from premises that are non-moral or non-external world. But what does it mean for a claim to be purely moral—that it contains only moral predicates? But then what about the following sentence: *For all x, if x is a murder, then x is wrong?*¹³ This is a claim that contains a non-moral predicate ("is a murder"), but which paraphrases a claim that should count as purely moral if any claim does: namely, *murder is wrong.*¹⁴ A similar problem crops up for saying what it is for a claim to be "purely external world." Consider the sentence: *For all x, if x is Obama, then there is a material being y such that x = y.* The sentence this paraphrases seems like it should count as purely external world if any sentence does: *Obama is a material being.*

Thus, the problem is actually quite tricky. But as I said, this is a problem that ethicists and logicians have spent a great deal of time thinking about. And they've come up with some results. One strategy (defended by Charles Pigden) is to note that all standard counterexamples to (Is-Ought Gap) have the

¹²A point discussed at length in Jackson [17], Pigden [28], Prior [31].

¹³I draw here on Brown [4], sect. 3.

¹⁴If one thinks that "is a murder" is a moral predicate, then we can simply substitute a different example. *Taking candy from babies is wrong* is a purely moral claim if any is, but *For all x, if x involves taking candy from babies, then x is wrong* will not count as a purely moral claim on the view we are currently considering.

following property: the moral predicates that occur in the conclusions of the counterexamples could be replaced by any other predicate, and the inference would be just as good.¹⁵ Thus, instead of arguing *grass is green; therefore, grass is green or murder is wrong*, we could argue that *grass is green; therefore, grass is green or murder is obligatory*. Pigden says that in inferences like this, the moral predicate in the conclusion occurs “vacuously.” What Pigden has shown is that in first-order logic “a predicate . . . cannot occur non-vacuously in the conclusion of a valid inference unless it appears among the premises.”¹⁶

Other philosophers have proven similar results in first-order logic. Campbell Brown, for example, has a recent paper that proves several results along these lines.¹⁷ Knowing about this work, one might think that results like these will allow us to avoid all counterexamples to the claim that ethics is autonomous (which is just another way of saying that there is an is-ought gap). And since these results are quite general, the epistemologist could also use these results to defend the existence of an external world gap.

If that’s right, then we could say that although someone who has suspended belief in all moral claims can be epistemically compelled to adopt some moral claims, it won’t be possible to convince them via a deductive argument formalizable in first-order logic to believe any “substantive” moral claims. Thus, we might try to modify the arguments given above in the following way.

First, we need to redefine what it means to be a skeptic. A skeptic can’t just be someone who has suspended belief in all moral or external world claims, since Prior-like counterexamples show that that type of skepticism is curable. Instead, we need to say that a moral skeptic is someone who suspends belief in all moral claims except those that follow as a matter of logic from her non-moral

¹⁵Pigden [28].

¹⁶Pigden [28], 136.

¹⁷Brown [4].

beliefs.¹⁸

Then, once we have redefined skepticism in that way, we can argue as follows:

(Fancy Is-Ought Gap) One can't reasonably infer any *substantive* moral conclusions from non-moral premises.

(Fancy Inference Assumption) If a moral skeptic is epistemically required to abandon her skepticism, that's because she is required to infer a *substantive* moral claim from some other belief(s) she has.

(Strong Conclusion) Therefore, no moral skeptic is epistemically required to abandon her skepticism.

To summarize: it appears that we can use the work of Pigden and others to defend our argument that all cases of moral skepticism are terminal from Prior-like counterexamples. However, if we have to make these moves in the moral case, we will also have to make them in the epistemological case, too. So, we will need to redefine external world skepticism as suspending belief in all external world claims except those that follow as a matter of logic from one's non-external world beliefs. Then we will need to run an argument analogous to the above. Thus, if Williamson were to take the obvious approach to defending the (Strong Conclusion), he would need to take a page out of the ethicist's book.

3.3 Counterexamples to autonomy

At this point, it might seem that we have answered the questions with which we began. All cases of moral skepticism and external world skepticism (in the

¹⁸On this understanding, a moral skeptic is no longer someone who is maximally noncommittal about first-order moral claims; instead, a moral skeptic is someone who is maximally noncommittal about *substantive* first-order moral claims.

senses we've just defined) are terminal—end of story. The ethicists and logicians have come to the rescue by showing us how to formulate the (Fancy Is-Ought Gap) claim, which saves the day.

Unfortunately, there is still a problem that remains to be dealt with. (Actually, there are several problems that remain to be dealt with, but I will take them one at a time.) The problem is that there is reason to believe that is-ought gap theses like the one defended by Pigden aren't strong enough to show that (Fancy Is-Ought Gap) is true.

Even if it is the case that in first-order logic, there is no way to derive a substantive moral claim from non-moral premises, there is an argument that is not formalizable in first-order logic, which seems to derive a substantive moral conclusion from non-moral premises. The argument in question relies on a modified version of an example due to Danny Mond.¹⁹ Suppose that Sam has a trusted friend Alfie, who has a fantastic track record with respect to all sorts of different intellectual pursuits. Whenever Sam consults Alfie about some topic he doesn't know much about, and which he later investigates, he always finds that Alfie's views were correct. One day, Sam, who is a moral skeptic, consults Alfie about morality. He asks, "is it really true that giving to charity is morally obligatory?" Alfie nods.

Is Sam epistemically required to believe Alfie? It might seem that he is. But note that Sam doesn't seem to be relying on any moral premises when he argues as follows: (1) Whenever Alfie says that p (and p isn't obviously crazy), p is true. (2) Alfie says that giving to charity is morally obligatory, and this claim isn't obviously crazy. Therefore, (3) giving to charity is morally obligatory.

What is happening here? Essentially, what is happening is that Pigden's results are not strong enough to rule out this type of counterexample. The reason

¹⁹This example was popularized by Karmo [20], 253.

why is that his results only rule out the existence of is-ought gap crossing arguments that are formalizable in first-order logic. But one of the premises of the above argument—namely, the claim that *Whenever Alfie says that p (and p isn't obviously crazy), p is true*—is not formalizable in first-order logic.

While the Alfie case does appear to pose a problem for (Fancy Is-Ought Gap), there is reason to think that appearances might be misleading here. The reason why is that, although Sam's argument to the claim that giving to charity is morally obligatory doesn't seem to rely on any moral premises, it's arguable that Sam must rely on moral premises in order to be *justified* in believing those premises. Suppose that Sam knows (somehow) that Alfie is an expert (or is even infallible) when it comes to all sorts of topics other than morality; the idea behind the Alfie example, then, is that Alfie's expertise on those topics generates an epistemic requirement for Sam to trust Alfie when it comes to moral matters. But it's not obvious that Sam is epistemically required to trust Sam on moral matters unless he has reason to believe that Alfie's expertise extends to morality; it's hard to see, though, how Sam could have reason to believe that Alfie's expertise extends to morality if Sam does not have any moral beliefs at all.

There's room to dispute, then, whether the Alfie case falsifies (Fancy Is-Ought Gap). But even if it does, it's possible to formulate an is-ought gap claim that allows us to obtain a conclusion that is somewhat similar to, though weaker than, (Strong Conclusion). To see how this would go, notice that there is an important similarity between the Alfie example and Prior's counterexamples to simple formulations of an is-ought gap claim: both feel like tricks. The reason why the Alfie case, in particular, feels like a trick is clear; it doesn't feel like a deep counterexample because it relies on Sam and Alfie finding themselves in a very unusual set of circumstances, which none of us are likely to ever find

ourselves in. Thus, it seems that even if the gap between is and ought can be bridged, it can be bridged only in certain very special cases. This tells us something important. What it tells us is that even if moral skepticism can be cured by putting skeptics in Alfie-like situations, it is really hard to do that. We can't do it at will. So, even if moral skepticism isn't terminal (because it is in principle curable), then it is still the case that moral skepticism might be *treatment resistant*—that is, it might be that it can be cured only by cures that are not always available.

What we can do, then, is to tweak our earlier argument for the conclusion that moral skepticism is terminal in order to get the following argument for the claim that moral skepticism is treatment resistant:

1. **(Weak Is-Ought Gap)** Only people in unusual situations can be epistemically required to infer a substantive moral claim from non-moral premises.
2. **(Fancy Inference Assumption)** If a moral skeptic is epistemically required to abandon her skepticism, that is because she is required to infer a substantive moral claim from some other belief(s).
3. **(Weaker Conclusion)** Therefore, no moral skeptics, except some in very unusual situations, are epistemically required to abandon their skepticism.

An interesting question to ask here is whether a similar tweak needs to be made to the argument that external world skepticism is terminal. I don't have the space to discuss this question here, though I do think that a similar tweak is probably necessary. Instead, in the remainder of this paper, I will focus primarily on the issues that arise for the argument about moral skepticism.

3.4 The (Fancy Inference Assumption)

As I have already hinted, even this last revision of our argument about moral skepticism does not get rid of every problem that it faces. A remaining problem is that some philosophers will take issue with the (Fancy Inference Assumption).

The easiest way to see why someone might reject this assumption is to consider the analogous claim about external world skepticism: If an external world skeptic is epistemically required to abandon her skepticism, that is because she is required to infer a substantive external world claim from some belief or beliefs she has.

Why would some reject this claim? One reason is that some believe that an agent's experiences can, by themselves, epistemically require her to adopt an external world belief.

This type of view is a natural extension of so-called "dogmatist" views in epistemology.²⁰ Dogmatists say that, in the absence of defeaters, you are justified in taking your experiences at face-value, even if you have no antecedent justification for believing that your experiences are reliable. (So, you are justified in taking your experience as of a table at face-value if you have no reason to think that your experiences are unreliable. But if you have reason to think, for example, that you are hallucinating, then you are not justified in taking that experience at face-value.) Dogmatism as it is usually understood is a view about what it is epistemically *permissible* for us to do. But we can easily extend the view so that it becomes a view about what we are epistemically required to do, by modifying it to say that you are epistemically *required* to take your experiences at face-value, even if you have no antecedent justification for believing

²⁰See, among others, Pryor [32].

that your experiences are reliable, so long as those experiences don't conflict with anything you believe.

Notice that if this view is correct, then it might be that some (if not all) external world skeptics are epistemically required to abandon their skepticism, even if that requirement was not generated by an argument. This necessitates a change to our definition of what it is to be a cure for skepticism. Remember that I said at the start of the paper that a cure (in the narrow sense) for *X*'s skepticism is an argument with the power to epistemically require *X* to abandon her skepticism. On that definition, if *X* is epistemically required to abandon her skepticism by an experience and not an argument, her skepticism counts as terminal (in the narrow sense), even though she is not epistemically permitted to remain a skeptic. This is an awkward consequence, which should be fixed.

Fixing this is easy. All we need to do is to introduce a broader notion of a cure: say that a *cure (in the broad sense) for X's skepticism* is an argument, experience, or other item that can epistemically require *X* to abandon her skepticism. Experiences of the right kind, then, might count as cures *in the broad sense* for a subject's skepticism, even if those experiences won't count as cures in the narrow sense.

From now on, when I use the term "cure," I intend to speak of cures in the broad sense.

Once we recognize the necessity of broadening our conception of what might generate a requirement to abandon skepticism, my earlier argument for the conclusion that moral skepticism is treatment resistant runs into trouble. One problem is that some people will undoubtedly be attracted to the moral analogue of the generalized form of dogmatism: this would be the view that, in the absence of defeaters, you are required to take your *moral* experiences at face-value, even

if you do not have antecedent reason to think that those experiences are reliable. Some philosophers hold that there are moral experiences.²¹ If they are right, the possibility presents itself of holding that such experiences might generate epistemic requirements for moral skeptics to abandon their skepticism. This view is sufficiently plausible that we should consider whether there is any way to reformulate the argument for the (Weaker Conclusion) so as to make it compatible with this position.

Considering whether the argument for the (Weaker Conclusion) can be made compatible with moral dogmatism can also help us evaluate another possible objection to that argument. For even if there are no moral experiences, it might be the case that moral intuitions can play much the same role in moral epistemology. Thus, one might be an “intuition dogmatist” and hold that one is rationally required to take one’s intuitions at face-value, even if one has no reason to think those intuitions accurate, so long as those intuitions do not conflict with any of one’s beliefs. If one finds this position plausible, then one might suggest that intuitions by themselves are able to epistemically compel every moral skeptic to give up her skepticism. Intuition dogmatism does seem to be well-motivated, since the basic methodology of ethics presupposes that one is required to take one’s moral intuitions seriously, at least until one has reason to doubt their reliability. So this view also needs to be addressed.

3.5 Moral experiences

About the first view—the view that all cases of moral skepticism can be cured by moral experiences—the first question to ask is: are there really any moral experiences?

²¹For recent examples, see McGrath [24] and Watkins and Jolley [44].

Well, what would a moral experience be if there were any? I take it that a moral experience is supposed to be an experience with a special sort of content—namely, a moral content, such as that some particular action or practice is morally permissible, obligatory, or forbidden. Of course, there are controversies over what the content of an experience is, and whether experiences even have contents, but I am going to set these debates aside to see how far we can get without diving into them.²² To make things more difficult for myself, I will also assume that there are moral experiences.

It might seem that if such experiences exist, and if moral dogmatism is correct, then it follows that all cases of moral skepticism will be curable. But, in fact, this does not follow.

One might suspect that this does not follow because some moral skeptics—the error theorists and the like—will have beliefs that conflict with their moral experiences. But it is not clear that merely having beliefs that conflict with one's moral experiences will prevent one's experiences from epistemically requiring one to form some moral beliefs. Generalized moral dogmatism states that one is required to take one's moral experiences at face-value, except if one possesses a defeater. But can any belief count as a defeater simply by conflicting with one's experiences? Dogmatists can go different ways here. There is a debate in epistemology about what it takes for one belief to defeat justification in another. Some philosophers (such as Plantinga and Bergmann) say that defeaters are very easy to come by; according to them, simply believing that one's perceptions are unreliable is enough to defeat any justification that those perceptions might otherwise provide.²³ If this is right, then it might be the case that many moral skeptics will have defeaters for their moral beliefs. However, others (such as

²²For an overview of controversies surrounding the idea that experiences have content, see Siegel [37].

²³Plantinga [29], 364-366; Bergmann [2], 161-168.

Alston) maintain that defeaters are much harder to come by.²⁴ This debate has to be settled in order to figure out whether error theorists, simply by being error theorists, have beliefs that will prevent their moral experiences from curing their skepticism. As I am more sympathetic to Alston's view, I am inclined to think that simply possessing beliefs that conflict with one's moral experiences is not enough to defeat the justification that those experiences provide.

However, there is a different way to show that even if moral experiences exist and moral dogmatism is correct, it does not follow that every case of moral skepticism will be curable. Even if we grant that all skeptics have moral experiences, it will still be the case that some moral skeptics will believe falsely, though *for good reason*, that they do not have moral experiences. The very fact that there is a debate about the existence of moral experiences shows that it is not clear to everyone that they have moral experiences. But if there is a debate about the existence of moral experiences, then I think we must recognize that some people will have good reason to believe that they don't have any—even if in fact they do. If one believes for good reason, though, that one does not have moral experiences, then that will be sufficient to defeat the prima facie justification that those experiences would otherwise provide. This, in short, is my argument, which I will now explain in more detail.

Suppose, for example, that someone who does have moral experiences looks inside to try to figure out whether she has some. If moral experiences are murky enough that there is a debate about their existence, it is conceivable that our subject might come up empty-handed. She might do her best to look for moral experiences, but she still might not find any—even if they are there. A comparison to a different sort of introspective task is helpful to think about here. There is a debate among philosophers about whether thoughts have a distinc-

²⁴Alston [1].

tive phenomenology, aside from that of the sensory images that accompany them. (What it is like to think that Rob Ford should resign is different from what it is like to think that it is starting to get cold outside, but some philosophers maintain that “this difference can arguably be pinned down to differences of sub-vocal speech, or other sorts of imagery that might go along with thinking these thoughts.”²⁵ The question is whether the phenomenology of thought is exhausted by the phenomenology of these accompanying states, or whether thought itself has a phenomenology of its own.) As Eric Schwitzgebel reports in a recent paper, at a conference devoted to this debate, philosophers were split, with 17 participants saying that thought does have a distinctive phenomenology and 8 saying that it doesn’t.²⁶ What is interesting about this split is that it is so clearly true that everyone on one side must be wrong: either thought does have a distinctive phenomenology or it does not; it seems highly unlikely that thought has a distinctive phenomenology for some people, but not for others. A sizable percentage, then, is making a mistake either way, even though the question about which they are mistaken should, in principle, be answerable via introspection.

Although a sizable percentage of Schwitzgebel’s philosophers are mistaken about whether thought has a distinctive phenomenology, there is pressure to say that all of them are justified in holding their respective views. If one looks inside and seems to see a phenomenology of thought, then, in my view, the deliverances of introspection give one *prima facie* reason to think that thought does have a distinctive phenomenology (in the absence of reason to think one’s introspection is unreliable). And if one looks inside and there does not seem to be a distinctive phenomenology of thought, then that should justify one in be-

²⁵Tye and Wright [43], 330.

²⁶Schwitzgebel [36], 257-258.

believing that thought does not have a distinctive phenomenology (again, in the absence of reason to think one's introspection unreliable). But if that is correct, then why shouldn't something similar be true of people who introspect and find that they do not have any moral experiences? If they look inside for some and come up empty-handed, then why shouldn't they be justified in believing that they don't have any? I can't imagine why such a person shouldn't trust the deliverances of introspection.²⁷ I think we must conclude, then, that such a person would be justified in believing that they don't have any moral experiences. Thus, some moral skeptics will be justified in believing that they do not have any moral experiences.

If I am right about this, then something very important follows. What follows is that for such people, moral experiences cannot play the role of a cure for their skepticism. The reason is simple. In general, if someone justifiably believes that their experience does not have content *p*, then their experience cannot epistemically compel them to believe *p*. This principle provides the best explanation of our reactions to certain cases involving introspective misidentification. Suppose, for example, that I am having an experience as of a hen with 15 speckles; I attempt to count the speckles that I am experiencing, but my quick count of the speckles goes wrong. Thus, I form the belief that I am not seeing 15 speckles. If my count was performed responsibly, then that belief is justified. But if I really am justified in believing that I am not seeing 15 speckles, then it is very hard to imagine how it could be that I am epistemically required to believe that I am seeing 15 speckles. The reason why I am not epistemically required to form that belief, I maintain, is that the above principle is true.

²⁷Maybe *some* people have sophisticated philosophical reasons not to trust the deliverances of introspection (perhaps because they've read Schwitzgebel [36]). But I take it that most people who haven't read Schwitzgebel don't have such reasons to doubt the reliability of introspection. So they won't count as having the relevant kind of defeaters.

But if it is true, as it seems to be, that an experience cannot epistemically require someone to believe p if that person justifiably believes that their experience does not have content p , then it follows that someone who justifiably believes that they are not having any moral experiences cannot be epistemically required by those experiences to form any moral beliefs. And that means that for such a person, moral experiences cannot constitute a cure (in the broad sense) for moral skepticism.

3.6 Moral intuitions

I have just argued that even if moral experiences do exist, there is reason to think that they will not be capable of curing all cases of moral skepticism. What I now want to do is to shift the focus of our discussion from moral experiences to moral intuitions, in order to explore whether intuitions can succeed where experiences failed.

To make things difficult for myself, I will assume what I earlier called “intuition dogmatism”—the view that we are epistemically required to take our intuitions at face value, at least so long as we don’t have reason to think those intuitions are unreliable. If it seems to one that some moral propositions are true, and one has no reason to think that one’s intuitions are unreliable, then one would be required to reject moral skepticism. The question remains, though: can moral intuitions cure every case of moral skepticism?

The answer, I suspect, is that they cannot, for two reasons. First, even if many people have moral intuitions, some people might not. If there really are any thoroughgoing error theorists for whom nothing appears to be wrong, their skepticism will not be cured via intuition, even if moral intuitions do suffice

to cure many other people's skepticism.²⁸ Second, even if such people do not exist—perhaps because everyone feels at least some twinge of sympathy for some moral claims—we must also take into account anti-morality intuitions.²⁹ The requirement to take one's intuitions at face-value can't be confined to moral intuitions. Consider, then, someone who has an intuition that murder is wrong, but who also has strong doubts that anything is wrong. Must this person take their moral intuitions at face-value, even if such doubts (which might be taken to be a kind of intuition) are present? It is hard to see why they would be so required. It seems, then, that such a person would not be required to abandon their skepticism.

Furthermore, it might very well be that there are some people who have only what I'll call "flickering" intuitions. Sometimes, it might seem to them that murder is wrong. But at other times, it might seem to them that it isn't. (They might have moods when it seems that "it's all just atoms and the void" and nothing is morally required.) It is unclear what intuitionism dogmatism would suggest about such a person. Are they epistemically required to believe in moral facts only when they have moral intuitions? But then what about the times when they are not having them? There seems to be no good reason to say they are epistemically compelled to take seriously intuitions they had in the past, the contents of which no longer seem compelling. So even if everyone does have moral intuitions, and thus all skeptics are cured at least some of the time, there will still be times when some people are not epistemically required to abandon their skepticism.

²⁸Richard Joyce makes a similar point in Joyce [19], 216.

²⁹Again, see Joyce [19], 216 for a similar observation.

3.7 The Transition Objection

So far, I have argued that even if one grants various assumptions that initially appear hostile to the view that that some moral skeptics will not be cured, it is still possible to show that this view is correct. But even if my argument so far shows that some possible cases of moral skepticism will not be cured by arguments, experiences, or intuitions, one might worry that those cases of moral skepticism might, nonetheless, fail to be *reasonable* or *epistemically permissible*. For it might be that in order for an agent's skepticism to be reasonable, the skeptic must have become a skeptic for a good reason. But for someone to become a skeptic for a good reason, there must be some good skeptical arguments. Thus, in order to show that there are reasonable skeptics, it would be necessary to pick a skeptical argument and to defend it.

I refer to the above objection (that a skeptic who isn't epistemically required to abandon her skepticism because of experience, intuition, or argument is not necessarily a reasonable skeptic) as *The Transition Objection*. The reason why is that the objection relies on the following *Transition Principle*:

A doxastic state can't be reasonable unless the transition to that doxastic state was epistemically respectable.

The Transition Principle implies that an agent's skepticism can't be reasonable unless the agent's transition to skepticism was epistemically respectable. But an agent's transition to skepticism can be epistemically respectable, one might think, only if the agent became a skeptic because she was exposed to a good argument for skepticism. Thus, it seems that in order to argue that skepticism can be reasonable, I need to pick an argument for skepticism and defend it.

It is worth noting that the Transition Objection not only threatens the view that some cases of *moral* skepticism are reasonable, but it also threatens the view

that some cases of *external world* skepticism are reasonable. The Transition Objection suggests that even if epistemologists are correct that it might not be possible to epistemically compel thoroughgoing external world skeptics to abandon their skepticism, it does not follow that external world skepticism is ever reasonable. That would follow only if there were some good arguments for external world skepticism.

My response to the Transition Objection is two-pronged. Either the Transition Principle faces counterexamples, or the objection at most forces me to make a minor concession—to acknowledge that the *reason* for which one becomes a skeptic matters. But acknowledging that the reason for which one becomes a skeptic matters does not force me to say that one can become a reasonable skeptic only if one is persuaded by a good skeptical argument.

To begin, I will explain why the Transition Principle seems to face counterexamples. A counterexample to the Transition Principle would involve a person who ends up in a reasonable doxastic state even though the person's transition to that doxastic state was not epistemically respectable. The following seems to be such a case:

Suppose that Tom believes that Rob Ford resigned today, but he has absolutely no evidence for or against that belief. Tom then decides for a bad reason (or maybe no reason) to give up that belief. However, once Tom has given up that belief, he comes to his senses and realizes that he never had any evidence for it. So he decides that he's right to suspend belief in the claim *Rob Ford resigned today*.

In this example, Tom transitions from believing that Rob Ford resigned today to suspending belief in the claim *Rob Ford resigned today*. Tom's transition to his final doxastic state (suspending belief in the claim *Rob Ford resigned today*)

was not epistemically respectable, since he had absolutely no reason to perform this transition. However, the doxastic state that Tom ends up in seems to be eminently reasonable.

Defenders of the Transition Principle might respond to the Tom case by arguing that I have misdescribed it. Really, there are *two* transitions involved in the case just described. First, there is the transition Tom makes (for no reason) to suspension of belief. But there is then a different kind of transition: Tom re-evaluates his epistemic state and changes the reason for which he suspends belief in the claim that *Rob Ford resigns today*, thereby coming to suspend belief in this claim because he recognizes that he lacks evidence for the claim. Defenders of the Transition Principle might maintain that the first transition, while not epistemically respectable, does not result in a reasonable doxastic state, and thus does not provide a counterexample to the Transition Principle. Defenders of the Transition Principle could also maintain that the second transition is epistemically respectable, and so cannot provide a counterexample to the Transition Principle either.

Even if we grant that epistemic transitions should be individuated in a way that allows one to say that there are two transitions in the Tom case, there is still a way to undercut much of the force of the Transition Objection. Suppose that someone becomes a skeptic because she has been convinced by a bad skeptical argument. However, once she becomes a skeptic, she decides that the argument is no longer so convincing after all. Even so, she decides to remain a skeptic, simply because she can't see any reason to abandon her skepticism. The position of such a skeptic is structurally analogous to the position of Tom. The skeptic's initial transition to skepticism was not epistemically respectable. However, after becoming a skeptic, the skeptic reevaluates her skepticism and undergoes

a second transition, changing her reasons for being a skeptic. If the agent really *does* lack experiences or intuitions sufficient to compel her to abandon her skepticism, then her skepticism, it would seem, must be reasonable.

Thus, an agent's reasons for becoming a skeptic might very well play a role in determining whether the agent's skepticism is reasonable. However, an agent led to skepticism by a bad skeptical argument can, nonetheless, end up a reasonable skeptic. What is required is that the agent reappraise her epistemic state once she becomes a skeptic. If she does, and she correctly recognizes that she is not epistemically required to abandon her skepticism, then her skepticism can be epistemically permissible after all.

3.8 Epistemically required skepticism

We have now seen that some mature agents can be reasonable moral skeptics, even if some agents are epistemically permitted to adopt a non-skeptical view of morality. But a line of reasoning much like the one that I have used to support this claim can be used to establish a somewhat stronger conclusion. Once supplemented by the compelling premise that a belief cannot be held reasonably in the absence of epistemic support of some kind, it can be shown not only that some mature agents can be epistemically *permitted* to suspend belief in all substantive moral claims, but it can also be shown that some mature agents can be epistemically *required* to suspend belief in all such claims. In this section, I will explain how to show this.

First, it is important to realize why the claim that some agents can be epistemically required to suspend belief in all substantive moral claims is stronger than the claim that some agents can be epistemically permitted to suspend belief in all such claims. The reason why the former claim is stronger than the

latter is that the latter claim is compatible with the idea that there is more than one epistemically permissible stance to take with respect to substantive moral claims. If there were just a single permissible stance that one could take with respect to substantive moral claims, then, arguably, one would be epistemically required to adopt that position.³⁰ But if there were several epistemically permissible stances available to one, then it could be that none of those positions taken individually is epistemically obligatory.

An analogy to rationality in action is helpful in bringing out this possibility. In some situations, there is a unique action that one is rationally required to perform. For example, imagine that, desperately wanting some ice cream, I face a choice between twenty five flavors that I hate and one that I love; surely, the only rational option for me—the option that I am rationally required to take—is to choose the flavor that I love. But in some situations, one is rationally required to take one of a set of options, each of which is permissible to take, but none of which is rationally required. Imagine, for example, that I am deciding between

³⁰I say “arguably” here because whether this is true depends on subtleties concerning the relationship between epistemic requirements and epistemic permissions.

The natural way to support the claim that *if one has just a single permissible stance that one can take, then one is rationally required to adopt that stance* would be to appeal to a principle of standard deontic logic. Standard deontic logic holds that for a state of affairs to be permitted is for the negation of that state of affairs not to be required (i.e., $Px \leftrightarrow \neg R\neg x$). This principle is, of course, equivalent to the principle that says that for the negation of a state of affairs to be impermissible is for that state of affairs to be required (i.e., $\neg P\neg x \leftrightarrow Rx$). Now, if one has only a single doxastic option that is epistemically permissible—call it x —then $\neg P\neg x$. But if $\neg P\neg x \leftrightarrow Rx$, then it follows that Rx . Thus, one’s single permissible option is epistemically required.

A potential problem with this argument is that the principle $\neg P\neg x \leftrightarrow Rx$ entails that there can be no epistemic requirement dilemmas—i.e., it entails that for any x , it cannot be true that $Rx \& R\neg x$ —when it is combined with the plausible principle that $Rx \rightarrow Px$. (To see this, note that if $R\neg x$, the former principle entails that $\neg Px$. But then Rx cannot also be true, as the latter principle would entail Px , which would yield a contradiction.) One might want to admit the possibility of epistemic requirement dilemmas. If one does, then one must reject either $\neg P\neg x \leftrightarrow Rx$ or $Rx \rightarrow Px$; it’s not clear, however, which principle should be rejected. (For discussion of this problem, as well as related issues in deontic logic, see McConnell [23], sec. 4.)

I want to remain neutral on the question of whether epistemic requirement dilemmas are possible, and, if so, whether one should reject the principle that $Px \leftrightarrow \neg R\neg x$ or the principle that $Rx \rightarrow Px$. Accordingly, I will content myself to say that it is “arguable” that an agent with a single epistemically permissible doxastic option is epistemically required to take that option.

vanilla, chocolate, and strawberry ice cream, and I have no preferences about which flavor to pick; in this case, I am rationally required to adopt one of a set of options—I must choose vanilla, chocolate, or strawberry—but none of these options taken individually is rationally required. If rationality in belief (or suspension of belief) is analogous to rationality in action, then it seems that we must admit that analogous scenarios could occur in which a person faces a doxastic choice between several possible doxastic stances that she might take with respect to a proposition (or set of propositions), any of which would be epistemically permissible for her to adopt, but none of which is epistemically required.

To admit the possibility of this kind of situation is to reject the thesis that Roger White has called “Uniqueness.” Uniqueness states that “given one’s total evidence, there is a unique doxastic attitude that one can take to any proposition.”³¹ As White has noted, many philosophers accept this thesis, though some have explicitly rejected it.³² I will not discuss here whether Uniqueness is correct; that is a complicated topic that I do not have space to properly address. It will suffice to note the following. If Uniqueness is true, then, arguably, establishing that it can be epistemically permissible to suspend belief in all claims of the form *Xing is obligatory* is enough to show that that stance is also epistemically *required* for some people.³³ Rejecting Uniqueness, on the other hand, makes my life harder, since it raises the complicating possibility described above. What I propose to do, then, is to work under the assumption that Uniqueness might be false, so as to consider the question of whether, on that assumption, my argument can be used to show that suspending belief in all claims of the form *Xing is obligatory* can be epistemically required for some agents.

³¹White [45], 445.

³²Ibid.

³³Again, this depends on the issues mentioned in footnote 31.

The answer, I believe, is that even if the Uniqueness thesis is false, the basic line of reasoning developed in earlier sections of this paper can be used to show that moral skepticism (that is, the stance of suspending belief in all claims of the form *Xing is obligatory*) can be epistemically required for some agents. To see how it can be used to show this, consider an agent of the sort that I have already argued can reasonably remain skeptical about all claims of the form *Xing is obligatory*. Such an agent satisfies three conditions: she lacks any substantive moral beliefs, she lacks undefeated moral experiences, and she has only flickering moral intuitions or conflicted moral intuitions. If the Transition Principle is true, she also must either have transitioned to skepticism for a good reason or have come to change her reason for maintaining her skepticism if she didn't. Satisfying these conditions is sufficient to permit an agent to suspend belief in all claims of the form *Xing is obligatory*. Notice, however, that satisfying these conditions is also sufficient to prevent her epistemic state from being able to support to any significant degree a belief of the form *Xing is obligatory*. If we assume that an agent can be epistemically permitted to form a belief only if she has some significant support for that belief—whether that support takes the form of argument, experience, or intuition—then it follows that any agent satisfying the three conditions previously mentioned is not only epistemically permitted to be a moral skeptic, but she is also epistemically *required* to be a moral skeptic.

The key assumption here—that an agent can be epistemically permitted to form a belief only if she has some significant support for that belief—is a quite modest claim, which can be accepted by a wide range of people. This assumption can be accepted by both coherentists—who will want to say that a belief's support is constituted by coherence with other beliefs that the relevant agent has—and by foundationalists, many of whom will want to say that the sup-

port that an agent's beliefs have ultimately bottoms out in some kind of non-doxastic state. And, importantly, this assumption can be accepted by those who reject Uniqueness, since one can maintain that a belief must have some significant support in order for it to be epistemically permissible to have it while also maintaining that a single body of evidence can, in some cases, significantly support more than one doxastic attitude towards the same content.

The only philosophers who cannot accept that a belief must have significant support in order for that belief to be epistemically permissible to possess are those who think that rationality is so permissive as to allow belief in the absence of any or only a minimal amount of support. Some philosophers do accept that belief can be epistemically permissible in such a case. Those who maintain "conservative" views in epistemology, according to which believing that p gives one a *prima facie* reason to continue to believe that p , fall into this camp.³⁴ So do those who agree with Pascal that beliefs can, in some cases, be epistemically permissible for pragmatic reasons alone.³⁵

These philosophers will not be convinced by the argument just given. Of course, not everyone can be convinced all of the time. But there is something to say even to these people in order to motivate the idea that moral skepticism can be an epistemically mandatory stance for some people to adopt. Given my earlier argument that moral skepticism can be an epistemically *permissible* stance for some people to adopt, to deny this claim would be to say that every moral skeptic is such that moral skepticism *and* belief in some claims of the form *Xing is obligatory* are permissible doxastic options. It is very hard to see, however, how this could be the case, even on a conservative or a Pascalian view. It is not clear that a non-skeptical view of morality has any pragmatic advantages

³⁴Harman [15], chap. 4.

³⁵Chignell [6], sect. 6.1.

over moral skepticism, so it is hard to see how it could be argued that a moral skeptic who satisfies the three conditions mentioned above could be justified on pragmatic grounds in adopting some belief of the form *Xing is obligatory*. It is also very hard to see how a conservative epistemology could justify the view that all moral skeptics would be epistemically permitted to abandon their skepticism. Indeed, if anything, a conservative epistemology would seem to support the idea that someone who has become a moral skeptic has reason *not* to abandon her moral skepticism.³⁶ Thus, even those who reject Uniqueness will have a hard time justifying resistance to the idea that some moral skeptics can be epistemically required to maintain their skepticism.

3.9 Conclusion: psychological issues

At this point, I have argued that both coherentist and broadly dogmatist moral epistemologies lead to the conclusions that some cases of moral skepticism will not be cured and that some moral skeptics will be epistemically required to maintain their skepticism. But a question remains: are there really any moral skeptics? That is, are there really any people who accept only those moral claims that follow as a matter of logic from their non-moral beliefs? And if there are, how likely is it that any of us will ever run into one of these people on the street?

One might think that if there are any moral skeptics, there are very few. The reason: it seems psychologically very hard to put oneself into the moral skeptic's position, as we can see if we think about what it takes to be a skeptic about the external world. An external world skeptic suspends belief in all external world

³⁶That is, so long as conservatism is generalized so as to become a view not only about *belief*, but also other doxastic attitudes (such as suspension of belief); a natural generalization along these lines is the following: any agent who has adopted doxastic attitude A towards proposition p has a prima facie reason to continue to adopt A towards p.

claims, except those that are entailed by her non-external world beliefs. But is it even possible to do that? Aren't our external world beliefs so firmly intertwined with the rest of our beliefs that they cannot be separated out and abandoned? If so, one might wonder whether our moral beliefs are just as hard to reject.

The worry here is a very deep one, which deserves an extended treatment. And it conceals an equally deep question, which also deserves attention: are our external world beliefs really as deeply held as they might at first appear? Just how hard is it, psychologically speaking, to become a skeptic about mid-sized dry goods?

One might think that adopting the skeptic's position is not that hard. Arguably, it doesn't require one to abandon the deliverances of one's senses; all it requires is a reinterpretation of them.³⁷ Similarly, one might maintain that adopting the moral skeptic's point of view doesn't require rejecting the sort of behavior endorsed by the moral point of view. All that skepticism requires is that one not hold that such behavior is a moral requirement. And it doesn't seem so hard to imagine that someone could adopt that stance, even if one ultimately disagrees with it.

On the other hand, however, one might think that adopting moral skepticism is more like becoming a solipsist.³⁸ It is also arguable that the solipsist needn't reject the deliverances of her senses. But it seems very hard to imagine realistic examples of agents who truly believe the solipsist's interpretation of the senses. If becoming a moral skeptic is more like becoming a solipsist than it is like becoming an idealist, then the interest my conclusion holds is not nearly as great as it might first have appeared.

³⁷Of course, to argue *this*, one has to take a stand on what the content of one's perceptions are. In order to maintain the view that adopting idealism needn't require the rejection of one's sensory perception, one must maintain that the contents of one's perceptions do not imply or presuppose the existence of an external world. This, however, is a controversial view.

³⁸I owe this comparison to Nicholas Sturgeon.

At the present time, I am not sure how to resolve this issue. It seems to me that at bottom it is closely related to the issue of what the conceptual presuppositions of moral discourse are—an issue that looms large in contemporary metaethics. The thinner the conceptual presuppositions of moral discourse are, the harder it will be to suspend belief in all moral claims (or at least those that don't follow as a matter of logic from one's non-moral beliefs). How thin those presuppositions are I won't venture a guess at here. It seems to me a very difficult question that demands further attention. That said, what I have done here is to argue that however hard it is to become a moral skeptic, some people who succeed in doing so will be very hard to rationally dislodge from their position.

CHAPTER 4
BLAME AND MORAL IGNORANCE

4.1 Introduction

It is widely acknowledged that non-culpable ignorance sometimes constitutes an excuse for bringing about a bad outcome. Suppose that I call my mother, thereby waking her from her nap. If I had no reason to think that she was sleeping, then, although the result of my action was unfortunate, I do not deserve to be blamed for acting the way I did.¹

Given that non-culpable ignorance sometimes excuses, it is natural to wonder whether it always does. As it turns out, it does not. Suppose that a driver intentionally hits a pedestrian whom the driver had never met before. Then even though the driver is non-culpably ignorant of the pedestrian's name, the driver's ignorance does not constitute an excuse.²

Still, even if non-culpable ignorance does not *always* excuse, it does excuse in some cases. One would like an account of the conditions under which non-culpable ignorance does constitute an excuse and the conditions under which it does not. Unfortunately, there is little agreement among philosophers about what these conditions are.

¹Throughout this paper, I assume that a person is blameworthy just in case she is the appropriate target of resentment and indignation. This is the view of blameworthiness assumed by Rosen [33], whose work is the explicit target of the authors I will discuss, such as Harman [14] and Talbert [42].

Not all philosophers accept this view of blame. (For an explanation of why some philosophers reject it, see the introductory chapter of Coates and Tognazzini [8].) However, even if one adopts a view according to which blameworthiness amounts to something other than being the appropriate target for resentment and indignation, the main thrust of this paper can be preserved. Instead of speaking of blame and blameworthiness, we can speak of resentment, indignation, and the conditions under which an agent is an appropriate target of these attitudes. The issues thus translated will remain just as important, and the arguments I offer will still go through.

²Compare Rosen's example of Jones and the Sikh (Rosen [33] f.n. 4), which illustrates the same point.

One reason why there is little agreement about when non-culpable ignorance excuses is that philosophers disagree about whether a particular type of non-culpable ignorance—namely, NCFPO ignorance—excuses. One group of philosophers defends what I will call the Forgiving View: the view that NCFPO ignorance constitutes an excuse for wrongdoing.³ Opposed to this group are the philosophers who defend what I will refer to as the Unforgiving View, which states that NCFPO ignorance does not constitute an excuse for wrongdoing.⁴ Despite several notable contributions in the past few years, the literature on the Forgiving and Unforgiving Views shows few signs of convergence.

That philosophers disagree about whether NCFPO ignorance excuses is not the only reason why there is little agreement about when non-culpable ignorance excuses. Another reason why there is little agreement about when non-culpable ignorance excuses is that the question of when non-culpable non-moral ignorance excuses has been largely neglected. This question has been addressed in passing in the literature on the Forgiving and Unforgiving Views, but it has not yet been discussed at length. Furthermore, many of the answers to this question that have been offered in passing are problematic.⁵

This chapter argues that by paying careful attention both to the question of when non-culpable non-moral ignorance excuses and to the question of what determines *how much* blame an agent deserves, one can find material for a compelling argument in favor of the Forgiving View. The argument I will present in favor of the Forgiving View has a number of advantages over the standard argument for the Forgiving View. One advantage is that my argument provides objections to two *prima facie* plausible alternatives to the Forgiving View that

³Recent defenders of the Forgiving View include Fitzpatrick [12] and Rosen [35].

⁴Recent defenders of the Unforgiving View include Harman [14] and Talbert [42].

⁵As I will show later. See section 4.6, where I discuss some of the shortcomings of Harman's and Talbert's accounts of non-culpable non-moral ignorance.

have been overlooked in the literature on moral ignorance. By exploring these views, I will not only be sharpening the case for the Forgiving View, but also showing that the range of possible views on non-culpable ignorance is wider than has been previously recognized.

The plan for this chapter is the following.

After some preliminary remarks, I will introduce the distinction between partial and complete excuses and argue that this distinction complicates the usual way of marking the difference between the Forgiving and Unforgiving Views. Once one recognizes the possibility of partial excuses, one can see that the distinction between the Forgiving and Unforgiving Views needs to be redrawn so as to leave room for an additional view, which I refer to as the Partially Forgiving View.

Next, I will formulate my “abductive argument” in favor of what I call the “Forgiving Package”—the conjunction of the Forgiving View with two other views that I will introduce in section 4.5. After outlining the abductive argument and explaining how the abductive argument improves on previous arguments in favor of the Forgiving View, I will explain what shape a successful objection to the abductive argument would need to take.

In the rest of this chapter, I will explore whether opponents of the Forgiving View can provide a convincing objection to the abductive argument. This occupies sections 4.6-4.8. Because the abductive argument has never before been discussed, these sections require me to do a bit of work on behalf of opponents of the Forgiving View. As we shall see, some remarks from opponents of the Forgiving View suggest the outlines of responses to the abductive argument, but these outlines need to be filled in before these responses can be properly evaluated. It is also necessary to correct some mistakes made by opponents of

the Forgiving View. I will argue that even after these mistakes are corrected and the outlines filled in, the objections that result do not succeed in refuting the abductive argument.

4.2 Epistemological preliminaries

As I have already indicated, not all ignorance is epistemically on a par. Some ignorance of obligations is *culpable*, and some is *non-culpable*. To a first approximation, we might say that if a person S is ignorant of the fact that she is obligated to do X, we can determine whether S's ignorance is culpable by answering the question, "does S's evidence adequately support believing that she ought to do X?"⁶ If the answer is "no," then her ignorance is non-culpable. If the answer is "yes," then her ignorance is culpable.⁷

We have already seen a case that arguably involves culpable moral ignorance. It's wrong for me to call my mother to chat if doing so will wake her from her nap. And if my mother has warned me not to call her at a certain time

⁶Here I am assuming that whether a person's ignorance is culpable or non-culpable is a purely *epistemic* matter, which means that I am working with a somewhat different notion of culpability than Rosen, who claims that the relevant notion is not epistemic, but instead *moral* (see f.n. 3 of Rosen [33]). Note, however, that Rosen never explains why he thinks that the relevant notion of culpability is moral, and this view of culpability seems to play no role in his argument.

⁷As my use above of the phrase "to a first approximation" suggests, there are several complications here. A first complication is that, in some cases, a person's ignorance of the fact that she is obligated to do X might be non-culpable even if her evidence supports thinking that she ought to do X. For example, if a person's evidence is sufficiently complicated that she lacks the cognitive resources necessary to appreciate the force of that evidence, then she will not be blameworthy for her ignorance. A second complication is that a person's ignorance of the fact that she is obligated to do X might be culpable even if her evidence does not support thinking that she ought to do X. If, for example, a person has purposefully isolated herself from relevant evidence, then she is to blame for the fact that she is ignorant of what she is obligated to do.

While these complications are important to note, neither endangers any of the arguments in this paper. None of the examples I will rely upon involve agents whose evidence is particularly complicated, nor will any of the examples in this paper involve people who have purposefully isolated themselves from relevant evidence. Thus, the rough test that I propose above will be sufficient for my purposes here.

because she'll be sleeping, then I have evidence that supports thinking that it would be wrong for me to call her at that time. Thus, if I fail to believe that it would be wrong to call her at that time, I am culpably morally ignorant of that fact.

It's much harder to give an example of *non-culpable* moral ignorance—especially NCFPO ignorance—for reasons that I'll explain when I discuss “epistemically-sensitive” views of obligation.⁸ For now, though, the reader need only keep in mind that any case in which a person is obligated to perform a certain action X, but she justifiably believes that she is not obligated to do X, will count as a case of non-culpable moral ignorance.

4.3 The state of the literature

Reading the literature on the Forgiving and Unforgiving Views can leave one with the impression that the debate has come to a stalemate. One reason why one can come away with this impression is that much of the literature focuses on formulating vivid examples of NCFPO ignorant agents who are supposed to be either clearly blameworthy or clearly blameless for acting wrongly—the hope being that some such example will persuade readers to adopt the author's preferred view. The problem is that so far neither side has formulated a case that has put much pressure on their opponents. And it seems unlikely that the dominant methodology will do any better in the future. Even some of the authors participating in the debate seem to recognize this point. For example, Gideon Rosen—a defender of the Forgiving View—says in a recent paper that what he has provided is “not so much an argument as it is a vivid restatement of [his] position” (Rosen [35] 609). And Elizabeth Harman—a defender of the

⁸See sections 4.5.1 and 4.7.

Unforgiving View—says in a recent paper of hers that she has “not aimed to establish” her view, but only “to set it out as an alternative to the picture Rosen offers” (Harman [14] 468).

Rosen’s appeal to his ancient slaveholder case provides a good illustration of the state of the debate. Rosen argues that an ancient Hittite slaveholder who seriously considered whether slavery was morally permissible, but was simply unable to find any compelling reason why it wouldn’t be, would count as non-culpably ignorant of the wrongness of slaveholding.⁹ Furthermore, Rosen claims that “when we focus on the act [i.e., enslaving people] and the objectionable attitude that underlies it, we are no doubt powerfully inclined to blame—so long as we ignore the stipulated fact that he is blameless for not knowing that his slave deserves much more. But when we bear this further fact in mind—when we ‘zoom out’, as it were—then (I claim) our sense of his culpability evaporates.”¹⁰ However, Rosen’s opponents have begged to differ with this characterization of “our” sense of the slaveholder’s culpability. Harman, for example, has suggested that even when we bear in mind the fact that the slaveholder is blameless for not knowing that slavery is wrong, it is implausible that this constitutes an excuse.¹¹

The cases that have been used to motivate the Unforgiving View suffer from a similar problem. For example, in the course of defending the Unforgiving View, Harman appeals to several cases of agents whom, she claims, “we would ordinarily take . . . to be blameworthy,” even if these agents came to believe that they were acting permissibly “without having behaved poorly in the management of [their] beliefs.”¹² But anyone who finds Rosen’s judgments about the

⁹Rosen [33], 64-66.

¹⁰Rosen [33], 73.

¹¹Harman [14], 465.

¹²Harman [14], 358. These examples are presented in Harman [14], sect. 3.

slaveholder plausible will surely deny that Harman's agents are blameworthy, and so will take issue with Harman's description of the view "we" would take of them. It seems, then, that the problematic dialectical features of Rosen's argument are possessed by Harman's as well.

Relying on cases of NCFPO ignorant agents who are supposed to be either clearly blameworthy or clearly blameless is unlikely to get us anywhere. Defenders of the Forgiving View will say what their view implies about such cases, and defenders of the Unforgiving View will say what their view implies. Furthermore, those unsure of what to say about the standard cases are unlikely to gain clarity by being presented with even *more* of them. In order to make progress, then, it seems that we must find some phenomenon that both defenders of the Forgiving View and defenders of the Unforgiving View acknowledge, or some assumption they share, which we can leverage into an argument for one of these views. Or, at the very least, we should find some phenomenon that a significant number of the uncommitted will acknowledge, which provides some dialectical leverage.

My abductive argument for the Forgiving View aims to do both of these things. The argument appeals to two different phenomena (or, as I will call them, *observations*), one of which is acknowledged by everyone who is party to the debate about the Forgiving and Unforgiving Views, and the other of which has a much better chance of being acknowledged by the uncommitted than the examples to which Rosen appeals. I will explain shortly what these observations are and why they support the Forgiving View, but I mention the basic structure of the argument now to indicate why my argument represents an advance over the traditional arguments found in the literature.

Before I explain the observations and outline the abductive argument, how-

ever, it will first be necessary to introduce the important distinction between partial and complete excuses.

4.4 Sharpening the Forgiving and Unforgiving Views

One surprising feature of the literature on the Forgiving and Unforgiving Views has been the literature's focus on the notion of an unqualified excuse. Every case discussed focuses on whether an agent is blameworthy or blameless *simpliciter*. Even the standard way of defining the views, which I followed earlier, focuses on what the views say about whether NCFPO ignorance is an unqualified excuse. What is surprising about this way of proceeding is that it neglects the widely recognized point that there are two different types of excuse: complete excuses and partial excuses.

The distinction between complete and partial excuses is commonly invoked in legal settings. A standard example of a partial excuse is provided by provocation.¹³ The fact that someone was provoked to perform an assault does not typically make the assault morally *permissible*. But the presence of provocation often reduces the amount of blame that the attacker deserves. Someone who commits an assault after being provoked will, in many cases at least, be less blameworthy than someone similarly situated who was not provoked. To recognize that the attacker still deserves blame, but a reduced amount, the agent is said to have a *partial* excuse. A complete excuse, by contrast, would make blame entirely inappropriate.

Once we recognize the distinction between complete and partial excuses,

¹³While it is standard to treat provocation as a partial excuse, treating provocation in this way is not entirely uncontroversial. See, for example, the exchange between McAuley [22]—who argues that provocation does not excuse (even partially)—and Dressler [9], who defends the standard view.

drawing the distinction between the Forgiving and Unforgiving Views becomes more complicated than defenders of these views have noticed. Earlier, I said that the Forgiving View states that NCFPO ignorance constitutes an excuse, and the Unforgiving View states that NCFPO ignorance does not constitute an excuse. But notice that there are several different ways of sharpening these views once we clarify whether the notion of excuse in play is complete or partial. Does the Forgiving View state that NCFPO ignorance constitutes a complete excuse for wrongdoing, or does it state only that it constitutes a partial excuse? Or does the Forgiving View instead state only that it constitutes an excuse of some kind or other? Similar questions also arise for the Unforgiving View. Is the Unforgiving View the view that NCFPO ignorance does not constitute *any* sort of excuse, or is it the view that NCFPO ignorance does not constitute a complete excuse for wrongdoing?

Because defenders of the Forgiving and Unforgiving Views do not appeal to the distinction between complete and partial excuses, the only way to sharpen these views while remaining faithful to their spirit is to attend to the context in which they have been defended. Given that defenders of the Forgiving View often speak as if NCFPO ignorant agents do not deserve *any blame*, and defenders of the Unforgiving View seem to think that NCFPO ignorant agents' ignorance does not constitute *any* kind of excuse for wrongdoing, the best way to draw the distinction between these views in the following:¹⁴

¹⁴There is textual evidence that Rosen, for example, believes that NCFPO ignorant agents do not deserve any blame for wrongdoing. Take his claim about his ancient slaveholder that when we remind ourselves of the slaveholder's epistemic situation "our resentment is properly blocked"; this suggests that the ancient slaveholder does not deserve to be the target of *any* amount of resentment, which suggests that, on his view, NCFPO ignorance constitutes a *complete* excuse (Rosen [33], 73).

The evidence that defenders of the Unforgiving View, such as Harman and Talbert, do not believe that NCFPO ignorance constitutes even a partial excuse for wrongdoing is more indirect. The strongest argument for reading Harman and Talbert in this way is that the factors that determine, on their views, whether an agent is blameworthy—which, I assume, are closely related to the factors that they believe determine *how much* blame an agent deserves—are entirely

(The Forgiving View) NCFPO ignorance constitutes a *complete* excuse for wrongdoing.

(The Unforgiving View) NCFPO ignorance does not constitute *any* kind of excuse for wrongdoing.

Notice, however, that if we define the Forgiving and Unforgiving Views in this way, the distinction between them is no longer exhaustive. The possibility of the following intermediate view presents itself:

(The Partially Forgiving View) NCFPO ignorance constitutes a *partial* excuse for wrongdoing.

As far as I can tell, the Partially Forgiving View has been completely overlooked in the literature on moral ignorance. That this view has been overlooked is surprising; one would expect that those put off by philosophical extremes would naturally gravitate to this philosophical compromise. And yet this view seems never to have been discussed.

Recognizing the Partially Forgiving View as an option is important not only because the view has never before been explored, but also because the view's existence suggests that in trying to decide between the possible positions on NCFPO ignorance we should attend to judgments not only about which agents are *completely* excused, but also about which agents are *partially* excused. This represents a richer array of philosophical data to draw upon than that which has been used in the past. I will return to this point shortly, as it plays a crucial role in this chapter's overall argument.

uninfluenced by the presence of NCFPO ignorance. For evidence of this, see my discussion of Harman and Talbert in sections 4.6.1 and 4.6.2.

4.5 The abductive argument

I have already indicated that I will be arguing that the Forgiving View, once supplemented with a few auxiliary assumptions, provides the best explanation of two observations about blame. In this section, I will outline this argument and explain what these observations are.

4.5.1 The first observation to be explained

The first observation that plays into my argument is one that we have already seen: namely, that non-culpable *non-moral* ignorance excuses some of the time, but not always. Everyone—defenders of the Forgiving View and defenders of the Unforgiving View alike—agrees that this is the case. But when exactly *does* non-culpable non-moral ignorance excuse, and when does it not?

It might seem that the Forgiving View has nothing to say about this question, but in fact it does, at least if it is supplemented with a modest auxiliary premise.

To see what the Forgiving View has to say about this question, let's return to an earlier example. Suppose that I call my mother on the phone; unbeknownst to me, however, she was sleeping, and my call wakes her. If I had no reason to think that my mother was sleeping, then I do not deserve to be blamed for waking her. At least part of the reason why I am not blameworthy is that I was non-culpably ignorant of a certain non-moral fact—namely, that my mother was sleeping. One might think, however, that this was not the whole story, for it is arguable that I was ignorant of a certain *moral* fact as well—namely, that it was morally wrong to call my mother. (I will explain shortly why I say that this is “arguable.”) My non-culpable ignorance of that fact might *also* play a role in explaining why I am blameless for calling my mother. Indeed, defenders of

the Forgiving View can say that the fundamental reason why I do not deserve blame for waking my mother is that I was non-culpably ignorant of the fact that calling my mother was wrong. My ignorance of the non-moral fact that calling my mother would wake her also plays a role in explaining why I am blameless, but the role it plays is indirect: it explains the source of my moral ignorance.

The Forgiving View can provide a similar account of why non-culpable ignorance of non-moral facts excuses in many other cases as well. Explanations with the same structure can be given: the agent's non-culpable ignorance of relevant non-moral facts ensures that she is non-culpably ignorant of her obligation to act otherwise than she does. But non-culpable moral ignorance of this kind—NCFPO ignorance—constitutes an excuse for wrongdoing. Thus, the agent's non-culpable ignorance of non-moral facts excuses precisely because it gives rise to NCFPO ignorance.

While the Forgiving View does tell us something about when non-culpable non-moral ignorance excuses, it tells us nothing about when it does *not*. But this shortcoming is easily overcome if the Forgiving View is supplemented with an additional assumption, which I will refer to as the "Reduction View." The Reduction View states that non-culpable ignorance excuses *only* when that ignorance either is or gives rise to NCFPO ignorance. The Reduction View, in other words, says that all excuses due to non-culpable ignorance (whether complete or partial) "reduce" to excuses due to NCFPO ignorance. The conjunction of the Forgiving View and the Reduction View, unlike the Forgiving View alone, can make sense of why non-culpable non-moral ignorance doesn't *always* excuse.

To see how it can, take, for example, the fact that a driver does not get off the hook for intentionally hitting a pedestrian simply because the driver was non-culpably ignorant of the victim's name. Why doesn't the driver's non-culpable

non-moral ignorance get him off the hook? According to the Reduction View, the reason why is that the driver's non-culpable non-moral ignorance does not give rise to NCFPO ignorance. This is the fundamental difference between the driver case and the previous case. If I am non-culpably ignorant of the fact that my mother is sleeping, then of course I will be non-culpably ignorant of the fact that it would be wrong for me to call her. But the driver's ignorance of the fact that his victim's name is Tom does not ensure that the driver is non-culpably ignorant of the fact that intentionally harming Tom is wrong.

The explanation that the conjunction of the Forgiving View and the Reduction View provides is quite powerful. However, there is an objection that might be raised to this explanation. Think again of the case in which I call my mother, thus waking her from the nap that, unbeknownst to me, she was taking. I said above that the Forgiving View can explain why it is that I do not deserve blame for waking her. The key to that explanation was the claim that, in this case, I suffer not only from non-culpable ignorance of a certain non-moral fact (that my mother was taking a nap), but I also suffer from non-culpable ignorance of a certain *moral* fact—namely, that calling my mother would be wrong. Some philosophers, however, would resist this claim. The people that I have in mind would say that I could not have been non-culpably ignorant of the fact that calling my mother was wrong because it was *not* wrong for me to call my mother.

The easiest way to motivate this view would be to appeal to the claim that moral obligations are *epistemically-sensitive* in the sense that one cannot be obligated to perform an action that one (epistemically) permissibly does not believe that one is obligated to perform.¹⁵ If obligations were epistemically-sensitive in this way, then it would follow that if I permissibly do not believe that calling my mother would be wrong, then calling my mother would *not* be wrong.

¹⁵For more on this view, see chapter two.

Thus, the above account of when non-culpable non-moral ignorance excuses would be flawed.

If obligations are epistemically-sensitive in this way, then that would not only derail the Forgiving and Reduction Views' explanation of when non-culpable non-moral ignorance excuses, but it would also trivialize the debate between the Forgiving, Unforgiving, and Partially Forgiving Views. If obligations are epistemically-sensitive in this way, then NCFPO ignorance is *impossible*. To see why, note that if one is non-culpably ignorant of the fact that one is obligated to do X, then (a) one must not believe that one is obligated to do X, and (b) one must be epistemically permitted to lack that belief. But if one satisfies both (a) and (b), then the epistemic-sensitivity of obligations (in the above sense) would imply that one is *not* obligated to do X; hence, there is no moral fact of which one can be ignorant. Thus, if moral obligations are epistemically-sensitive in the sense under consideration, then the debate about whether NCFPO ignorance excuses has been entirely misconceived.

Luckily, there is good reason to believe that obligations are not epistemically-sensitive in this way.¹⁶ It is important to note, however, that establishing that obligations are not epistemically-sensitive in this way leaves open whether they are epistemically-sensitive in a weaker sense. This possibility needs to be addressed, which I will do in section 4.8. However, until that point, I will continue to assume that obligations are not epistemically-sensitive in a way that would cause problems for the examples that I have been using so far (and will continue to use throughout this paper).

¹⁶See chapter two.

4.5.2 The second observation

We have now been introduced to the first observation and have seen how the conjunction of the Forgiving View and the Reduction View explains it. Next, I will introduce the second observation.

While the first observation concerned non-culpable *non-moral* ignorance, the second phenomenon concerns *moral* ignorance—both culpable and non-culpable. In particular, the second observation is that increasing the amount of evidence that an agent has that his action is permissible seems to reduce the amount of blame that he deserves for acting wrongly.

The following three cases, which I will refer to as the *Comparison Cases*, illustrate this phenomenon.

Assume that it's wrong to yell at women entering abortion clinics. (If you disagree, you can change the details of the cases to reflect your moral views.) Suppose that Tom is a pro-life activist, who spends his time yelling at women entering abortion clinics. What Tom does is wrong, but is he blameworthy? It's hard to tell without further details, so consider how the following three stipulations would effect your view of Tom's action:

(Case 1) The reason why Tom believes that abortion is morally wrong and that yelling at women entering abortion clinics is permissible is that he takes his *extremely unreliable* uncle Bill to be a moral expert, and Bill says that abortion is wrong.

(Case 2) The reason why Tom believes that abortion is morally wrong and that yelling at women entering abortion clinics is permissible is that Tom takes his *relatively reliable* uncle Bill to be a moral expert.

(Case 3) Tom believes that abortion is morally wrong and that yelling at women entering abortion clinics is permissible because *everyone*

he has ever met has told him that abortion is morally wrong. In fact, *everyone in the entire world* (except abortion providers and the women who get abortions) believes that abortion is morally wrong, and Tom knows this to be the case.

When I consider my reactions to Tom's conduct in Cases 1-3, I find myself less inclined to blame Tom as I proceed through the cases. I find it hard to wrap my head around the idea that Tom in Case 2 would be just as blameworthy as Tom in Case 1. And I find it much, much harder to wrap my head around the idea that Tom in Case 3 would be just as blameworthy as Tom in Cases 2 or 1.

Let the *Comparison Judgments* be the judgments that Tom in Case 2 is less blameworthy than Tom in Case 1 and that Tom in Case 3 is less blameworthy than Tom in Cases 2 or 1. The idea that these judgments are correct—that, as we proceed through the comparison cases, the amount of blame that Tom deserves decreases—is the second observation that plays into the abductive argument.

It is important to note that accepting the Comparison Judgments does *not* require one to say, on pain of inconsistency, that Tom is blameless in any of the cases. All that is required is that one accept that Tom deserves *less* blame in Case 2 than he does in Case 1, and that he deserves less blame in Case 3 than he does in Cases 2 or 1. This is important because it shows that the Comparison Judgments are considerably weaker than the judgments that Rosen relies upon in arguing for the Forgiven View, as Rosen's judgments all concern agents who are supposed to be *entirely* blameless for acting wrongly.

If one accepts the Comparison Judgments, the natural question to ask is: what explains why the Comparison Judgments are true? Why is it that Tom in Case 3 deserves less blame than Tom in Cases 2 or 1? And why does Tom in Case 2 deserve less blame than Tom in Case 1?

What I will now argue is that the Forgiving View provides a compelling explanation of the Comparison Judgments, at least if it is supplemented with a closely related view.

Let's start with the judgment that Tom in Case 3 deserves less blame than Tom in Cases 2 and 1. The Forgiving View explains this judgment as follows: Tom in Case 3 is non-culpably ignorant of his obligation not to yell at women entering abortion clinics. But, in Cases 1 and 2, he is arguably *culpably* ignorant of his obligation. Thus, the Forgiving View will imply that Tom in Case 3 is completely excused for acting wrongly; however, the Forgiving View does not imply that Tom in Cases 1 and 2 is excused on account of his ignorance. Presumably, Tom in Cases 1 and 2 does not have any other complete excuse for acting wrongly (if he did, what would that be?). Thus, Tom in Cases 1 and 2 does deserve at least *some* blame for acting wrongly, while Tom in Case 3 does not.

Now, one objection to this explanation must be addressed at the outset. An opponent of the Forgiving View might argue at this point that Tom in Case 3 is in fact *culpably* ignorant of his obligation. The only reason to think that Tom in Case 3 is *non-culpably* ignorant of his obligation not to yell at women entering abortion clinics is that he has a great deal of testimonial evidence that yelling at women entering abortion clinics is permissible. But why think that moral testimony can provide justification? Don't some philosophers deny that it can?¹⁷ And even if one grants that moral testimony can provide justification, why think that in this case Tom's testimonial evidence is sufficiently strong that it outweighs any other moral evidence that he has?

I have several replies to offer here. First, as far as I can see, there is no com-

¹⁷Robert Paul Wolff, for example. See Wolff [47], 13.

elling argument that moral testimony cannot provide justification.¹⁸ It may be true that it is harder for moral testimony to provide justification than it is for other types of testimony to provide justification (perhaps because, as Markovits suggests, it is rare to be in a position to reliably determine who is a moral expert about a particular issue without being in a position to form a belief about that subject matter ourselves), but that does not mean that it is impossible for it to do so.¹⁹ Second, even if it is controversial whether moral testimony can provide justification, it is still the case that many people believe that it does.²⁰ That is all that is needed for the Comparison Cases to provide a significant amount of dialectical leverage. Third, even if one thinks that someone like Tom has strong non-testimonial evidence in favor of the wrongness of his action, it is still possible that a huge, consistent, and compelling body of testimony can undermine that evidence. There are standard examples of testimonial evidence overcoming the justification provided by basic arithmetical reasoning, sense perception, and more.²¹ If these types of evidence can be overcome by testimonial evidence, I see no reason to think that Tom's evidence against the permissibility of his action can't be overcome by the testimony to which he is exposed.

One further point to make here is that even if one rejects the idea that moral testimony can provide justification, it is likely still possible to generate another

¹⁸Even those who think that there are moral reasons not to rely on moral testimony typically admit that moral testimony can provide knowledge or justification. See, for example, Hills [16] and Nickel [27]. Wolff is a rare example of a philosopher who has denied in print that moral testimony provides justification (Wolff [47], 13). But his argument against moral testimony relies on the idea that mathematical testimony never provides justification except insofar as it prompts the recipient to think through a chain of reasoning that is itself a source of justification. As has been pointed out before by others, this view of mathematical testimony is highly implausible, since it rules out the possibility that someone who is not a professional mathematician may come to justifiably believe on the basis of a mathematician's say-so that Fermat's Last Theorem is true.

¹⁹Markovits [21], f.n. 38.

²⁰For a start, see Hills [16], Jones [18], Nickel [27], Sliwa [38].

²¹See, for example, Christensen [7]'s mental math example (757), or Elga [10]'s horse racing example (486).

series of cases that can serve the same function as the Comparison Cases. Whatever one thinks that moral evidence is, one can imagine a series of cases constructed using the following schema:

Let action A be wrong.

(Case 1) Person P thinks that A is permissible, but P doesn't have much evidence that it is.

(Case 2) Person P thinks that A is permissible, and P has some, but not very strong, evidence that it is.

(Case 3) Person P thinks that A is permissible, and P has a *massive* body of evidence that points to the permissibility of A.

I conjecture that for any story about moral evidence, a series of cases can be constructed along these lines, and in any such series of cases, it will be overwhelmingly plausible that P in Case 3 deserves less blame than P in Cases 1 or 2 and that P in Case 2 deserves less blame than P in Case 1. If this conjecture is true, then these cases can be substituted for the Comparison Cases and my argument run in much the same way. Thus, to derail my argument it is not enough to attack the view that moral testimony provides justification.

So far, I have argued that the Forgiving View explains why Tom in Case 3 deserves less blame than Tom in Case 1 or Case 2. But this is only one of the two Comparison Judgments. How does the Forgiving View explain why Tom in Case 2 deserves less blame than Tom in Case 1? The answer is that it cannot *by itself* explain this, since Tom in Case 2 is *culpably* ignorant of his obligation, which means that the Forgiving View tells us nothing about him. Instead, what is necessary is to combine the Forgiving View with a natural *extension* of that view. Given that the Forgiving View says that ignorance of one's obligations that rises

to the level of non-culpability constitutes a complete excuse for wrongdoing, it is natural to believe that ignorance of one's obligations that meets a lower standard (that is culpable, but still the result of *some* kind of reason to believe that one was acting permissibly) constitutes a partial excuse whose "strength" is proportional to the strength of the reasons for the ignorance. Call this extended view the Forgiven Extension.

Appealing to the Forgiven Extension provides a natural explanation of why Tom in Case 2 is less blameworthy than Tom in Case 1. Tom in Case 2, while culpably ignorant, at least has better reasons to believe that his action is permissible than Tom in Case 1. Thus, Tom in both cases has a *partial* excuse for wrongdoing. But the strength of those reasons, according to the Forgiven Extension, will vary according to the reasons each person has for his ignorance. Since Tom in Case 2 has better reasons for his ignorance than Tom in Case 1 does, Tom in Case 2 is less blameworthy than Tom in Case 1.

4.5.3 Formulating the abductive argument

Now that we have seen how the Forgiven View, when combined with the Reduction View and the Forgiven Extension, can explain the first and the second observations, we are in a position to outline the abductive argument. The abductive argument is quite simple. It says that the Forgiven View, when combined with these additional views, doesn't just explain the first and second observations; this "package" of views (which I will refer to as the *Forgiven Package*) provides the *best* explanation of these observations and therefore should be accepted:

1. The Forgiven Package (that is, the Forgiven View plus the Reduction View and the Forgiven Extension) provides the best explanation of both

- (a) when non-culpable non-moral ignorance excuses (and when it does not) as well as (b) the Comparison Judgments.
- 2. If 1., then the Forgiving Package is probably true.
- C. Therefore, the Forgiving Package is probably true.

This is the abductive argument in favor of the Forgiving Package. It is, I believe, the best argument for the Forgiving View. If the argument succeeds, then it provides strong, though defeasible, support for the Forgiving View.

Over the course of the next several sections, I will explain why one should think that the Forgiving View really does provide the best explanation of the two observations. To make my case, I will explore how defenders of the Unforgiving and Partially Forgiving Views might offer rival explanations. This will occupy the remainder of the paper.

4.6 The Quality of Will objection

In order to make the abductive argument stick, it is crucial to establish that the Forgiving View provides the best explanation of our two observations. And in order to establish that, it must be shown that no alternative view can provide or be combined with an explanation of these observations that is at least as good as the explanation provided by the Forgiving View. Our question, then, is whether such an explanation can be provided. I will start by considering how defenders of the Unforgiving View might try to explain the facts about when non-culpable non-moral excuses and when it does not.

Note that if such an explanation can be provided, it will not follow *directly* from the Unforgiving View; the Unforgiving View will have to be *combined* with such an explanation. The Forgiving View was able to explain when non-

culpable non-moral ignorance excuses by linking this type of ignorance with a type that the Forgiving View already held to be exculpatory. However, the Unforgiving View says nothing about when (or even whether) non-culpable ignorance *does* excuse; it says only that a certain type of non-culpable ignorance *does not* excuse. Thus, the Unforgiving View by itself cannot possibly explain when and why non-culpable non-moral ignorance excuses. That said, the possibility remains that an explanation can be wheeled in to supplement the Unforgiving View.

One way in which defenders of the Unforgiving View might try to explain when non-culpable non-moral ignorance excuses would be to appeal to the Quality of Will View often associated with Strawson: that is, the view that whether an agent deserves blame for acting wrongly is determined by the moral quality of her will.²² Appealing to the Quality of Will View in this way would be a natural move, since two of the most visible defenders of the Unforgiving View, Elizabeth Harman and Matthew Talbert, have appealed to the Quality of Will View to explain *why* non-culpable non-moral ignorance sometimes excuses.²³ The question is whether their explanations hold water, and whether those explanations can be generalized to explain *when* non-culpable non-moral ignorance excuses and when it does not.

The basic form of Harman's and Talbert's explanations is the same. First, each starts with the Quality of Will View. This view on its own does not imply, or even support, the Unforgiving View. In fact, the Quality of Will view is compatible with the Forgiving View, as one might maintain that any agent who suffers from NCFPO ignorance necessarily has a good will.²⁴ To support the Unforgiv-

²²Strawson [39].

²³Harman [14], Talbert [42].

²⁴This important point is missed by Rosen, who assumes in his discussion of the Quality of Will View that that view is incompatible with the Forgiving View. See Rosen [33], 72-73.

ing View, the Quality of Will View needs to be supplemented. Thus, the second ingredient of Harman's and Talbert's explanations: an account of what makes an agent's will good, bad, or objectionably indifferent to moral considerations. Combining this account with the Quality of Will View, the thought goes, will both support the Unforgiving View and tell us why non-culpable non-moral ignorance sometimes excuses.

What I will do now is explain how Harman and Talbert implement this strategy. I will also argue that their explanations of *why* non-culpable non-moral ignorance sometimes excuses do not generalize to provide explanations of *when* non-culpable non-moral ignorance excuses.

4.6.1 Talbert

I will begin by quoting an important passage from Talbert:

Even if a wrongdoer is ignorant of the fact that her behavior is wrong, and even if this ignorance is not her fault, her actions may still express the contemptuous judgment that certain others do not merit consideration, that their interests do not matter, and that their objections can be overlooked. If one is injured by a wrongdoer who is moved by such judgments, then the attitudes and responses involved in moral blame are reasonable regardless of what the wrongdoer thinks about the moral status of her behavior.²⁵

This short paragraph presents the key idea behind Talbert's approach to both defending the Unforgiving View and explaining why non-culpable non-moral ignorance sometimes excuses. As I indicated above, Talbert, like Harman, ac-

²⁵Talbert [42], 234.

cepts the Quality of Will View and combines it with a view about what determines the quality of an agent's will. In this paragraph, Talbert provides the latter. According to Talbert, whether an agent's will is good, bad, or objectionably indifferent to morality is determined by whether her action expresses "an objectionable moral judgment". A few examples will help to illustrate this view.

Start by imagining that I call my mother, thus waking her from the nap that she was taking; this time, however, imagine that I knew full well beforehand that she would be napping. (Suppose that I call anyway because I am mad that she forgot my birthday, and I'm trying to get back at her.) Obviously, in this case I would deserve blame for calling my mother. Talbert would agree, and he would explain this fact as follows. Whether I deserve blame for calling my mother is determined by whether my action expressed "an objectionable moral judgment." In this case, my action does express an objectionable moral judgment—namely, that my mother's interest in sleeping does not matter. This judgment, Talbert would say, provides sufficient grounds for blame.

Contrast this case with one in which I call my mother, thus waking her from the nap that I had no reason to think she was taking. If I had no reason to think that my mother was sleeping, I obviously do not deserve blame. Why don't I? Talbert would answer that I do not deserve blame because my action does *not* express an objectionable moral judgment.

These two cases illustrate how Talbert explains why non-culpable non-moral ignorance sometimes excuses. Before we explore whether his view can provide a fully general account of when non-culpable non-moral ignorance excuses, it is worth pausing to recount how Talbert tries to apply his view to moral ignorance.

Talbert denies that NCFPO ignorance constitutes an excuse for wrongdoing. The key to understanding why he denies this is the view that I have just

illustrated—namely, the view that what determines the quality of an agent’s will (and so whether they deserve blame) is whether the agent’s action expressed “an objectionable moral judgment.” Above, I quoted Talbert’s claim that “[even] if a wrongdoer is ignorant of the fact that her behavior is wrong, and even if this ignorance is not her fault, her actions may still express the contemptuous judgment that certain others do not merit consideration, that their interests do not matter, and that their objections can be overlooked.”²⁶ In other words, Talbert maintains that NCFPO ignorance does not constitute an excuse for wrongdoing because NCFPO ignorance does not prevent an agent’s action from expressing objectionable moral judgments (and so does not prevent an agent’s will from being bad). So, for example, Talbert would say that Rosen’s Hittite slaveholder does deserve blame for holding slaves, even if the slaveholder was non-culpably ignorant of the wrongness of slavery, simply because his actions expressed the objectionable judgment that his slaves’ interests did not matter.

We have now seen the outlines of Talbert’s account of non-culpable non-moral ignorance, and we have seen why he denies that NCFPO ignorance constitutes an excuse for wrongdoing. It is now time to evaluate his approach and to consider whether his approach might be able to explain when non-culpable non-moral ignorance excuses and when it does not.

In my view, Talbert’s approach to non-culpable ignorance is plausible on its face, but problems begin to surface when one digs deeper into Talbert’s notion of “an objectionable moral judgment.” Talbert gives us some examples of what he calls “objectionable moral judgments” in the passage quoted above: “the contemptuous judgment that certain others do not merit consideration, that their interests do not matter, and that their objections can be overlooked.”²⁷ How-

²⁶Talbert [42], 234.

²⁷Talbert [42], 234.

ever, he never indicates whether these examples are supposed to be exhaustive. But whether they are intended to be exhaustive or not, Talbert faces a problem.

If the examples are supposed to be exhaustive, then Talbert's view is false. An aggressor can deserve blame even if he judges that his victim deserves consideration, that his victim's interests matter, and that his victim's objections cannot be overlooked. It might simply be that the aggressor judges that his victim's interests and objections, while important, are overruled by something else. Imagine, for example, someone who believes for bad reasons that he must assault the next person he sees in order to stop an imminent alien invasion of the Earth. Such a person is blameworthy even if he believes that his victim's interests and objections matter. Thus, it is not true that only agents whose actions express "the contemptuous judgment that certain others do not merit consideration, that their interests do not matter, and that their objections can be overlooked" deserve blame.

Now, one might respond that my objection to Talbert's view relies on too narrow an understanding of Talbert's examples. Perhaps Talbert is not claiming that the only "objectionable moral judgments" that can ground blame are those he lists. But if this is Talbert's response, then my complaint is that Talbert gives us no indication of how to generalize his examples of "objectionable moral judgments," which leaves us with, at best, a very incomplete picture of blame. Indeed, it leaves us with a picture on which it is unclear why non-culpable non-moral ignorance *ever* constitutes an excuse for wrongdoing. Consider, for example, the following question: doesn't *every* wrong action express an objectionable judgment—namely, that the agent had sufficient reason to perform the action? It's not clear on what grounds Talbert could deny this. But it is absolutely essential to his project that he does; otherwise, blameless wrongdoing will be *impossi-*

ble on his view. This is clearly not a conclusion he would embrace, but without further development of his view, it is unclear how he can avoid it.

Talbert's approach does not provide a satisfying alternative to the Forgiving View's explanation of when non-culpable non-moral ignorance excuses. If defenders of the Unforgiving View who accept the Quality of Will View are to formulate a convincing objection to the abductive argument, they will need a far more systematic account of when an agent's will counts as good.

4.6.2 Harman

Harman's approach to non-culpable non-moral ignorance promises to provide a more systematic account of what determines whether an agent's will is good or bad, which at first looks to fill the hole left in Talbert's account. As we shall see, though, Harman's view faces problems of its own.

Harman's view is difficult to state, as a careless formulation makes it sound as if it is a restricted form of the Forgiving View. The reason why is that Harman says that she believes that "a person [cannot be] blameworthy for a wrong action he believed to be permissible but is blameless for his false belief."²⁸ The key to understanding Harman's view, though, is to recognize that Harman is working with a very different notion of blameless or non-culpable belief/ignorance than I am.

The notion of non-culpable ignorance or belief that I have been working with so far is a *purely epistemic* notion—whether one's ignorance or belief is non-culpable is determined by one's evidence. Harman's notion of non-culpable ignorance or false belief, on the other hand, is not purely epistemic; hers is moral. She provides a helpful summary of her notion of non-culpable ignorance in the

²⁸Harman [14], 459.

following passage:

... one can be blameworthy for having false moral beliefs although one has not been in any way irresponsible in the management of one's opinion ... Rather, one has violated some moral norms that apply to *beliefs* themselves, not to the management of one's beliefs. The view holds that we morally ought to believe the moral truth.²⁹

The first sentence here is illuminating: it implies that an agent's false belief that is non-culpable in my sense can be culpable in Harman's sense. In light of the significant difference between Harman's notion of non-culpable false belief and my own notion of non-culpable ignorance, I propose to use the notation "non-culpable_{harman}" when referring to her notion.

With this notation in hand, it is possible to state the difference between Harman's view and the Forgiving View quite simply. The Forgiving View says that a certain kind of non-culpable moral ignorance—NCFPO ignorance—excuses. Harman, on the other hand, denies that NCFPO ignorance excuses, though she does believe that non-culpable_{harman} false moral beliefs about one's own obligations excuse.

What, then, is it for a false moral belief to be non-culpable_{harman}? This is where things start to get a bit tricky. Harman says that we are morally obligated to believe the moral truth, which might lead one to believe that, according to Harman's view, *any* false moral belief is blameworthy. But this is not correct. While we are morally obligated to believe the moral truth, Harman does say that some false moral beliefs are blameless. She writes,

... beliefs (and failures to believe) are blameworthy if they involve inadequately caring about what is morally significant. Believing a

²⁹Harman [14], 459.

certain kind of behavior is wrong on the basis of a certain consideration is a way of caring about that consideration. Some failures to believe moral truths relevant to one's actions are not blameworthy. For example, if one blamelessly falsely believes a non-moral claim, and this leads to one's false moral belief, then one's false moral belief does not involve inadequately caring about what is morally significant.³⁰

To understand this passage, it is helpful to consider an example. Suppose I call my mother to chat, thus waking her from the nap that, unbeknownst to me, she was taking. I thought that it was permissible for me to call, since I thought that she would be awake. But I was wrong. Even so, I am obviously not to blame for calling my mother. Harman would agree. She would say that I could be blamed only if I was culpable_{harman} for believing that calling was permissible. However, since my belief that calling was permissible was the result of my blameless false belief that my mother would be awake (a blameless false belief about a non-moral claim), my belief that calling was permissible was not culpable_{harman}. Thus, I am not a proper target for blame.

Harman's treatment of this case seems plausible enough, but does her account generalize to provide an account of when non-culpable non-moral ignorance excuses? The answer, I believe, is that it does not, at least not without serious modifications. As I see it, there are two problems with Harman's account.

First, while Harman gives us a sufficient condition for a false moral belief to count as non-culpable_{harman}, which gives us a sufficient condition (when combined with the rest of her view) for wrongdoing to be blameless, Harman never

³⁰Harman [14], 460.

gives us any indication of whether this condition is also supposed to be *necessary*. Perhaps we are supposed to read this as necessary as well (in which case I believe that there are still problems—see below). But if we are not, then she simply hasn't explained when non-culpable non-moral ignorance excuses *and when it does not*. Thus, she has not shown that defenders of the Unforgiving View have the resources necessary to challenge the abductive argument.

The second problem with Harman's account is that she is simply wrong when she says that "if one blamelessly falsely believes a non-moral claim, and this leads to one's false moral belief, then one's false moral belief does not involve inadequately caring about what is morally significant [and so one is not blameworthy]."³¹ Suppose that I blamelessly falsely believe that a certain person's name is John, and that leads me to believe that it is permissible to kill John because I have the strange belief that everyone named John can permissibly be killed. Harman's claim implies that I would not be blameworthy if I killed John, but clearly this is false.

These problems lead me to believe that Harman's spin on the Quality of Will View does not provide the resources to explain when non-culpable non-moral ignorance excuses. Thus, defenders of the Unforgiving View still stand in need of a response to the abductive argument.

4.6.3 An alternative

Given the failure of Talbert's and Harman's accounts of non-culpable non-moral ignorance, one might wonder whether it *is* possible to explain when non-culpable non-moral ignorance excuses (and when it does not) by appealing to the Quality of Will View. I think that it probably is possible to do this; I will now

³¹Harman [14], 460.

sketch how this might be done.

The most promising approach to using the Quality of Will View to explain when non-culpable non-moral ignorance excuses and when it does not takes off from a claim that Harman makes in passing. The claim is the following: “believing a certain kind of behavior is wrong on the basis of a certain consideration is a way of caring about that consideration.”³² In context, this claim seems to be offered to help explain how a belief can fail to be culpable_{harman}, and thus can render an agent’s wrongdoing blameless. How exactly this explanation is supposed to go is complicated; however, this needn’t detain us since there is good reason to believe that Harman’s claim, as stated, is false. What one believes and what one cares about are separate matters. Indeed, it seems entirely possible for an agent to believe that a certain kind of behavior is wrong on the basis of a certain consideration and yet still fail to *care* about avoiding such actions. Harman’s claim, however, seems to imply that this is impossible.

Even so, I do believe that the rough idea that Harman is gesturing towards here—that what determines whether an agent is blameworthy for wrongdoing is whether she is appropriately responsive to the *wrong-making features* of her action—is appealing. Where Harman goes astray, I believe, is in spelling out this idea more precisely.

A better way to spell out this basic idea, I believe, is the following:

(The Improved Quality of Will View) an agent is blameworthy for performing a wrong action A if and only if, for some non-moral feature F of A that made A wrong, the agent (a) had sufficient reason to believe that A had F, and (b) the agent intended to perform action A anyway.³³

³²Harman [14], 460.

³³Note that, on this view, an agent can be blameworthy even if he did not realize that possess-

The Improved Quality of Will View captures Harman's basic idea, but it avoids conflating what an agent believes and what an agent cares about.

One significant attraction of the Improved Quality of Will View is that it provides the right verdicts about all of the cases involving non-moral ignorance that I have discussed so far. For example, it explains why, if I was not aware that my mother was napping, I am not blameworthy for calling her and thereby waking her. It is natural to think that calling my mother in this situation would be wrong because it would annoy her. (If you think that calling would be wrong for some other reason, simply substitute that reason in the following explanation. The explanation will still work.) Thus, according to the Improved Quality of Will View, in order to be blameworthy for wrongly waking my mother, it would be necessary that (a) I had sufficient reason to believe that calling would wake my mother, and (b) I called anyway. But (a) is false. I did not have sufficient reason to believe that calling would wake my mother. Therefore, I do not deserve blame for calling.

The view also explains why a driver does not get off the hook for intentionally hitting a pedestrian simply because he was non-culpably ignorant of his victim's name. Hitting the victim was wrong because it harmed the victim. (Again, if you think that the action was wrong for some other reason, simply substitute that reason in the following explanation.) The Improved Quality of Will View therefore says that the driver is blameworthy iff (a) he had sufficient reason to believe that hitting the victim would harm the victim, and (b) he did it anyway. Both (a) and (b) are true, so the driver deserves blame.

As currently stated, the view seems promising. That said, it is likely that it needs further work. For one thing, the view might be too simple to handle cases

ing F made his action wrong. All that is necessary is that the agent intentionally performed an action with a feature that, it turns out, did make his action wrong. Whether he recognized that wrong-making feature as wrong-making does not matter.

in which, for some wrong-making feature *F*, an agent has reason to believe that her action *might* have *F*, though she lacks sufficient reason to believe that it *does*. Some cases that fit this mould seem to be cases in which blame is appropriate. (Suppose, for example, that I have reason to believe that calling *might* wake my mother, but I lack sufficient reason to believe that it *will*. If I decide to call and this does end up waking her, blame seems appropriate. The Improved Quality of Will View, however, cannot explain why it is.) That said, accommodating these cases seems to be detail work, which there is no reason to believe cannot be completed. Thus, it appears that a view much like the Improved Quality of Will View probably can explain when non-culpable non-moral ignorance excuses and when it does not.

Now, even if a view in the neighborhood of the Improved Quality of Will View can explain when non-culpable non-moral ignorance excuses and when it does not, that would not by itself show that the abductive argument fails. For the abductive argument doesn't say only that the Forgiving Package provides the best explanation of when non-culpable non-moral ignorance excuses and when it does not. The abductive argument says that the Forgiving View constitutes the best explanation of that *and* the Comparison Judgments. The question, then, becomes: if we adopt the Improved Quality of Will View, is it possible to explain the Comparison Judgments?

4.6.4 Can defenders of the Quality of Will View explain the Comparison Judgments?

Unfortunately for defenders of the Unforgiving View, the Improved Quality of Will View is not up to the challenge of explaining the Comparison Judgments.

Furthermore, there is reason to believe that the view cannot be supplemented in a way that will help it to meet this challenge.

First, consider what the Improved Quality of Will View implies about the Comparison Cases. To figure out what it implies, we need to know why Tom's actions in the Comparison Cases are wrong. Presumably, yelling at women entering abortion clinics is wrong because it upsets them. (As usual, if you think that yelling at women entering abortion clinics is wrong for some other reason, please insert that reason in what follows.) Now, Tom in Cases 1-3 knows that yelling at these women will upset them, and he does so anyway. Thus, according to the Improved Quality of Will View, he is blameworthy, since he intentionally performed an action with a feature that, it turns out, made his action wrong.

One problem for the Improved Quality of Will View is that this is *all* the view implies about the Comparison Cases. It provides us with no account of why Tom in Case 2 is less blameworthy than Tom in Case 1, and it provides us with no account of why Tom in Case 3 is less blameworthy than Tom in Cases 1 or 2. Now, this alone is not necessarily a fatal problem. As we saw earlier, even the Forgiving View needed to be supplemented in order to explain the Comparison Judgments. The problem is that there is little chance that even a supplemented Improved Quality of Will View will be able to provide the necessary explanation.

Here's why. The Improved Quality of Will View is a view about the factors that determine whether an agent is blameworthy. But to explain the Comparison Judgments, a defender of the Improved Quality of Will View would need to appeal to an additional view about what determines the *amount* of blame that an agent deserves—a view that implies the Comparison Judgments. Presumably,

the factors that determine whether an agent deserves blame are closely related to the factors that determine how much blame an agent deserves. But the factors that the Improved Quality of Will View says determine whether an agent deserves blame do not vary between the Comparison Cases. (In all of the cases, Tom's action is wrong for the same reason, his non-moral beliefs and evidence do not vary, and he intends to perform the same action.) Thus, it is hard to imagine a view about what determines the amount of blame an agent deserves that both (a) is consistent with the Improved Quality of Will View and (b) implies the Comparison Judgments. It seems, then, that defenders of the Improved Quality of Will View will be unable to explain the Comparison Judgments.

If I am right that the Improved Quality of Will View cannot be supplemented in a way that will yield an explanation of the Comparison Judgments, then the view will not help defenders of the Unforgiving View to respond to the abductive argument. Defenders of the Unforgiving View, then, will need to look elsewhere.

One possible response for defenders of the Unforgiving View would be to move to a different type of Quality of Will View. The obvious problem with this approach is that defenders of the Quality of Will View would then need to come up with a different explanation of when non-culpable non-moral ignorance excuses and when it does not. However, it is not obvious how to do so.

Another response would be to deny the Comparison Judgments. To someone who doesn't find the Comparison Judgments at all plausible, there is not much that can be said. But there is a way that defenders of the Unforgiving View might try to *motivate* denying the Comparison Judgments. It might be suggested that Tom deserves just as much blame in Cases 2 and 3 as he does in Case 1, although as we proceed through the cases his *moral luck* becomes worse.

We are already forced to accept the existence of moral luck, someone who pursues this line of defense would suggest. So why not write off Tom as another example of bad moral luck?³⁴

I have two concerns about this sort of response. First, as Nagel has pointed out, the idea of moral luck is *troubling*, since the existence of moral luck seems to be inconsistent with an intuitive control principle.³⁵ To appeal, then, to moral luck as if doing so is philosophically unproblematic would be quite odd. Second, even if there are cases of bad moral luck, it is not clear that Tom in the Comparison Cases is one of them. After all, we don't want to say that *every* agent who acts wrongly because she is in a bad epistemic situation is an example of bad moral luck (and so deserves blame). It would be absurd to appeal to moral luck in order to defend the claim that I deserve blame if my phone call wakes my mother from a nap that I had no reason to believe she would be taking. There are very real limitations, then, to this strategy. But someone who wants to appeal to the notion of bad moral luck in order to motivate denying the Comparison Judgments needs to explain *what those limitations are*. Until we hear a good account of why it doesn't make sense to treat *all* of the non-culpably non-morally ignorant as victims of bad moral luck, we should be suspicious of this response.

³⁴Compare Harman [14]'s treatment of her examples of NCFPO ignorance: "I see their cases as instances of constitutive moral bad luck: they are unlucky to have found themselves in circumstances that have caused them to be (in some respects) morally bad people" (462).

³⁵Nagel [26].

4.7 The second response on behalf of the Unforgiving View: epistemic sensitivity, again

Earlier, I mentioned that some might respond to the abductive argument by appealing to the view that obligations are epistemically-sensitive in the sense that any agent who permissibly does not believe that she is obligated to do X is not obligated to do X. Call this view (Epistemic-Sensitivity). (Epistemic-Sensitivity) appears to throw a wrench in the Forgiven View's explanation of when non-culpable non-moral ignorance excuses. That explanation relied on the idea that non-culpable non-moral ignorance often gives rise to NCFPO ignorance, which, according to the Forgiven View, excuses. But if obligations are epistemically-sensitive, NCFPO ignorance is impossible. Thus, the Forgiven View's explanation of when non-culpable non-moral ignorance excuses would fail if (Epistemic-Sensitivity) were true.

In section 4.5.1, I argued that this way of responding to the abductive argument is not promising, as there are good reasons to reject (Epistemic-Sensitivity).³⁶ Furthermore, anyone who accepts (Epistemic-Sensitivity) will have to accept the same conclusions about which agents are blameworthy as defenders of the Forgiven View do.

There is, however, a different kind of epistemic sensitivity thesis to which defenders of the Unforgiving View might appeal, which avoids these difficulties. The thesis that I have in mind says that what is required in order for one's action to be wrong is that one have a sufficient amount of information about the non-moral features of one's action that make the action wrong. (This idea should sound familiar after having reading about the Improved Quality of Will View.) One might say, for example, that an action A can be wrong for an agent S to per-

³⁶I present these reasons in chapter two.

form only if, for one of the non-moral features of A that grounds A's wrongness, S is epistemically required to believe that A has that feature (though S needn't think of that feature as wrong-making). Call this the *weak epistemic-sensitivity view*.³⁷

An example will help to illustrate the weak epistemic-sensitivity view. Once again, consider a case in which I call my mother to chat, thus waking her from a nap that I had no reason to think she was taking. Is my action wrong? The weak epistemic-sensitivity view says that if my action is wrong, then, for some non-moral feature of my action that grounds its wrongness, I must be epistemically required to believe that my action has that feature. But if calling my mother is wrong, then it is wrong because waking her would annoy her. (As usual, please substitute a different wrong-making property if you think that necessary.) The weak epistemic-sensitivity view then implies that if calling my mother is wrong, I must be epistemically required to believe that calling my mother will annoy her. But I am *not* epistemically required to believe that calling my mother will annoy her, since I have no reason to believe that she is sleeping. Thus, calling my mother is not wrong, according to the weak epistemic-sensitivity view.

What makes the weak epistemic-sensitivity view interesting is that it does not entail that NCFPO ignorance is impossible, which means that it is not refuted by my argument for the possibility of NCPFO ignorance. All that the weak epistemic-sensitivity view entails is that a *specific kind* of non-culpable *non-moral* ignorance is impossible: namely, non-culpable ignorance of all of the non-moral features of one's (wrong) action that ground its wrongness.

The weak epistemic-sensitivity view provides defenders of the Unforgiving

³⁷Notice that the weak epistemic-sensitivity view allows for the possibility that an action can be wrong for a person to perform even if she is non-culpably ignorant of *some* of the non-moral features that make her action wrong. What the view rules out is that an action can be wrong for a person if she is non-culpably ignorant of *all* of the non-moral features that ground the wrongness of her action.

View with an interesting response to the abductive argument. The response is to deny the first observation that plays into the abductive argument—that non-culpable non-moral ignorance sometimes excuses. A defender of the weak epistemic-sensitivity view can maintain that non-culpable non-moral ignorance *never* excuses; however, they can say that it *appears* to excuse precisely when an agent’s non-moral ignorance prevents her action from actually being wrong. For example, we have already seen that a defender of the weak epistemic-sensitivity view would say that if I call my mother to chat, thus waking her from the nap that I had no reason to believe she was taking, my action was not wrong. Thus, this case need not be held up as an example in which non-culpable non-moral ignorance *excuses*. Something similar can be said about the case in which I slap my sunburned friend on the back, causing him excruciating pain. If it is wrong for me to greet my friend with a warm slap on the back, that would be because it causes my friend pain. But I am not epistemically required to believe that slapping my friend will cause him pain. Thus, the weak epistemic-sensitivity view entails that it is not wrong for me to slap him on the back. Again, it is possible to explain why I am not blameworthy without appealing to the idea that non-culpable non-moral ignorance excuses me.

While the weak epistemic-sensitivity view does provide defenders of the Unforgiving View with a way to argue that non-culpable non-moral ignorance never excuses, the view does not provide a compelling response to the abductive argument. For the only reason to accept the weak epistemic-sensitivity view is that the weak epistemic-sensitivity view can explain the appearance that non-culpable non-moral ignorance excuses. But the fact that it can explain this does not by itself show that it is a serious rival to the Forgiven View. For the Forgiven View doesn’t just explain why non-culpable non-moral ignorance ex-

cuses (or, to put this more neutrally, it doesn't just explain why non-culpable non-moral ignorance *appears* to excuse); the Forgiving View also explains the Comparison Judgments. If the weak epistemic-sensitivity view is to ground a response to the abductive argument, it must also be possible for this view to provide or be combined with an explanation of the Comparison Judgments.

Unfortunately for defenders of the Unforgiving View, the weak epistemic-sensitivity view cannot explain the Comparison Judgments. This much is clear, as the view is designed specifically so that it says *only* that certain kinds of non-culpable *non-moral* ignorance are impossible. It says nothing about cases of moral ignorance at all.

Again, though, the fact that a view does not by itself explain both of the observations that play into the abductive argument does not mean that the view fails. It is important to consider whether the view could be supplemented to explain the Comparison Judgments. Perhaps it can be, but whatever view is wheeled in to help explain the Comparison Judgments, it will have little to do with the weak epistemic-sensitivity view. It will be entirely independent. And it is, at this moment, unclear what this additional view could be.

4.8 The Partially Forgiving View

Perhaps the problem with the previous two ways of proceeding were that they were too committed to the Unforgiving View. Perhaps the thing to do is to admit that NCFPO ignorance does constitute *some* kind of excuse for wrongdoing, but to avoid committing to the full-strength Forgiving View. The Partially Forgiving View provides an attractive option for someone considering this approach.

The Partially Forgiving View's attractions become even more pronounced when one considers how the view explains the Comparison Judgments. The

Partially Forgiving View provides an easy explanation of why Tom in Case 3 is less blameworthy than Tom in Cases 1 and 2: Tom in Case 3 is non-culpably ignorant of his obligation not to yell at women entering abortion clinics; thus, he is partially excused for his wrongdoing. The trickier explanation to provide is why Tom in Case 2 is less blameworthy than Tom in Case 1, but here the Partially Forgiving View can appeal to the Forgiving Extension, just as the Forgiving View did; defenders of the Partially Forgiving View can say that ignorance of one's obligations that meets a lower standard (that is culpable, but still the result of *some* kind of reason to believe that one was acting permissibly) constitutes a partial excuse whose "strength" is proportional to the strength of the reasons for the ignorance. This allows one to say that Tom in Case 2 deserves less blame than Tom in Case 1, on account of his stronger evidence that he is acting permissibly.

The Partially Forgiving View's attractions do not end there. This view can not only explain the Comparison Judgments, but it appears to be able to accommodate some of the intuitions that motivate defenders of the Unforgiving View. Many people are attracted to the Unforgiving View, I believe, because they find it very implausible to say, for example, that Rosen's Hittite slaveholder deserves no blame at all for holding slaves. They feel that the slaveholder is a proper target of anger and indignation. Notice, though, that this can be true even if the Partially Forgiving View is correct. For the Partially Forgiving View can say that although the amount of blame that the slaveholder deserves is *reduced* on account of his NCFPO ignorance, he is not entirely blameless. His excuse, in other words, is *merely* partial.

The Partially Forgiving View, then, seems to represent an appealing, ecumenical alternative to the Unforgiving and Forgiving Views—an alternative that

has so far been overlooked in the literature on moral ignorance. But for the Partially Forgiving View to be taken seriously, it needs to show not only that it can explain the Comparison Judgments, but that it can either provide or be combined with an explanation of when non-culpable non-moral ignorance excuses (and when it does not). There is reason to believe, however, that the Partially Forgiving View cannot explain when non-culpable non-moral ignorance excuses. And it is not clear what additional view could be wheeled in to fill this gap.

The natural way to use the Partially Forgiving View to explain when non-culpable non-moral ignorance excuses would be to mimic the strategy used by the Forgiving View. First, suggest that non-culpable non-moral ignorance excuses when non-culpable non-moral ignorance gives rise to non-culpable *moral* ignorance. Second, appeal to the Reduction View to show that when non-culpable non-moral ignorance does not give rise to non-culpable moral ignorance, it does not excuse. The problem with this two-part strategy, however, is that it only results in an explanation of when non-culpable non-moral ignorance *partially* excuses (and when it does not). But non-culpable non-moral ignorance does not only provide a *partial* excuse when, say, I slap my sunburnt friend on the back because I did not realize that he was sunburnt. In that case, my non-culpable non-moral ignorance provides a *complete* excuse for my action. The Partially Forgiving View, however, cannot explain why this is.

This is a serious problem for the Partially Forgiving View. Without supplementing the Partially Forgiving View's explanation of when non-culpable non-moral ignorance excuses, the view cannot defend itself against the abductive argument. As far as I can see, however, there is no obvious candidate to do this supplementing. What would be needed is a view about what determines the

amount of blame that an agent deserves for acting wrongly—a view that entails that in the standard cases of exculpation by non-culpable non-moral ignorance, the ignorance constitutes a complete excuse. But it is unclear what that view would be.

4.9 Conclusion

I have argued that the Forgiving Package (the conjunction of the Forgiving View, the Reduction View, and the Forgiving Extension) not only provides *an* explanation of the two observations introduced in section 4.5, but it provides the *best* explanation. To make this claim stick, I explored three different attempts to explain these observations without appealing to the Forgiving View.

First, I explored whether defenders of the Unforgiving View might be able to explain the two observations by appealing to a Quality of Will View. I showed that while the particular Quality of Will views defended by Harman and Talbert suffer from shortcomings, it is possible for an improved Quality of Will View to explain the *first* observation. However, it seems to be impossible for this view to explain the second observation.

Next, I discussed whether the “weak epistemic-sensitivity view” provides a more promising approach for defenders of the Unforgiving View. I argued that it does not. While the weak epistemic-sensitivity view can explain the first observation, it falls short in explaining the second.

Finally, I examined whether the Partially Forgiving View might be able to improve on the previous two views. I argued that it does not. The Partially Forgiving View does provide an attractive approach to the Comparison Judgments, but it fails to explain when non-culpable non-moral ignorance excuses. And it is unclear how the view could be supplemented to fix this problem.

Given that the three best alternatives to the Forgiving Package fail, I believe that there is good reason to conclude that the Forgiving Package provides the best explanation of the two observations. Of course, it is possible that further work will uncover a better explanation of the observations, which does not appeal to the Forgiving Package.³⁸ But until that explanation has been uncovered, the Forgiving Package's explanation should be accepted.

³⁸One promising alternative explanation, which I have become aware of since writing this chapter, would be to *combine* the Partially Forgiving View with the Weak Epistemic-Sensitivity View. This combination could potentially repair the defects of each view, while explaining both the first observation and the second. Evaluating this combination, and weighing it against the Forgiving Package, is a task that I do not have space to complete here. But note that even if this combination did provide as good an explanation of the two observations, a version of the abductive argument could still be run for the disjunction of the Forgiving and Partially Forgiving Views. Thus, appealing to this potential alternative explanation would do nothing to help defuse the threat the abductive argument poses to the Unforgiving View.

BIBLIOGRAPHY

- [1] Alston, William. 2002. "Plantinga, Naturalism, and Defeat." In James Beilby (ed.), *Naturalism Defeated? Essays on Plantinga's Evolutionary Argument Against Naturalism*. Ithaca: Cornell University Press.
- [2] Bergmann, Michael. 2006. *Justification Without Awareness: A Defense of Epistemic Externalism*. New York: Oxford University Press.
- [3] Blackburn, Simon. 1999. "Is Objective Moral Justification Possible on a Quasi-Realist Foundation?" *Inquiry* 42:213–227.
- [4] Brown, Campbell. forthcoming. "Minding the Is-Ought Gap." *Journal of Philosophical Logic* .
- [5] Byrne, Alex. 2004. "How Hard Are the Sceptical Paradoxes?" *Nous* 38:299–325.
- [6] Chignell, Andrew. 2010. "The Ethics of Belief." In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy (Fall 2010 Edition)*.
- [7] Christensen, David. 2009. "Disagreement as Evidence: The Epistemology of Controversy." *Philosophy Compass* 4:756–767.
- [8] Coates, D. Justin and Tognazzini, Neal A. (eds.). 2013. *Blame: Its Nature and Norms*. New York: Oxford University Press.
- [9] Dressler, Joshua. 1988. "Provocation: Partial Justification or Partial Excuse?" *The Modern Law Review* 51:467–480.
- [10] Elga, Adam. 2007. "Reflection and Disagreement." *Nous* 41:478–502.
- [11] Feldman, Fred. 2006. "Actual Utility, the Objection from Impracticality, and the Move to Expected Utility." *Philosophical Studies* 129:49–79.

- [12] Fitzpatrick, William J. 2008. "Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge." *Ethics* 118:589–613.
- [13] Graham, Peter. 2010. "In Defense of Objectivism about Moral Obligation." *Ethics* 121:88–115.
- [14] Harman, Elizabeth. 2011. "Does Moral Ignorance Exculpate?" *Ratio* XXIV.
- [15] Harman, Gilbert. 1986. *Change in View*. Cambridge: MIT Press.
- [16] Hills, Alison. 2009. "Moral Testimony and Moral Epistemology." *Ethics* 120:94–127.
- [17] Jackson, Frank. 1974. "Defining the Autonomy of Ethics." *The Philosophical Review* 83:88–96.
- [18] Jones, Karen. 1999. "Second-Hand Moral Knowledge." *The Journal of Philosophy* 96:55–78.
- [19] Joyce, Richard. 2009. "The Skeptic's Tale." *Philosophy and Phenomenological Research* 78:213–221.
- [20] Karmo, Toomas. 1988. "Some Valid (but no Sound) Arguments Trivially Span the 'Is'-'Ought' Gap." *Mind* 97:252–257.
- [21] Markovits, Julia. 2010. "Acting for the Right Reasons." *The Philosophical Review* 119:201–242.
- [22] McAuley, Finbarr. 1987. "Anticipating The Past: The Defence of Provocation in Irish Law." *The Modern Law Review* 50:133–157.
- [23] McConnell, Terrance. 2014. "Moral Dilemmas." In *The Stanford Encyclopedia of Philosophy* (Fall 2014 Edition).

- [24] McGrath, Sarah. 2004. "Moral Knowledge By Perception." *Philosophical Perspectives* 18:209–228.
- [25] Nagel, Thomas. 1979. "Death." In *Mortal Questions*. Cambridge: Cambridge University Press.
- [26] —. 1979. "Moral Luck." In *Mortal Questions*, 24–38. Cambridge: Cambridge University Press.
- [27] Nickel, Philip. 2001. "Moral Testimony and Its Authority." *Ethical Theory and Moral Practice* 4:253–266.
- [28] Pigden, Charles R. 1989. "Logic and the Autonomy of Ethics." *Australasian Journal of Philosophy* 67:127–151.
- [29] Plantinga, Alvin. 2000. *Warranted Christian Belief*. New York: Oxford University Press.
- [30] Prichard, Harold Arthur. 2002. "Duty and Ignorance of Fact." In Jim MacAdam (ed.), *Moral Writings*, 86–101. New York: Oxford University Press.
- [31] Prior, A. N. 1960. "The Autonomy of Ethics." *Australasian Journal of Philosophy* 38:199–206.
- [32] Pryor, Jim. 2000. "The Skeptic and the Dogmatist." *Nous* 34:517–549.
- [33] Rosen, Gideon. 2003. "Culpability and Ignorance." *Proceedings of the Aristotelian Society* 103:61–84.
- [34] —. 2004. "Skepticism about Moral Responsibility." *Philosophical Perspectives* 18:295–311.

- [35] —. 2008. “Kleinbart the Oblivious and Other Tales of Ignorance and Responsibility.” *Journal of Philosophy* 105:591–610.
- [36] Schwitzgebel, Eric. 2008. “The Unreliability of Naive Introspection.” *The Philosophical Review* 117:245–273.
- [37] Siegel, Susanna. 2011. “The Contents of Perception.” In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2011 Edition).
- [38] Sliwa, Paulina. 2012. “In defense of moral testimony.” *Philosophical Studies* 158:175–195.
- [39] Strawson, P.F. 2008. “Freedom and Resentment.” In *Freedom and Resentment and Other Essays*. New York: Routledge.
- [40] Sturgeon, Nicholas L. 1986. “What Difference Does It Make Whether Moral Realism Is True?” *The Southern Journal of Philosophy* XXIV (Supplement):115–149.
- [41] —. 2002. “Ethical Intuitionism and Ethical Naturalism.” In *Ethical Intuitionism: Re-evaluations*, 184–211. Oxford: Clarendon Press.
- [42] Talbert, Matthew. 2013. “Unwitting Wrongdoers and the Role of Moral Disagreement in Blame.” In David Shoemaker (ed.), *Oxford Studies in Agency and Responsibility*, volume 1. Oxford: Oxford University Press.
- [43] Tye, Michael and Wright, Briggs. 2011. “Is there are phenomenology of thought?” In *Cognitive Phenomenology*, 326–344. New York: Oxford University Press.
- [44] Watkins, Michael and Jolley, Kelly Dean. 2002. “Pollyanna Realism: Moral

Perception and Moral Properties." *Australasian Journal of Philosophy* 80:75–85.

[45] White, Roger. 2005. "Epistemic Permissiveness." *Philosophical Perspectives* 19:445–459.

[46] Williamson, Timothy. 2002. *Knowledge and Its Limits*. New York: Oxford University Press.

[47] Wolff, Robert Paul. 1970. *In Defense of Anarchism*. New York: Harper.

[48] Zimmerman, Michael. 1997. "Moral Responsibility and Ignorance." *Ethics* 107:410–426.