

Spring 1-1-2017

Moral Obligation, Evidence, and Belief

Jonathan Trevor Spelman

University of Colorado at Boulder, jontspelman@gmail.com

Follow this and additional works at: https://scholar.colorado.edu/phil_gradetds



Part of the [Applied Ethics Commons](#)

Recommended Citation

Spelman, Jonathan Trevor, "Moral Obligation, Evidence, and Belief" (2017). *Philosophy Graduate Theses & Dissertations*. 60.
https://scholar.colorado.edu/phil_gradetds/60

This Dissertation is brought to you for free and open access by Philosophy at CU Scholar. It has been accepted for inclusion in Philosophy Graduate Theses & Dissertations by an authorized administrator of CU Scholar. For more information, please contact cuscholaradmin@colorado.edu.

MORAL OBLIGATION, EVIDENCE, AND BELIEF

by

JONATHAN TREVOR SPELMAN

B.A., Ashland University, 2008

M.A., University of Missouri–St. Louis, 2010

A dissertation submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirement for the degree of
Doctor of Philosophy
Department of Philosophy

2017

This dissertation entitled:
Moral Obligation, Evidence, and Belief
written by Jonathan Trevor Spelman
has been approved for the Department of Philosophy

Chris Heathwood

David Boonin

Alastair Norcross

Graham Oddie

Date_____

The final copy of this dissertation has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

ABSTRACT

Spelman, Jonathan (Ph.D., Philosophy)
Moral Obligation, Evidence, and Belief
Dissertation directed by Associate Professor Chris Heathwood

Objectivism about moral obligation is the view that an agent's moral obligations do not depend on her beliefs or her evidence. The three leading theories in normative ethics (viz., consequentialism, Kantianism, and virtue ethics) have all traditionally been formulated as versions of objectivism. (For example, the traditional formulation of consequentialism requires agents to maximize value, not to do what they believe maximizes value or to do what their evidence suggests maximizes value.) In my dissertation, I argue that if we pay closer attention to how we use the phrase "moral obligation" and reflect more carefully on the nature of obligation more generally, we will find ourselves with good reasons to reject objectivism. Furthermore, I contend that our reasons for rejecting objectivism also speak against *prospectivism*, the view that an agent's moral obligations depend on her evidence, and in favor of *subjectivism*, the view that an agent's moral obligations depend on her beliefs. Finally, I argue that none of the most common objections to subjectivism are successful. Thus, we have most reason to be subjectivists about moral obligation.

A number of significant implications follow from this. For example, if we accept subjectivism about moral obligation, as well as the standard view that *moral rightness* and *wrongness* can be defined in terms of moral obligation, then we must also accept subjectivism about moral rightness and wrongness. And if we accept subjectivism about moral rightness and wrongness, then given some additional plausible assumptions, we must also accept (i) that an agent's act is morally

wrong if and only if she is blameworthy for performing that act, (ii) that agents rarely act morally wrongly, (iii) that we are rarely justified in believing that others have acted morally wrongly, and (iv) that we are rarely justified in saying that others have acted morally wrongly. Even if we can somehow resist these claims, however, the fact that we have most reason to accept subjectivism about moral obligation forces us to rethink not only how our moral concepts fit together but also how to talk about the morality of one another's acts.

DEDICATION

To my parents, Jeff and Jody, for their support and encouragement all these years,
to my wife, Anne Marie, for her love and companionship, and
to our children, Junia and Aquinas

ACKNOWLEDGEMENTS

There are many people whose support over the years has enabled me to write this dissertation. That list includes my parents (Jeff and Jody), my brother (Jordan), and my best friend (Joe D'Andrea). It also includes my undergraduate philosophy professors (Louis Mancha, Jeffrey Tiel, and William Vaughn), my undergraduate thesis committee (Kyle Fedler, David Foster, and Louis Mancha), and my master's thesis committee (Berit Brogaard, John Brunero, and Eric Wiland).

I also want to thank CU Boulder's philosophy department for its generous financial support, including a semester-long fellowship that allowed me to focus on my research, and my dissertation committee (David Boonin, Elinor Mason, Alastair Norcross, and Graham Oddie). I cannot thank my committee members enough for their helpful feedback along the way.

I am especially grateful to my wife, Anne Marie, for the support she has given me these past six years. While I may have completed the dissertation even if I hadn't met her, it wouldn't have been as good, and it surely would have taken a greater toll on my mental health.

Finally, I want to extend special thanks to my dissertation advisor, Chris Heathwood, for his help and encouragement these past seven years. I could not have asked for a better mentor.

CONTENTS

INTRODUCTION.....	1
CHAPTER	
1. THE MORAL ‘OUGHT’ AND MORAL OBLIGATION.....	6
1.0 Introduction.....	6
1.1 Objectivism and Its Competitors.....	8
1.2 The Drug Example and Jackson’s Argument against Objectivism.....	8
1.3 Objections to Jackson’s Argument against Objectivism.....	10
1.4 My Jackson-Style Argument against Objectivism.....	13
1.5 Objections to My Jackson-Style Argument against Objectivism.....	17
1.6 Other Responses to My Jackson-Style Argument against Objectivism.....	19
1.6.1 The Argument from Objective Moral Obligations.....	19
1.6.2 The Argument from Wrongness.....	21
1.6.3 The Argument from Advice.....	21
1.6.4 The Argument from Moral Advice.....	27
1.6.5 The Argument from Moral Monsters.....	28
1.7 Conclusion.....	34
2. MORAL OBLIGATION AND PUNISHMENT.....	36
2.0 Introduction.....	36
2.1 An Argument against Objectivism.....	38
2.2 An Argument for (7).....	40
2.2.1 An Argument for (8).....	41
2.2.2 In Defense of (8).....	44
2.2.3 An Argument for (9).....	48
2.3 Objections to the Argument for (2).....	50
2.4 Conclusion.....	51
3. AGAINST OBJECTIVISM ABOUT MORAL OBLIGATION.....	53
3.0 Introduction.....	53
3.1 The Simple Argument against Objectivism.....	54
3.2 Graham’s Objection and a Response.....	55
3.3 Graham’s Argument from New Information.....	61
3.4 A Response.....	62
3.5 Graham’s Argument from <i>Promise</i>	63
3.6 Graham’s Argument from <i>Promise*</i>	65

3.7	The Moral Obligation to Keep One’s Promises.....	66
3.8	The Moral Obligation to Track the Time.....	69
3.9	The Moral Obligation to Play It Safe.....	70
3.10	Conclusion.....	73
4.	IN DEFENSE OF SUBJECTIVISM ABOUT MORAL OBLIGATION.....	74
4.0	Introduction.....	74
4.1	Motivating Subjectivism.....	74
4.2	Does “Ought” Imply “Can”?.....	79
4.3	Did Hitler Act Wrongly?.....	81
4.4	Could One Avoid Wrongdoing by Failing to Attend to One’s Situation?.....	85
4.5	Could One Be Morally Infallible?.....	86
4.6	Conclusion.....	88
5.	IS THERE BLAMELESS WRONGDOING?.....	90
5.0	Introduction.....	90
5.1	Blameworthiness.....	91
5.2	Blameless Wrongdoing.....	97
5.2.1	Objectivism and Blameless Wrongdoing.....	99
5.2.2	Prospectivism and Blameless Wrongdoing.....	103
5.3	Conclusion.....	108
	BIBLIOGRAPHY.....	110

INTRODUCTION

Objectivism about moral obligation (hereafter, *objectivism*) is the view that an agent's moral obligations do not depend on her beliefs or her evidence. One might say that according to objectivism, an agent's moral obligation depend on the facts. The three leading theories in normative ethics (viz., consequentialism, Kantianism, and virtue ethics) have all traditionally been formulated as versions of objectivism. For example, the classic formulation of consequentialism requires agents to maximize value, not to do what they believe maximizes value or to do what their evidence suggests maximizes value. Similarly, the classic formulation of Kantianism requires agents to act in accordance with the categorical imperative, not to do what they believe accords with the categorical imperative or what their evidence suggests accords with the categorical imperative. Finally, the classic formulation of virtue ethics requires agents to act virtuously, not to do what they believe is virtuous or what their evidence suggests is virtuous.

In Chapters 1-3 of my dissertation, I argue against objectivism and the standard formulations of consequentialism, Kantianism, and virtue ethics. In Chapters 4-5, I argue for *subjectivism*, the view that an agent's moral obligations depend on her beliefs, and I defend it against objections. More specifically, in Chapter 1, "The Moral 'Ought' and Moral Obligation," I argue that facts about how we use the phrase "moral obligation," specifically facts about when we would be willing to admit to having violated a moral obligation, give us a good reason to think that an agent's moral obligations depend on her beliefs or her evidence. Imagine, for example, a case in which a doctor believes, in accordance with her evidence, that a certain drug is almost sure to cure her patient. Moreover, she believes, in accordance with her evidence, that this drug is low-risk, that

is, that if it does not cure her patient, it certainly will not harm him. Finally, however, assume that the doctor is wrong; the drug in question kills her patient. In this case, we might admit that the doctor failed to do what was best, but we would not say that she violated a moral obligation or, alternatively, that she failed to meet her moral obligations. This suggests that agents are not morally obligated to do what is best, but rather something like what they believe is best or what their evidence suggests is best. If that is right, then objectivism is false.

In Chapter 2, “Moral Obligation and Punishment,” I argue that this conclusion is corroborated by facts about the nature of *obligation* more generally. In particular, I argue that obligations are such that if an agent violates a certain kind of obligation, then he is subject to the relevant kind of punishment. (An agent who violates a *legal* obligation, for example, is subject to *legal* punishment.) This, in turn, entails that agents who violate *moral* obligations *deserve* punishment, for moral obligations, unlike other kinds of obligations, are necessarily just. If this is right, and objectivism is true, then the doctor described above deserves punishment. But that is implausible, so we should reject objectivism.

In Chapter 3, “Against Objectivism about Moral Obligation,” I respond to Peter A. Graham’s recent defense of objectivism.¹ More specifically, I argue that the phenomena he leverages in favor of objectivism can be explained equally well by both subjectivism, the view that an agent’s moral obligations depend on her beliefs, and *prospectivism*, the view that an agent’s moral obligations depend on her evidence. In particular, I argue that both subjectivism and prospectivism can explain why we often think, after changing our moral beliefs in light of new information, that we have figured out what we were morally obligated to do all along. Moreover, I argue that both subjectivism and prospectivism can explain why we might tell someone that she is morally obligated

¹ Graham 2010.

to perform some action, even when we know that she neither believes nor has any evidence suggesting that that action is best.

In Chapter 4, “In Defense of Subjectivism about Moral Obligation,” I point out that one of the central reasons for moving from objectivism to prospectivism (viz., the fact that objectivism is inconsistent with the claim that it is morally wrong to take unreasonable risks) also speaks in favor of moving from prospectivism to subjectivism. Moreover, I argue that none of the most common objections to subjectivism are successful. In particular, I contend that, contrary to what some have suggested, the truth of subjectivism is consistent with the claim that ‘ought’ implies ‘can.’² Additionally, while I admit subjectivism entails (i) that Hitler’s commanding genocide was morally permissible (assuming he believed it was best for him to command genocide), (ii) that agents with no beliefs about what is deontically best have no moral obligations, and (iii) that it is possible for agents to be morally infallible, I contend that none of these implications provides us with a good reason to reject the view.

Finally, in Chapter 5, “Is There Blameless Wrongdoing?” I respond to a further objection to subjectivism. According to that objection, we should reject subjectivism on account of the fact that it does not allow for cases of blameless wrongdoing. After sketching an account of blameworthiness on which an agent is not blameworthy for performing an action unless she believed she had a deontically better alternative, I argue that neither objectivist nor prospectivist arguments for blameless wrongdoing are convincing. Thus, the fact that subjectivism does not allow for cases of blameless wrongdoing does not give us a good reason to reject it.

If my arguments are sound, then we have most reason to accept subjectivism about moral obligation. That is, we have most reason to accept the view that an agent’s action is morally

²Zimmerman 2008: 14.

obligatory if and only if she believes it is best “in terms of what matters morally, insofar as [her] moral obligations are concerned.”³ A number of significant implications follow from this. For example, if we accept subjectivism about moral obligation, as well as the standard view that *moral rightness* and *wrongness* can be defined in terms of moral obligation, then we must also accept subjectivism about moral rightness and wrongness. And if we accept subjectivism about moral rightness and wrongness, then given some additional plausible assumptions, we must also accept (i) that an agent’s act is morally wrong if and only if she is blameworthy for performing that act, (ii) that agents rarely act morally wrongly, (iii) that we are rarely justified in believing that others have acted morally wrongly, and (iv) that we are rarely justified in saying that others have acted morally wrongly.

However, even if we have most reason to be subjectivists about *moral obligation*, *moral rightness*, and *moral wrongness*, it does not follow that we have most reason to be subjectivists about what we might call *moral bestness*, the property had by those acts that are not merely believed to be best but that are, in fact, best. Consider again the example I described earlier, in which a doctor falsely believes, in accordance with her evidence, that a certain drug, let us call it drug A, is best for her patient. In cases like this, I argue that the doctor is morally obligated to prescribe drug A to her patient, even if it will kill him, because that is what she believes is best. (Moreover, assuming that moral rightness and wrongness can be defined in terms of moral obligation, my arguments also entail that it is morally right for the doctor to prescribe drug A to her patient and morally wrong for the doctor to do otherwise.) But even if facts about what the doctor is *morally obligated* to do (and, therefore, facts about what it is morally right and morally wrong for the doctor to do) turn on the doctor’s beliefs, it does not follow that facts about what it would be *morally best* for the doctor to do

³ Zimmerman 2014: 4.

turn on the doctor's beliefs. For, presumably, it would be morally better for the doctor to do nothing than for her to do what she believes is best (i.e., what she is morally obligated to do) given that her doing what she believes is best will kill her patient.⁴ And if, as this example suggests, facts about what it would be morally best for an agent to do do not turn on what the agent believes, this blunts the significance of (i)-(iv), for we remain capable of saying things like (i*) that an agent's failing to do what is morally best is neither necessary nor sufficient for that agent's being blameworthy, (ii*) that agents frequently fail to do what is morally best, (iii*) that we are frequently justified in believing that others have failed to do what is morally best, and (iv*) that we are frequently justified in saying that others have failed to do what is morally best.

But even if the truth of (i)-(iv) does not rule out the truth of (i*)-(iv*), the truth of (i)-(iv) is still significant. Accordingly, some will surely try to resist (i)-(iv). The problem with this, however, is that if my arguments are sound, resisting (i)-(iv) most likely requires one to reject the standard view that moral rightness and wrongness can be defined in terms of moral obligation. This strikes me as too a high price to pay.

In the end, regardless of whether one is inclined to accept (i)-(iv) and admit the concept of *moral bestness* into one's moral vocabulary or, alternatively, to deny (i)-(iv) by rejecting the standard view that moral rightness and wrongness can be defined in terms of moral obligation, the fact that we have most reason to accept subjectivism about moral obligation forces us to rethink not only how our moral concepts fit together but also how we talk about the morality of one another's acts.

⁴ This, of course, leaves open the possibility that facts about what is morally best for an agent to do might turn, in part, on facts about what that agent believes is best.

CHAPTER 1

THE MORAL ‘OUGHT’ AND MORAL OBLIGATION⁵

1.0 Introduction

We often move from the fact that an agent *should* or *ought* to perform a certain action to the conclusion that he or she is *morally obligated* to perform that action. For example, when someone says that an agent *should* keep his or her promises and *ought* to refrain from torturing babies for fun, we often infer that the agent is morally obligated to do these things. But while we can *often* move from claims about what one *should* or *ought* to do to conclusions about what one is *morally obligated* to do, we cannot *always* do so. For example, that an agent *should* donate more money to charity does not entail that he or she is morally obligated to donate more money to charity (because this might be supererogatory rather than obligatory).⁶ Additionally, that an agent *ought* to refrain from putting his or her elbows on the table does not entail that he or she is morally obligated to refrain from putting his or her elbows on the table (because this might be required by etiquette rather than morality).

So, although *shoulds* and *oughts* sometimes entail moral obligations, they don’t always do so. Nevertheless, many philosophers use our intuitions about how we use these words to draw

⁵ I am grateful to Chris Heathwood, Alastair Norcross, Elinor Mason, Michael J. Zimmerman, an anonymous referee, and audiences at both CU Boulder and the 2016 Pacific Division Meeting of the American Philosophical Association for their helpful feedback on earlier drafts of this chapter.

⁶ This, I take it, is the thesis of Elizabeth Harman’s recent paper (2016).

conclusions about what agents are *morally obligated* to do.⁷ This, I argue, is a significant mistake that matters for certain disputes in ethical theory. If we pay closer attention to how we use the phrase ‘moral obligation,’ we will see not only that objectivism about moral obligation, the view that our moral obligations do not depend on beliefs or our evidence, is false, but also that we should be skeptical of the weaker view that there is an objective sense of ‘moral obligation.’

In Section 1.1, I describe objectivism about moral obligation (hereafter, *objectivism*) and contrast it with two competing views of moral obligation, subjectivism and prospectivism. In Section 1.2, I present Frank Jackson’s well-known drug example and corresponding argument against objectivism, and in Section 1.3, I outline the most popular objection to that argument. In Section 1.4, I admit that this objection succeeds in refuting Jackson’s argument against objectivism but argue that if we pay attention to our intuitions about how to use the phrase ‘moral obligation,’ we can construct a similar argument against objectivism that does not fall prey to that objection. In Section 1.5, I discuss three objections to my Jackson-style argument against objectivism. In responding to the third of those objections, I suggest that my argument threatens not only objectivism but also the weaker view that there is an objective sense of ‘moral obligation.’ In Section 1.6, I discuss numerous arguments for this weaker view and contend that none of them are decisive. I conclude, therefore, not only that we should reject objectivism,⁸ but also that we should be skeptical of this weaker view.

⁷ Some advance arguments where our intuitions about how to use ‘should’ or ‘ought’ are used to draw conclusions about our moral obligations (see Feldman 1986, Jackson 1991, and Dorsey 2012). Others use ‘should’ or ‘ought’ to mean ‘morally obligated’ (see Zimmerman 2006, 2008, 2014, and Parfit 2011). And still others suggest that there is a conceptual link between ‘ought’ and ‘moral obligation’ (see Andrić 2011a).

⁸ On the assumption that *moral obligatoriness* is definable in terms of *moral rightness* or *wrongness*, my argument entails that we should reject objectivism about moral rightness and wrongness as well. Alternatively, on the assumption that we should accept objectivism about moral rightness and wrongness, my argument entails that we should deny that moral obligatoriness is definable in terms of moral rightness or wrongness. Although I find the former assumption more plausible than the latter, those who disagree with me are

1.1 Objectivism and Its Competitors

Objectivism is the view that an agent's moral obligations do not depend on her beliefs or her evidence; they depend on the facts. This view is typically contrasted with views on which an agent's moral obligations *do* depend on her beliefs or her evidence. One such view is *subjectivism*, on which an agent's moral obligations depend on her beliefs. Another such view is *prospectivism*, on which an agent's moral obligations depend on her evidence.

The three leading theories in normative ethics (viz., consequentialism, Kantianism, and virtue ethics) have traditionally been formulated as versions of objectivism. The classic formulation of consequentialism, for example, requires agents to maximize value, not to do what they believe maximizes value or to do what their evidence suggests maximizes value. Similarly, the classic formulation of Kantianism requires agents to act in accordance with the categorical imperative, not to do what they believe accords with the categorical imperative or what their evidence suggests accords with the categorical imperative. Finally, the classic formulation of virtue ethics requires agents to act virtuously, not to do what they believe is virtuous or what their evidence suggests is virtuous. If objectivism is false, then we should abandon the traditional way of formulating these theories.

1.2 The Drug Example and Jackson's Argument against Objectivism

In "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," Jackson presents the following example.

welcome to interpret this chapter in the second way. (I am grateful to Alastair Norcross for encouraging me to highlight this point.)

The Drug Example

Jill is a physician who has to decide on the correct treatment for her patient, John, who has a minor but not trivial skin complaint. She has three drugs to choose from: drug A, drug B, and drug C. Careful consideration of the literature has led her to the following opinions. Drug A is very likely to relieve the condition but will not completely cure it. One of drugs B and C will completely cure the skin condition; the other though will kill the patient, and there is no way that she can tell which of the two is the perfect cure and which the killer drug. What should Jill do?⁹

The answer to Jackson's closing question is obvious. Because both drugs B and C have a fifty percent chance of killing John, it would be too risky for Jill to prescribe either of them. Jill *should* prescribe drug A.

Jackson infers from this that objective versions of consequentialism, on which an agent's moral obligations depend on their actual consequences, must be false.¹⁰ Although he does not explicitly state his argument for this conclusion, here is one way it might go:

Jackson's Argument against Objectivism

- (1) Jill should prescribe drug A.
- (2) If Jill should prescribe drug A, then she is morally obligated to prescribe drug A.¹¹
- (3) Therefore, Jill is morally obligated to prescribe drug A.
- (4) If an objective version of consequentialism is true, then Jill is not morally obligated to prescribe drug A.
- (5) Therefore, it is not the case that an objective version of consequentialism is true.

It is plausible that Jill *should* prescribe drug A, and it is plausible that this 'should' entails a moral obligation. Thus, it is plausible that Jill is morally obligated to prescribe drug A. But if an objective version of consequentialism is true, then Jill is not morally obligated to prescribe drug A because

⁹Jackson 1991: 462-463. See Regan 1980 for an earlier example along the same lines.

¹⁰Jackson 1991: 465-466.

¹¹Jackson does not explicitly make this claim, but he clearly moves from the fact that Jill *should* prescribe drug A to either (i) the conclusion that Jill is morally obligated to prescribe drug A or (ii) the conclusion that Jill's prescribing drug A would be morally right (i.e., that it is at least morally permissible, if not morally obligatory). Either way of formulating the argument is acceptable for my purposes.

objective versions of consequentialism require agents to maximize value, and Jill's prescribing drug A would not do that. Her prescribing one of the other drugs would. Thus, it is not the case that an objective version of consequentialism is true.

While Jackson's argument was only meant to refute objective versions of consequentialism, it actually threatens objective versions of any theory on which the rightness of some actions depends on their consequences. So, on the plausible assumption that the correct theory of right action is one on which the rightness of some actions depends on their consequences, it follows that Jackson's argument actually threatens objectivism more generally. I will assume as much going forward.

1.3 Objections to Jackson's Argument against Objectivism

In response to Jackson's argument, one might admit that one speaks truly when one says, 'Jill should prescribe drug A' while also pointing out that there are similar situations in which one could truly say, 'Jill should have prescribed drug B.' Here is an example:

Jill's Retrospective Judgment

Sometime after prescribing drug A, a more knowledgeable doctor informs Jill that it was drug B that would have cured her patient's condition. At that point, Jill says, "I guess that means I *should have* prescribed drug B."¹²

¹² I would encourage those readers who are hesitant to admit that Jill could truly admit that she *should have* prescribed drug B to consider similar cases in which the claim might be even more plausible. First, we could imagine a case like *Jill's Retrospective Judgment* in which it is not Jill who says that she should have prescribed drug B, but rather the more knowledgeable doctor. For example, imagine that the more knowledgeable doctor, knowing that drug B would have cured Jill's patient, finds out that Jill prescribed drug A. In that case, it seems that he or she could truly say, "That's too bad. She *should have* prescribed drug B." Additionally, we could imagine a case in which we are jurors who convicted a man of murder only to later find out that he had been framed. Even if all our evidence spoke in favor of convicting the man, it seems that we could, nevertheless, truly admit, upon finding out that the man was framed, that we should not have convicted him (Dorsey 2012: 4).

This suggests that we can truly say both that Jill *should* prescribe drug *A* and that Jill *should have* prescribed drug *B*. One might conclude from this that ‘should’ must be ambiguous, in which case Jackson may be guilty of equivocating. The sense of ‘should’ he uses in (1) might not be the sense of ‘should’ that entails a moral obligation.¹³

Here is another way to put the objection. *The Drug Example* suggests that we can truly say that Jill should prescribe drug *A*, thereby giving us a reason to think that she is morally obligated to prescribe drug *A*. However, *Jill’s Retrospective Judgment* suggests that we can truly say that Jill should have prescribed drug *B*, thereby giving us a reason to think that she is morally obligated to prescribe drug *B*. Unless we have a principled way of determining when *shoulds* entail moral obligations, we cannot be sure of (2), the claim that if Jill should prescribe drug *A*, then she is morally obligated to prescribe drug *A*.

As this objection reveals, ‘should’ appears to be ambiguous, which might lead one to think that ‘moral obligation’ is similarly ambiguous. For example, one might think that insofar as there is a subjective (or prospective) sense of ‘should,’ on which what one *should* do depends on one’s beliefs (or evidence), there must also be a subjective (or prospective) sense of ‘moral obligation’ on which what one is *morally obligated* to do depends on one’s beliefs (or evidence). Similarly, one might think that insofar as there is an objective sense of ‘should’ on which what one *should* do depends on the facts, there must also be an objective sense of ‘moral obligation’ on which what one is *morally obligated* to do depends on the facts. On this view, both of the *shoulds* in question entail moral obligations, but they entail different senses of moral obligation, thereby avoiding any contradiction.

¹³ Although a number of philosophers, including Brian Berkey, have raised this line of argument in conversation, Feldman 1986 and Dorsey 2012 are the best examples of it in the literature.

Among those who take this view, that ‘moral obligation’ is ambiguous, many suggest that the objective sense of ‘moral obligation’ is the primary or most fundamental sense of that phrase. This group includes Henry Sidgwick, Fred Feldman, Dale Dorsey, and Julia Driver.¹⁴ Although these philosophers are sometimes called objectivists, here I will call them *weak objectivists* (because, unlike *objectivists*, they are open to there being multiple senses of ‘moral obligation’). Others who think that ‘moral obligation’ is ambiguous suggest that it is the prospective sense of ‘moral obligation’ that is the primary or most fundamental sense of that phrase. This group includes Allan Gibbard and Elinor Mason.¹⁵ I will call these philosophers *weak prospectivists*. Still others (e.g., Holly Smith) have either ignored the question of whether there is a primary or most fundamental sense of ‘moral obligation’ or, like Derek Parfit, denied that there is a primary or most fundamental sense of ‘moral obligation.’¹⁶ I will call these philosophers *ambiguity theorists*.

There is, however, another group of philosophers who deny that ‘moral obligation’ is ambiguous. This group includes G. E. Moore, H. A. Prichard, W. D. Ross, Peter A. Graham, and Michael J. Zimmerman.¹⁷ Of these five, only Moore and Graham are objectivists,¹⁸ and only Graham responds to Jackson’s argument.¹⁹ In doing so, Graham follows those above in admitting that we can truly say of Jill that she *should* prescribe drug A. However, unlike those above, he contends that the ‘should’ in question is not a moral ‘should’ but a pragmatic ‘should’ and therefore does not tell us anything about Jill’s moral obligations.²⁰

¹⁴ See Sidgwick 1981/1907, Feldman 1986, Dorsey 2012, and Driver 2012a, 2012b.

¹⁵ See Gibbard 1990 and Mason 2013.

¹⁶ See Smith 2010 and Parfit 2011.

¹⁷ See Moore 1912, Prichard 2002/1932, Ross 1939, Graham 2010, and Zimmerman 2008. Zimmerman is more tentative about this in his more recent work (2014).

¹⁸ Prichard and Ross are subjectivists while Zimmerman is a prospectivist.

¹⁹ Obviously, Moore does not respond to Jackson’s argument because he predates Jackson.

²⁰ Graham 2010: 103. I take this suggestion to be rather implausible. The Jackson-style argument I present in the following section illustrates why.

In the next section, I present a Jackson-style argument that sidesteps these worries. In particular, I argue that if we consult our intuitions about what Jill is *morally obligated* to do rather than our intuitions about what she *should* do, we will see that objectivism is false.

1.4 My Jackson-Style Argument against Objectivism

To recap, Jackson's drug example provides us with a case in which it seems obvious that an agent *should* perform a certain action despite knowing that doing so will not maximize value (or, more generally, despite knowing that doing so is *not best relative to the facts*). It seems equally obvious, however, that if that agent later found out which of his or her alternatives would have maximized value, he or she could truly say that he or she *should have* performed that action instead. Although the conclusions philosophers have drawn from this are varied, the important point is that philosophers have not taken Jackson's argument against objective versions of consequentialism as decisive.

I agree that Jackson's argument, as stated above, is not decisive, and I believe that the reason for that is that (2) is false; *shoulds* and *oughts* do not always entail moral obligations. As I mentioned earlier, even if an agent *should* donate more money to charity, it does not follow that he or she is *morally obligated* to donate more money to charity because such an act may be supererogatory rather than obligatory. Moreover, even if an agent *ought* to refrain from putting his or her elbows on the table, it does not follow that he or she is *morally obligated* to refrain from putting

his or her elbows on the table because such an act might be required by etiquette rather than morality.²¹

What I want to argue, however, is that there is another kind of case in which *shoulds* and *oughts* do not entail moral obligations. Consider again *Jill's Retrospective Judgment*. In that case, Jill can truly say that she *should have* prescribed drug B. Of course, Jill would not say this because she thinks that prescribing drug B was supererogatory, for she would not think that her prescribing drug B was above and beyond the call of duty. Moreover, Jill would not say that she *should have* prescribed drug B because she thinks that prescribing drug B was required by some standard other than morality, for she would not think that prescribing drug B was required by etiquette, the law, prudence, rationality, etc. Instead, Jill would say that she *should have* prescribed drug B because she thinks that prescribing drug B was required by morality. That is, Jill would say that she *should have* prescribed drug B for *moral reasons*. Thus, we would expect that this 'should' would entail a moral obligation.

But this 'should' does not entail a moral obligation. To see this, notice that in a case like *Jill's Retrospective Judgment*, where Jill has prescribed drug A only to find out that drug B would have cured her patient's condition, Jill can truly say that she *should have* prescribed drug B. However, she would not say that she *violated a moral obligation* in prescribing drug A. Let me reiterate that point. Jill would not say that she violated a moral obligation in prescribing drug A. What does this tell us? Well, the fact that Jill *would not say* that she violated a moral obligation in prescribing drug A suggests that she *cannot truly say* that she violated a moral obligation in prescribing drug A. And if Jill cannot truly say that she violated a moral obligation in prescribing drug A, it follows that she was not

²¹ Note that I do not mean to rule out the possibility of its being morally wrong to put one's elbows on the table in at least some circumstances.

morally obligated to prescribe drug A. So, the fact that Jill would not say that she violated a moral obligation in prescribing drug A suggests that she was not morally obligated to prescribe drug A, in which case objectivism is false.

Note that I do not mean to simply insist on this claim that Jill would not say that she violated a moral obligation in prescribing drug A. Instead, I am inviting you, the reader, to consult your linguistic intuitions. If Jill were to look back on her decision to prescribe drug A (after learning that drug B would have cured her patient's condition), would she say that she *violated a moral obligation* in prescribing drug A? Of course not! Moreover, it is implausible that she would even make the weaker claim that she *failed to satisfy her moral obligations*.

Jill might *wish* that she had prescribed drug B or *regret* having prescribed drug A, but she would not say that she was *morally obligated* to prescribe drug B or that she *violated a moral obligation* in prescribing drug A. Similarly, although she can admit that her prescribing drug B would have been *better* than her prescribing drug A and therefore that she *should have* prescribed drug B or *should not have* prescribed drug A, she cannot admit, at least not plausibly, that she was *morally obligated* to prescribe drug B or that she *violated a moral obligation* in prescribing drug A.

This linguistic data tells us something about the nature of moral obligation. One thing it tells us is that, even if it is true that an agent *should* perform some (non-supererogatory) act for moral reasons, it does not follow from this that she is *morally obligated* to perform that act. In fact, the agent in question can be *morally prohibited* from performing that act. This is a surprising result, and it suggests that 'should' is ambiguous in a way that 'morally obligated' is not. Just because there is an objective sense of 'should' does not mean there is an objective sense of 'morally obligated.' If this is right, then *contra* Zimmerman and Parfit,²² it is a mistake to use 'should' or 'ought' to mean

²² See Zimmerman 2006, 2008, 2014, and Parfit 2011.

‘morally obligated.’ Moreover, *contra* Feldman, Jackson, and Dorsey,²³ we cannot use our intuitions about what agents ‘should’ or ‘ought’ to do to draw conclusions about what they are morally obligated to do.

On top of this, the fact that Jill cannot truly say either that she was morally obligated to prescribe drug B or that she violated a moral obligation in prescribing drug A seems to be decisive evidence against the claim that Jill was morally obligated to prescribe drug B. She was morally obligated to prescribe drug A. Of course, she may have done something that she, in some sense, *should not have* done. But she did not violate a moral obligation. If this is right, then *contra* Moore and Graham,²⁴ objectivism must be false.

This line of reasoning suggests the following Jackson-style argument:

My Jackson-Style Argument against Objectivism

- (6) Jill could not truly admit to having violated a moral obligation in prescribing drug A.
- (7) If Jill could not truly admit to having violated a moral obligation in prescribing drug A, then Jill was not morally prohibited from prescribing drug A.
- (8) Therefore, Jill was not morally prohibited from prescribing drug A.
- (9) If objectivism is true, then Jill *was* morally prohibited from prescribing drug A.
- (10) Therefore, objectivism is false.

If the premises of this argument are true, we must reject objectivism.

²³ See Feldman 1986, Jackson 1991, and Dorsey 2012.

²⁴ See Moore 1912 and Graham 2010.

1.5 Objections to My Jackson-Style Argument against Objectivism

It is hard to see how objectivists could deny (7) or (9). So, I am inclined to think that they will respond to the argument above by either denying (6) or arguing that the conclusion of the argument has unacceptable implications. In this section, I will raise and respond to a few objections to (6). In the following section, I will raise and respond to five arguments to the effect that the conclusion of my Jackson-style argument has unacceptable implications.

First, an objectivist might attempt to deny (6) by arguing that people are prideful and therefore that they refuse to admit to having violated moral obligations even when they *could*. But even if that is true, we need not think that Jill's refusal to admit that she violated a moral obligation is attributable to pride. We can assume that Jill is more than willing to admit when she has violated a moral obligation. Even then, it does not seem that Jill could *truly* admit to having violated a moral obligation in prescribing drug A. Furthermore, when *we* consider whether Jill violated a moral obligation in prescribing drug A, we judge that she did not, and presumably that judgment is not attributable to pride.²⁵

Second, one might attempt to deny (6) by arguing that there is something special about *The Drug Example*. For example, one might contend that there is something about doctors that makes them especially unlikely to admit to having violated their moral obligations or that there is something about three-option cases that makes objectivism seem less plausible. The problem with this objection, however, is that we can use very different cases to run the same argument. Consider the following example, which does not involve doctors and can be thought of as a two-option case.

The Bomb Example

Amanda is returning to her hotel room late one night when she sees smoke escaping a nearby room. She tries to open the door to help anyone who might be inside, but

²⁵ I am grateful to Chris Heathwood for drawing my attention to this point.

it is locked. In making her way out of the hotel, she notices a fire alarm and pulls it. Unbeknownst to Amanda, however, an evildoer has rigged that fire alarm to detonate a bomb in the hotel's basement. As it turns out, the fire on its own would not have resulted in any casualties, but the bomb explosion kills a number of innocent people.²⁶

In this case, we are inclined to say, prior to Amanda's pulling the fire alarm, that she *should* do so, and after the fact, we are inclined to say that she *should not have* done so. *The Bomb Example* is perfectly analogous to *The Drug Example* in this way. Moreover, when we consider whether Amanda would admit that she violated a moral obligation in pulling the fire alarm, the answer is obvious; she would not. This, of course, speaks just as strongly against objectivism as the fact that Jill would not admit that she violated a moral obligation in prescribing drug A. This suggests that there is nothing special about *The Drug Example*.

Third, one might attempt to deny (6) by arguing that although there is a subjective sense of 'moral obligation' on which Jill *cannot* truly say that she violated a moral obligation, there is an objective sense of 'moral obligation' on which Jill *can* truly say that she violated a moral obligation in prescribing drug A.²⁷ This objection is plausible only if it is plausible that there is an objective sense of 'moral obligation,' and that is plausible only if it is plausible that there is some sense in which Jill could truly say that she violated a moral obligation. As I see it, however, *there is no sense in which Jill can truly say that she violated a moral obligation*. If there were an objective sense of 'moral obligation' on which Jill violated a moral obligation in prescribing drug A, then her prescribing drug A was in some sense *morally impermissible* (or *prohibited*). But there is no sense in which Jill's prescribing drug A was *morally impermissible*, as evidenced by the fact that Jill would not say that her prescribing drug A was *morally impermissible*. It might be true that Jill's prescribing drug A was

²⁶ This example is inspired by an example that Caleb Pickard described in conversation.

²⁷ Weak objectivists, weak prospectivists, and ambiguity theorists presumably have to take this line.

not *best*, but that does not entail that it was *morally impermissible*. This suggests not only that objectivism is false, but also that, *contra* Sidgwick, Feldman, Gibbard, Smith, Parfit, Dorsey, Driver, and Mason,²⁸ there is no objective sense of ‘moral obligation.’

1.6 Other Responses to My Jackson-Style Argument against Objectivism

I have argued that none of the preceding objections to my Jackson-style argument is particularly plausible. If I am correct, then not only is objectivism false, but there is not even an objective sense of ‘moral obligation.’ Given the intuitive plausibility of my claim that Jill cannot truly admit to having violated a moral obligation in prescribing drug A, I think we ought to accept both of these conclusions, provided that they do not have any unacceptable implications. Below, I respond to five reasons for thinking that these conclusions do have unacceptable implications.

1.6.1 The Argument from Objective Moral Obligations

First, one might argue that if I am right that there is no sense in which Jill can truly say that she violated a moral obligation, then we would not be able to make sense of claims about one’s ‘objective moral obligations.’²⁹

This claim is clearly false. We can make sense of claims about one’s ‘objective moral obligations’ by reference to what one’s *actual* moral obligations would be if one knew all the facts. On this view, to say that someone has an objective moral obligation to ϕ is just to say that he or

²⁸ See Sidgwick 1981/1907, Feldman 1986, Gibbard 1990, Smith 2010, Parfit 2011, Dorsey 2012, Driver 2012a, 2012b, and Mason 2013.

²⁹ I am grateful to Alastair Norcross for suggesting this worry.

she *would* have an *actual* moral obligation to ϕ if he or she knew all the morally relevant facts. So, for example, on the assumption that drug B is the perfect cure, we can say that Jill was objectively morally obligated to prescribe drug B on account of the fact that she *would* have had an *actual* moral obligation to prescribe drug B *if* she had known all the morally relevant facts.

Given that I have just defined what an objective moral obligation is, one might wonder how I can deny the existence of an objective sense of ‘moral obligation.’ For presumably, if there are objective moral obligations, then there is an objective sense of ‘moral obligation.’

That does not follow. Imagine that we disagree over whether there are any unicorns. You point to a toy unicorn as evidence for the conclusion that there are unicorns. In response, I note that toy unicorns are no more unicorns than rubber ducks are ducks. In doing so, I do not deny that there are toy unicorns, but I deny that there is a sense of ‘unicorn’ on which toy unicorns count as unicorns. We could, of course, stipulate into existence a sense of ‘unicorn’ on which toy unicorns are unicorns, but in doing so we would not have shown that there is a sense of the English term ‘unicorn’ on which toy unicorns are unicorns.

Objective moral obligations, I want to suggest, are like toy unicorns in that there is no sense of ‘moral obligation’ on which an ‘objective moral obligation’ counts as a moral obligation. We could, of course, stipulate into existence a sense of ‘moral obligation’ on which ‘objective moral obligations’ are moral obligations, but in doing so we would not have shown that there is a sense of the English phrase ‘moral obligation’ on which objective moral obligations are moral obligations. Therefore, the fact that we can talk about objective moral obligations does not commit us to the existence of an objective sense of the English phrase ‘moral obligation.’

Before moving on, however, I should reiterate that my argument against objectivism does not depend on there being no objective sense of ‘moral obligation.’ Instead, it depends on there being a non-objective (e.g., subjective or prospective) sense of ‘moral obligation.’ I take it that *The*

Drug Example and *The Bomb Example* prove that there is such a sense. Whether they also suggest that there is no objective sense of ‘moral obligation’ is a further question. I think that they do, but if it turns out that they do not, then my argument against objectivism still goes through.

1.6.2 The Argument from Wrongness

A second way one might respond to my Jackson-style argument is to point out that after Jill learns that drug B would have cured her patient, she can truly say, “Oh no, I picked the wrong drug!” This suggests that Jill’s prescribing drug A was ‘wrong,’ which suggest that she *did* violate a moral obligation in prescribing drug A.

Note, however, that ‘wrong,’ like ‘should’ and ‘ought,’ is ambiguous. Although we can certainly use ‘wrong’ to acknowledge our moral mistakes, we can also use ‘wrong’ to acknowledge our non-moral mistakes. For example, when I attribute my being late to a party to the fact that I made a wrong turn, I am not admitting to having made a moral mistake; I am admitting to having made a non-moral mistake. The same is true of Jill. When she says that she picked the wrong drug, she is not admitting to having made a moral mistake, but to having made a non-moral mistake. This is obvious once we note that Jill would not admit to having violated a moral obligation in prescribing drug A.

1.6.3 The Argument from Advice

Third, one might use the following case to argue against my conclusion.

Advice

Prior to deciding which drug to use, Jill asks Jack, a more knowledgeable doctor (who knows that drug B will cure John), what she ought to do. In response, Jack tells Jill that she ought to prescribe drug B.³⁰

According to my opponent, the fact that Jack seems to speak truly in telling Jill that she *ought* to prescribe drug B provides support for the claim that she is morally obligated to prescribe drug B. Thus, *Advice* provides support for the ambiguity theory at least, if not weak or full-blown objectivism.

Zimmerman, a prospectivist, has responded to this argument in a couple of places.³¹ In each, Zimmerman suggests that Jack may be morally obligated to tell Jill that she ought to prescribe drug B on account of the fact that his doing so may be best given his evidence, but Zimmerman denies that Jack's claim is *true*, at least on the assumption that the 'ought' in question expresses moral obligation.

This is not how I would respond to this example. One of the main lessons of this chapter is that there are, in addition to the well-known cases in which *shoulds* and *oughts* do not entail moral obligations,³² other cases in which *shoulds* and *oughts* do not entail moral obligations. More specifically, there are cases in which one (i) *should* or *ought* to perform some (non-supererogatory) action and (ii) *should* or *ought* to perform that action for moral reasons, but (iii) is not morally obligated to perform that action. Thus, we should avoid using our intuitions about what agents *should* or *ought* to do to draw conclusions about what they are *morally obligated* to do. Even if Jack

³⁰ A similar case appears in Zimmerman 2006: 344-345.

³¹ See Zimmerman 2006: 344-345 and 2008: 32.

³² By "the well-known cases in which *shoulds* and *oughts* do not entail moral obligations," I mean those in which an act is supererogatory rather than obligatory and those in which an act is required by some standard (e.g., etiquette, the law, prudence, rationality, etc.) other than morality.

speaks truly in telling Jill that she *ought* to prescribe drug B, it does not follow that Jill is *morally obligated* to prescribe drug B.

Hence, rather than responding as Zimmerman does, we would be better off admitting that Jack's claim, that Jill ought to prescribe drug B, is true while denying that the 'ought' in question expresses moral obligation. To see why this is the more plausible interpretation, consider the following case:

Moral Advice

Prior to deciding which drug to use, Jill asks Jack, a more knowledgeable doctor (who knows that drug B will cure John), what she is *morally obligated* to do. In response, Jack tells Jill that she is *morally obligated* to prescribe drug B.

If the 'ought' in *Advice* expresses moral obligation, then our intuitions about *Moral Advice* should mirror our intuitions about *Advice*. But they do not. While it is easy to imagine someone like Jill asking someone like Jack what she *ought* to do, it is more difficult to imagine her to asking someone like Jack what she is *morally obligated* to do.³³ Furthermore, even if we can imagine Jill asking Jack what she is morally obligated to do, it is not obvious that Jack would be willing to give Jill a direct answer to her question. Instead, he might respond by telling her that she *ought* to prescribe drug B, thereby leaving aside the question of what she is *morally obligated* to do. And finally, even if we assume that Jack would be willing to tell Jill that she is morally obligated to prescribe drug B, it is not clear how plausible that claim is. I should, of course, admit that once Jack has told Jill that she is morally obligated to prescribe drug B, it is quite plausible to think that she is morally obligated to do so, for in telling Jill that she is morally obligated to prescribe drug B, Jack presumably makes

³³ On my view, Jill is more likely to ask Jack what she *ought* to do than what she is *morally obligated* to do because questions about what one *ought* to do are questions about what is *best*, and that is what Jill is unsure about, what is best.

it the case that Jill knows (or is at least justified in believing) that prescribing drug B is best with respect to what matters morally. But objectivists are committed to the even stronger view on which Jill was morally obligated to prescribe drug B even before Jack said so. That claim is much less plausible.

But let me grant my opponent that Jack's claim is *plausible* (i.e., that it is *plausible* that Jill is morally obligated to prescribe drug B) and that its plausibility is, to at least some degree, independent of the effect Jack's claim has on Jill's epistemic position. Certainly that gives us a reason to prefer objectivism to its competitors, does it not? I do not think that it does. Even if we assume that Jack's claim is plausible independent of the effect it has on Jill's epistemic position, it may very well get its plausibility, not from its semantic content, but from its pragmatic content.

To see why I think that, consider the following example:

Jack's Retrospective Judgment

Jill must decide whether to prescribe drug A, B, or C on her own. Afterward, however, she asks Jack whether she violated a moral obligation in prescribing drug A.

In this case, can Jack truly say that Jill violated a moral obligation in prescribing A? I do not think so. He might be able to truly say that it would have been *better* for Jill to prescribe drug B than for her to prescribe drug A, that Jill *ought to* have prescribed drug B, or that Jill *ought not to* have prescribed drug A, but he cannot truly say that Jill *was morally obligated to* prescribe drug B or that she *violated a moral obligation in* prescribing drug A. And, of course, if this is correct, then we should not think that he can truly say that Jill is morally obligated to prescribe drug B in *Moral Advice*.

If we grant that in *Moral Advice*, Jack's claim that Jill is morally obligated to prescribe drug B does not seem false, then we are left wondering why that is. To answer that question, I want to begin by noting that whereas in *Jack's Retrospective Judgment*, Jack has a single goal in mind, namely

answering Jill's question (by telling her whether he thinks she violated a moral obligation), in *Moral Advice*, he has at least two goals in mind. First, he wants to answer Jill's question (by telling her what she is morally obligated to do), but more importantly, he wants to give her a recommendation about what to do.³⁴

Given these two distinct goals, it should come as no surprise that Jack tells Jill that she is morally obligated to prescribe drug B in *Moral Advice*, for in telling her that, he accomplishes his primary goal of giving Jill a recommendation about what to do, without obviously failing to answer her question. Because Jack's claim accomplishes both of these things, it is apt and therefore seems true even though it is, strictly speaking, false.

Now, what I have said assumes that Jill's goal in asking the more knowledgeable doctor what she is morally obligated to do is not to find out what she is morally obligated to do but rather to find out what he would recommend that she do. If it were common for doctors in Jill's position to ask more knowledgeable doctors about their moral obligations, then this assumption would not be plausible. But it is not common for doctors in Jill's position to ask more knowledgeable doctors about their moral obligations. As I suggested earlier, it is much more common for doctors in Jill's position to ask more knowledgeable doctors what they *should* or *ought* to do to treat their patients' illnesses. This is not to say that it is not common for people to ask for moral advice. People ask for moral advice all the time. But it is not common for doctors in situations like Jill's to do so. Jill presumably knows what she is morally obligated to do given the medical information she has. What she wants is more medical information, not moral advice.

³⁴ On top of this, Jack probably has not spent much time reflecting on the nature of moral obligation and therefore is not in a position to draw a distinction between what it would be *best* for Jill to do, what she *ought* to do, and what she is *morally obligated* to do. Or, even if he has spent time reflecting on the nature of moral obligation, he might recognize that the issue is contentious and therefore not worth getting into.

In support of this suggestion, consider some paradigmatic cases in which people ask for moral advice. In one kind of case, an agent thinks that he is biased in favor of taking some course of action rather than another, so he seeks out a third-party who he trusts to be impartial (or at least less partial than he is). Then, the agent lays out the details of his situation as he sees them and asks the third party what he is obligated to do given those details. For example, if an agent finds something valuable lying on the street, he might consult a friend to determine the lengths to which he must go to track down its owner. In another kind of case, an agent seeks out a religious authority (e.g., an imam, pastor, priest, rabbi, etc.) or even an ethicist in hopes that this person will be able to provide the agent with insight into the nature of morality. For example, an agent might consult a religious authority (or an ethicist!) to determine whether (or how much) he or she is morally obligated to donate to charity.

In neither of these cases does an agent seek out a moral advisor for her practical knowledge about what sorts of consequences a particular action will have.³⁵ This highlights an important distinction between moral advice and practical advice. We rarely, if ever, go to doctors for moral advice; we go to trusted, impartial third-parties, religious authorities, or ethicists. Conversely, we rarely, if ever, go to trusted, impartial third-parties, religious authorities, or ethicists for practical advice about how to treat skin conditions; we go to doctors.

So, when Jill asks Jack, a doctor, what she is morally obligated to do, it is difficult to imagine that she is seeking moral advice. But even if we can get ourselves to imagine that she is, it is implausible that Jack's response to her question serves as good evidence for objectivism. My explanation for why Jack tells Jill that she is morally obligated to prescribe drug B is much more

³⁵ I am grateful to Anne Marie Hardy Spelman for pointing out that some applied ethicists (e.g., Peter Singer) might be exceptions to this rule.

plausible than one on which his claim is, strictly speaking, true, especially given the fact that Jack would not say that Jill was morally obligated to prescribe drug B in *Jack's Retrospective Judgment*.

1.6.4 The Argument from Moral Advice

Fourth, one might contend that if there is no objective sense of 'moral obligation,' then it is difficult to make sense of the kinds of cases I have just described in which individuals look for help in determining what their moral obligations are. Presumably, they know what they believe their moral obligations are. What they want to know is what their moral obligations *really* are. Given that this worry is common, it gives us good reason for thinking that there is an objective sense of 'moral obligation.'

In response, I do not want to deny that we sometimes wonder what we are morally obligated to do or that we do so with some frequency. However, it is not obvious to me that when we say that we are wondering what we are morally obligated to do, we are, as a matter of fact, wondering what we are morally obligated to do. For example, I think that there are a fair number of occasions on which we say that we are wondering what we are morally obligated to do when in fact we are wondering what we *should* or *ought* to do. (*Moral Advice* is one such case.) But setting those cases aside, there remain cases in which we do wonder what we are morally obligated to do. This speaks in favor of objectivism, does it not?

Although this might seem strange if objectivism were false, I do not think it actually speaks in favor of objectivism. Let us return to the two kinds of cases in which people seek moral advice. In the first, the person seeking the advice is worried that he is biased in favor of one course of action over others, so he lays out the details of his situation as he sees them and asks a less-biased third party what he is obligated to do given those details. Note that in this case, the person seeking the

advice is not operating under the assumption that his moral obligations are independent of his epistemic position. Instead, insofar as he lays out the details of his situation as he sees them, he seems to be operating under the assumption that his moral obligations are *not* independent of his epistemic position. Thus, cases like this do not speak in favor of objectivism and may even speak in favor of subjectivism or prospectivism.

But now consider the other kind of case in which people seek moral advice. In that case, the person seeking the advice seeks out a religious authority or ethicist in hopes that he or she will be able to provide him or her with insight into the nature of morality. In doing so, it is important to note that the person seeking the advice does not obviously take a position with respect to whether his moral obligations depend on his epistemic position. Presumably, he is open to the possibility that the religious authority or ethicist will tell him that his moral obligations depend on his epistemic position! There will, of course, be people, both those who seek advice as well as those who think themselves in a position to give it, who think that an agent's moral obligations are completely independent of his epistemic position, but these people could, of course, be wrong. In fact, that is precisely what *The Drug Example* and *The Bomb Example* are supposed to show, that those who think that one's moral obligations are completely independent of one's epistemic position are wrong.

In the end, therefore, the fact that we sometimes wonder what we are morally obligated to do does not give us a reason to be objectivists or even weak objectivists.

1.6.5 The Argument from Moral Monsters

Finally, consider someone like Heinrich Himmler. Let us make the plausible assumption that he was a thoughtful person with abhorrent moral views. In particular, let us assume that even

after reflection, he believed that his facilitating genocide was at least morally permissible, if not morally obligatory. One might think that if objectivism is false, then we cannot truly say that Himmler violated a moral obligation in facilitating genocide. At most, we can truly say that *there is a sense in which* he violated a moral obligation in facilitating genocide. But surely we can truly say at least that. Therefore, there must be an objective sense of ‘moral obligation.’³⁶

This, I think, is the best argument that can be given for objectivism, but even then, I do not think it is a particularly good one.³⁷ It is, of course, plausible that Himmler violated a moral obligation in facilitating genocide. But that does nothing to weaken our intuition that, in *The Drug Example*, Jill did not violate a moral obligation in prescribing drug A. So, the most we can get from this argument is the conclusion that there is an objective sense of ‘moral obligation.’ We cannot get anywhere near the truth of objectivism.

But I want to suggest that Himmler’s case cannot even get us to the conclusion that there is an objective sense of ‘moral obligation,’ though the argument for this conclusion is complex. While a thorough treatment of this question would require its own chapter, here, I will be satisfied to simply point my reader in the direction of a response.

First, note that we do not need an objective sense of ‘moral obligation’ in order to say what objectivists want to say, namely that Himmler violated a moral obligation in facilitating genocide. For on the plausible assumption that Himmler’s evidence suggested that his facilitating genocide *was not* morally permissible, prospectivism, the view that one’s moral obligations depend on one’s

³⁶ Zimmerman 2008: 14 raises this worry for subjectivism, but I am grateful to Alastair Norcross for noting that this worry has the potential to threaten prospectivism as well.

³⁷ Graham 2010 presents two arguments for objectivism. Of the two, one seems to be better than the other, and the strength of that better argument seems to depend on features it shares in common with this one. I respond to both of Graham’s arguments in Chapter 3.

evidence, entails that Himmler violated a moral obligation in facilitating genocide.³⁸ Thus, we do not need an objective sense of ‘moral obligation’ to say what we want to say about Himmler’s case; a prospective sense will do.

It is more difficult to deny that there is an objective sense of ‘moral obligation,’ however, if we assume that Himmler’s evidence suggested that his facilitating genocide *was* morally permissible. For even on this assumption, it still seems (at least to some) that Himmler did, in fact, violate a moral obligation in facilitating genocide. This suggests that there is an objective sense of ‘moral obligation.’

In response to this further argument, I should emphasize that while it may, in principle, be *possible* for someone’s evidence to suggest that his or her facilitating genocide is morally permissible, it is not obvious that any actual adult human being’s evidence ever has, as a matter of fact, suggested that his or her facilitating genocide was morally permissible. It is one thing for someone to *think* that his or her evidence suggests that his or her facilitating genocide is morally permissible. It is another thing for someone’s evidence to actually do that.

But let us assume that Himmler’s evidence did suggest that his facilitating genocide was morally permissible. At this point, we should distinguish between two kinds of cases, those in which Himmler is epistemically *responsible* (i.e., formed his beliefs responsibly) and those in which Himmler is epistemically *irresponsible* (i.e., formed his beliefs irresponsibly). In cases where Himmler is epistemically *responsible*, the claim that he violated a moral obligation in facilitating genocide is not particularly plausible. For example, in a case in which absolutely all of Himmler’s evidence speaks

³⁸ Subjectivists cannot get this result, but I do not think this refutes subjectivism. Because defending subjectivism is beside my purposes in this chapter, I will leave this task as an exercise for my reader.

in favor of his facilitating genocide, it does not seem that Himmler violates a moral obligation in doing so.

In those cases in which Himmler is epistemically *irresponsible*, however, the claim that he violated a moral obligation in facilitating genocide remains plausible. So, for example, in a case in which Himmler ignores or actively avoids evidence against the moral permissibility of facilitating genocide, it seems that Himmler's facilitating genocide *does* violate a moral obligation. This, of course, is a problem for prospectivists insofar as their view entails that in a case like this, Himmler's facilitating genocide does *not* violate a moral obligation.

In response, prospectivists can argue that an epistemically irresponsible version of Himmler *seems* to violate a moral obligation in facilitating genocide, not because he violates a moral obligation *in facilitating genocide*, but because he violates a moral obligation earlier on, *in being epistemically irresponsible*. Then, prospectivists can argue that while Himmler's epistemically irresponsibility makes his facilitating genocide *blameworthy*, it does not make his facilitating genocide *morally wrong*.

To see how this argument would go, consider the following case:

Negligent Doctor

Edward, a dermatologist at a large practice, is responsible for keeping up with the latest skin research. Unfortunately, however, he has a hard time doing that. He loves watching sports and often prioritizes that over reading scientific studies. Having fallen significantly behind, Edward gets a patient with a relatively rare skin condition, condition X. He remembers learning about condition X in medical school, and at that time, all the evidence suggested that drug D would benefit those with condition X while drug E would harm them. Edward believes, in accordance with his evidence, that he cannot comb through all the skin research he has missed before treating his patient, so he goes ahead and prescribes drug D. Had Edward

kept up with the latest skin research, however, he would have known that whereas drug D will harm this particular patient, drug E would have benefitted her.³⁹

Regarding *Negligent Doctor*, prospectivists can say that Edward violated a moral obligation in failing to keep up with the latest skin research. Moreover, they can say that Edward is blameworthy for prescribing drug D on account of the fact that had he fulfilled his earlier moral obligations (i.e., had he kept up with the latest skin research), his evidence would have spoken against his prescribing drug D. So, they can say that Edward is blameworthy for prescribing drug D on account of the fact that he violated his earlier moral obligation to keep up with the latest skin research. What prospectivists cannot say, however, is that Edward violated a moral obligation in prescribing drug D, and it is this fact, that prospectivists cannot say that Edward violated a moral obligation in prescribing drug D, that weak objectivists and ambiguity theorists would use to support their claim that there is an objective sense of ‘moral obligation.’

But prospectivists can cast doubt on the claim that Edward violated a moral obligation in prescribing drug D. “It is true,” they can admit, “that Edward violated a moral obligation, but the moral obligation he violated was not a moral obligation to prescribe drug E, but an earlier moral obligation to keep up with the latest skin research. Consider the situation in which Edward found himself. He believed, in accordance with his evidence, that drug D would benefit his patient and that drug E would harm her. Certainly a person in that position is morally obligated to prescribe drug D. We are, of course, not suggesting that Edward’s prescribing drug D was best or that Edward does not deserve blame for prescribing drug D. We are simply suggesting that Edward did not violate a moral obligation when he did so.”

³⁹ This example is based on a case that Phil Robichaud and Jan Willem Wieland discuss in their paper “Blame Transfer,” which they presented at the 2015 Rocky Mountain Ethics Congress, and their case is based on an earlier case by Smith 1983: 543.

If the prospectivist's description of Edward's case is plausible, and Edward's case is sufficiently analogous to that of an epistemically irresponsible version of Himmler,⁴⁰ then we do not need an objective sense of 'moral obligation' in order to give a plausible description of Himmler's case. The question is whether the prospectivist's description of Edward's case is plausible.

I think it is. While there is something plausible about the claim that an epistemically irresponsible version of Himmler violated a moral obligation in facilitating genocide, it is not obvious that its plausibility is attributable to the intuition that an epistemically irresponsible version of Himmler *violated a moral obligation in facilitating genocide* rather than the intuition that an

⁴⁰Note that while I do not think the two cases are perfectly analogous, I *do* think that if the prospectivist can make plausible the claim that Edward does not violate a moral obligation in prescribing drug D, then he or she can also make plausible the claim that an epistemically irresponsible version of Himmler does not violate a moral obligation in facilitating genocide.

Here is why. As I have formulated prospectivism, it is the view that an agent's moral obligations depend, to some extent, on his or her evidence. There are two ways one might flesh out this view. First, there is *half-hearted prospectivism*. On this view, an agent's moral obligations depend on his or her evidence about the descriptive facts combined with *the actual* moral facts. So, on a utilitarian version of half-hearted prospectivism, an agent whose evidence suggests that killing a particular innocent person will save three other innocent people would be morally obligated to kill that innocent person, even if his or her evidence also suggests that killing innocent people is always morally wrong. In contrast to this view, there is *full-blown prospectivism* on which an agent's moral obligations depend on his or her evidence about the descriptive facts combined with *his or her evidence about* the moral facts. So, on full-blown prospectivism, an agent whose evidence suggests that killing a particular innocent person will save three others would *not* be morally obligated to kill that innocent person (and, in fact, would be morally obligated *not* to kill that innocent person) if his or her evidence also suggests that killing innocent people is always morally wrong.

In *Negligent Doctor*, Edward's epistemic irresponsibility causes his evidence about the *descriptive facts* to be defective. Thus, regardless of whether a prospectivist goes in for half-hearted prospectivism or full-blown prospectivism, he or she will have to say that Edward did not violate a moral obligation in prescribing drug D. So, insofar as prospectivists can make it plausible that Edward did not violate a moral obligation in prescribing drug D, they thereby show that both versions of prospectivism are defensible.

But the prospectivist need not show that both versions of prospectivism are defensible in order to respond to the problem that arises when we consider an epistemically irresponsible version of Himmler, for in that case, Himmler's epistemic irresponsibility causes his evidence about the *moral facts* to be defective. Thus, only full-blown prospectivists will have to say that Himmler did not violate a moral obligation in facilitating genocide. Nevertheless, because I think that full-blown prospectivism is the more plausible version of prospectivism, I have defended that version of prospectivism.

epistemically irresponsible version of Himmler *is blameworthy for* facilitating genocide. It is difficult, I submit, for our intuitions to track this subtle distinction.

If I am right, then we do not need to be weak objectivists or ambiguity theorists in order to give a plausible account of the behavior of an epistemically irresponsible version of Himmler. And if that is true, then neither weak objectivists, weak prospectivists, nor ambiguity theorists can use this case to establish the existence of an objective sense of ‘moral obligation.’

1.7 Conclusion

I began this chapter by presenting Jackson’s well-known drug example and corresponding argument against objective versions of consequentialism. After noting that Jackson’s argument actually threatens any plausible version of objectivism, I outlined how objectivists have responded to Jackson’s argument and acknowledged that Jackson’s argument falls prey to that response. At that point, however, I suggested that we can construct a better argument against objectivism if, rather than attending to what one *should* do in *The Drug Example*, we attend to what one is *morally obligated* to do in that example. After raising a few objections to my Jackson-style argument, I showed that all of them fail. Then, after suggesting not only that objectivism is false, but also that there is no objective sense of ‘moral obligation,’ I considered a number of arguments for the existence of such a sense and showed that none of them are compelling.

In doing these things, I take myself to have established three things. First, and most importantly, I take myself to have shown that not even *moral shoulds* or *oughts* entail moral obligations. Consequently, *contra* Zimmerman and Parfit,⁴¹ we should not say that an agent *morally*

⁴¹ See Zimmerman 2006, 2008, 2014, and Parfit 2011.

should or *ought* to ϕ when we mean to say that he or she is *morally obligated* to ϕ . Similarly, *contra* Feldman, Jackson, and Dorsey,⁴² we should not use our intuitions about what agents *morally should* or *ought* to do to draw conclusions about what they are morally obligated to do. Second, I take myself to have shown, *contra* Moore and Graham,⁴³ that we should reject objectivism about moral obligation. Even if there is an objective sense of ‘moral obligation,’ there is also a subjective or prospective sense of that phrase. Finally, I take myself to have shown, *contra* Sidgwick, Feldman, Gibbard, Smith, Parfit, Dorsey, Driver, and Mason,⁴⁴ that we should be skeptical that there is an objective sense of ‘moral obligation.’ Insofar as prospectivism (and potentially subjectivism) can give a plausible account of (i) ‘objective moral obligations,’ (ii) how we use the word ‘wrong,’ (iii) practical advice, (iv) moral advice, and (v) moral monsters, we should take seriously the possibility that there is no objective sense of ‘moral obligation.’

⁴² See Feldman 1986, Jackson 1991, and Dorsey 2012.

⁴³ See Moore 1912 and Graham 2010.

⁴⁴ See Sidgwick 1981/1907, Feldman 1986, Gibbard 1990, Smith 2010, Parfit 2011, Dorsey 2012, Driver 2012a, 2012b, and Mason 2013.

CHAPTER 2

MORAL OBLIGATION AND PUNISHMENT⁴⁵

2.0 Introduction

The debate discussed in Chapter 1, between objectivists, subjectivists, and prospectivists, has been especially prominent amongst consequentialists. Some consequentialists endorse objective versions of consequentialism.⁴⁶ They believe that actions are morally obligatory if and only if they maximize *actual* value. Others endorse subjective versions of consequentialism.⁴⁷ They believe that actions are morally obligatory if and only if they maximize *expected* value.⁴⁸ Still others endorse prospective versions of consequentialism.⁴⁹ They believe that actions are morally obligatory if and only if they maximize *expectable* value.⁵⁰

In this chapter, I present another argument against objective versions of consequentialism (hereafter, simply *objectivism* unless otherwise noted). To do this, I argue first that the concept of an obligation is such that those who violate *obligations* are *subject to punishment*. Then, I argue that if those

⁴⁵ I am grateful to Chris Heathwood for his helpful feedback on earlier drafts of this chapter as well as to those graduate students and faculty at CU Boulder who gave me valuable comments on this chapter when I presented it in 2015.

⁴⁶ See, e.g., Sidgwick 1981/1907, Moore 1912, Feldman 2006, Graham 2010, and Driver 2012a, 2012b.

⁴⁷ See, e.g., Jackson 1991. Prichard 2002/1932 and Ross 1939 are not consequentialists, but they are subjectivists.

⁴⁸ The *expected* value of an agent's action is the product of how likely *she believes* it is to have certain outcomes and how valuable *she believes* those outcomes to be.

⁴⁹ See, e.g., Mason 2013, 2014. Zimmerman 2006, 2008, 2014, though not obviously a consequentialist, is a prospectivist.

⁵⁰ The *expectable* value of an agent's action is the product of how likely *her evidence suggests* it is to have certain outcomes and how valuable *her evidence suggests* those outcomes are.

who violate obligations are subject to punishment, then those who violate *moral* obligations *deserve* punishment. Finally, I argue that if those who violate moral obligations deserve punishment, then objectivism is false.

I am not the first person to suggest that agents who violate obligations are subject to punishment. In Plato's *Euthyphro*, for example, Socrates notes that those who are accused of committing some crime do not deny that wrongdoers are to be punished but, instead, deny that they have done anything wrong.⁵¹ Similarly, in *Utilitarianism*, John Stuart Mill writes, "We do not call anything wrong, unless we mean to imply that a person ought to be punished in some way or other for doing it; if not by law, by the opinion of his fellow creatures; if not by opinion, by the reproaches of his own conscience."⁵² But whereas neither Plato's Socrates nor Mill argues for these claims, I do, in hopes that consequentialists will take them more seriously. For if consequentialists *do* take them more seriously, they will see that they must reject objective version of consequentialism in favor of subjective or prospective versions of consequentialism.

Of course, it is also worth noting that this same debate could arise between adherents of deontology or virtue ethics. We can imagine, for example, deontologists disagreeing about our moral obligations with respect to killing. There could be objectivists, according to whom it is morally wrong to kill others, subjectivists, according to whom it is morally wrong to perform acts that you believe will kill others, and prospectivists, according to it is morally wrong to perform acts that your evidence suggests will kill others. So, while my attention in this chapter will be focused on arguing against objective versions of consequentialism, my argument has broader implications.

⁵¹ *Euthyphro* 8d.

⁵² Mill 2001/1861: 48-49.

2.1 An Argument against Objectivism

Imagine a doctor who believes that drug A will cure her patient. Let us also assume that she is mistaken. As a matter of fact, it will kill him. Unfortunately, she prescribes drug A, and her patient dies. At this point, we can ask whether the doctor has violated a moral obligation.

Before answering this question, you might want more information. For example, you might wonder whether the doctor could have known that drug A would kill her patient. And if she could have, you might wonder whether the doctor's belief that drug A would cure her patient was attributable to negligence.

Let us assume that although the doctor could have known that drug A would kill her patient, her belief that drug A would cure her patient was not attributable to negligence.⁵³ Given this information, you might conclude that the doctor did not violate a moral obligation. If so, you will be sympathetic to the following argument.⁵⁴

The Simple Argument against Objectivism

- (1) If objectivism is true, then the doctor has violated a moral obligation.
- (2) The doctor has not violated a moral obligation.
- ∴ (3) Objectivism is false.

This argument is valid, and assuming that the doctor's patient would have been better off had she not prescribed drug A, (1) is true. Moreover, (2) is intuitively plausible. For even if it seems that, after hearing of her patient's death, the doctor would admit that she *ought* not to have prescribed

⁵³ This might be the case if, for example, (i) drug A is only deadly to a very small percentage of the population and (ii) this fact has only recently been discovered by a small research team whose findings have yet to be published.

⁵⁴ This argument is not original to me. A similar argument appears in Moore 1912: 191-192, in Ross 1939: 163-164, and in numerous articles and books since then. What is unique to me is the way in which I defend it.

drug A, it does not seem that she would admit that she was *morally obligated* not to have prescribed drug A. Objectivism appears to be sunk.

In response to this argument, objectivists like G. E. Moore and Peter Graham have attempted to deny (2). In particular, they have suggested that (2) owes its plausibility to the plausibility of the following argument.⁵⁵

The Argument from Blameworthiness

- (4) The doctor is not blameworthy.
- (5) If an agent is not blameworthy, then that agent has not violated a moral obligation.
- ∴ (2) The doctor has not violated a moral obligation.

In *Ethics*, for example, Moore responds to an argument like this one by admitting (4) but denying (5). More specifically, while acknowledging that someone like our doctor is not blameworthy, Moore contends that there is no reason to think that an agent who is not blameworthy has not violated a moral obligation.⁵⁶ Thus, on Moore's view, the fact that the doctor is not blameworthy does not give us a reason to think that she has not violated a moral obligation.

It is worth noting that Moore has not given us a reason to deny (5). Instead, he has claimed that there are no good reasons to accept it. Thus, his objection to (5) is relatively weak. If we can come up with a good reason to accept (5), we will have provided a sufficient response to Moore's objection.⁵⁷

My goal in this chapter, however, is not to defend (5) but to note that the argument that Moore and Graham attribute to their opponents is not the only argument that their opponents can

⁵⁵ Moore 1912: 192-193 and Graham 2010: 93-94.

⁵⁶ Moore 1912: 192-193.

⁵⁷ One might think that the intuitive plausibility of (5) gives us a good reason to accept (5), even in light of Moore's objection. While I am sympathetic to that view, I worry that such a victory comes a bit too cheap.

give for that premise. There are, in fact, a myriad of additional arguments that can be given for (2), and here I want to raise one that is closely related to the one mentioned by Moore and Graham.

That argument is as follows:

The Argument from Punishment

- (6) The doctor does not deserve punishment.
- (7) If an agent does not deserve punishment, then that agent has not violated a moral obligation.
- ∴ (2) The doctor has not violated a moral obligation.

Unlike the previous argument, this argument does not appeal to the fact that the doctor is blameless to support the conclusion that she has not violated a moral obligation. Instead, it appeals to the fact that the doctor does not deserve punishment. I imagine that this claim, (6), is relatively uncontroversial. The second premise of the argument, however, (7), almost certainly is controversial, so I will argue for it below. To begin, I argue that agents who violate *obligations* are *subject to punishment*. Then, I argue that if agents who violate obligations are subject to punishment, then agents who violate *moral* obligations *deserve* punishment. If those two arguments are compelling, then we will have a good reason to accept (2) and to reject objectivism.

2.2 An Argument for (7)

My argument for (7) is as follows:

- (8) If an agent has violated an obligation, then that agent is subject to punishment.
- (9) If it is true that if an agent has violated an obligation, then that agent is subject to punishment, then it is true that if an agent has violated a moral obligation, then that agent deserves punishment.
- ∴ (10) If an agent has violated a moral obligation, then that agent deserves punishment.

∴ (7) If an agent does not deserve punishment, then that agent has not violated a moral obligation.

If (8) and (9) are true, then (10) must be true, and if (10) is true, its contrapositive, (7), must also be true. In Section 2.2.1, I argue for (8). In Section 2.2.2, I defend (8) from objections. And in 2.2.3, I argue for (9).

2.2.1 An Argument for (8)

As rational beings, we have moral obligations, but we have other obligations as well. For example, as children of particular parents, we have certain filial obligations. As citizens of particular states, we have certain legal obligations. And finally, many believe that we, as “children of God,” have certain religious obligations. To better understand what it means to have a *moral* obligation, it is worth considering what it means to have an obligation more generally, and to do that, it is worth considering what filial, legal, and religious obligations have in common.

One thing these three kinds of obligations have in common is that they must be met. Of course, they may have other things in common, but for now, let us focus our attention on this feature. Given that obligations must be met, we might wonder, “What does it mean to say that obligations must be met?” Below, I provide three examples of *non-moral* obligations, each of which suggests that to say that an obligation must be met is to say that it must be met in order to avoid being subject to punishment.⁵⁸

⁵⁸ If we looked at a wider range of cases, we might conclude that obligations must be met *either* to avoid being subject to punishment *or to be subject to receive some reward*. However, given that it is possible to understand one’s not being subject to receive some reward as a kind of punishment, I will simply talk as if obligations must be met in order to avoid being subject to punishment.

Consider the following example. A child is sitting at a dinner table. She has finished everything on her plate except for her Brussels sprouts. “You must finish your Brussels sprouts,” her parents say. “Or else?” the daughter asks. “Or else you will spend the rest of the evening in your room,” her parents reply. This child’s obligation to finish her Brussels sprouts would seem to be a paradigm case of a filial obligation, and the obligation is this: she must finish her Brussels sprouts *or else* filial norms permit her parents to require her to spend the rest of the evening in her room.

Here is another example. In Colorado, theft is a class 1 petty offense provided that the stolen item is valued at less than fifty dollars.⁵⁹ So, if someone steals a DVD from an electronics store in Colorado, he or she will have committed a class 1 petty offense, which is punishable by a fine of less than \$500 or imprisonment for up to 6 months.⁶⁰ This would seem to be a paradigm case of a legal obligation, and the obligation is this: one must refrain from stealing anything under fifty dollars in Colorado *or else* legal norms permit the state of Colorado to require one to pay a fine of less than \$500 or imprison one for up to 6 months.

Here is a final example. According to the Catholic Church, All Saints’ Day is a holy day of obligation. Consequently, one must attend mass on All Saints’ Day. Those who deliberately fail to do so commit mortal sin.⁶¹ It seems to me that this is a paradigm case of a religious obligation, and the obligation is this: one must attend mass on All Saints’ Day *or else* religious norms permit God to send one to hell.

⁵⁹ C.R.S. § 18-4-401.

⁶⁰ C.R.S. § 18-1.3-503

⁶¹ A mortal sin is such that “If it is not redeemed by repentance and God's forgiveness, it causes exclusion from Christ's kingdom and the eternal death of hell” (*Catechism of the Catholic Church* 1861).

Each of the obligations I have mentioned is such that if an agent has violated it, then that agent is subject to punishment. This suggests that (8) is true; an agent's having violated an obligation entails that agent's being subject to punishment.

Before moving on, I need to clarify what I mean when I say that an agent is subject to punishment. First, when I say that an agent is subject to punishment, I do not mean to suggest that that agent will be punished. Just because someone violates a legal obligation (e.g., by smoking marijuana where it is illegal to do so), it does not follow that he or she will be punished. Second, when I say that an agent is subject to punishment, I do not mean to suggest that anyone is subjectively justified in punishing him or her. Just because a particular agent smokes marijuana illegally, it does not follow that anyone is subjectively justified in punishing him or her, for an agent might smoke marijuana in private (while being too intoxicated to remember doing so). In such a case, no one would be subjectively justified in punishing that agent for smoking marijuana, not even the agent him or herself. Finally, when I say that an agent is subject to punishment, I do not mean to suggest that that agent deserves punishment. Just because someone violates a legal obligation by smoking marijuana, it does not follow that he or she deserves punishment, for the law against smoking marijuana might be unjust.

Instead, when I say that an agent is subject to punishment, I mean to suggest that some norm permits someone (or something like the state) to punish the agent. For example, if someone is subject to punishment for stealing a DVD from an electronics store in Colorado, it follows that some norm permits someone (or something like the state) to punish him or her for stealing a DVD from an electronics store in Colorado. With these distinctions in mind, I contend that we should accept (8), the claim that if an agent violates an obligation, then that agent is subject to punishment.

2.2.2 In Defense of (8)

Although (8) seems innocent enough, there are a number of ways in which one might object to it. In this section, I will consider and respond to those objections.

First, one might try to provide a counterexample to (8) by identifying an obligation that is such that those who violate it are not subject to punishment. For example, a recent *Snapple* commercial suggests that it is against the law to sing off-key in North Carolina. If that is right, then those in North Carolina are obligated either to sing on-key or to not sing at all. But, of course, we cannot imagine the state of North Carolina punishing someone for singing off-key. This appears to be a case in which an agent can violate an obligation without being subject to punishment, in which case (8) is false.

The strength of this objection depends on a particular way of understanding what it means for an agent to be subject to punishment. In particular, it assumes that if those who violate a particular obligation are never punished, then those who violate that obligation are not subject to punishment. This, however, does not follow from my account of what it means for an agent to be subject to punishment. As indicated above, to say that an agent who violates a moral obligation is subject to punishment is to say that some norm permits someone (or something like the state) to punish him or her. So, the fact that those who violate a particular obligation are never punished does not entail that there is no norm permitting someone (or something like the state) to punish him or her, and therefore the fact that those who violate a particular obligation are never punished does not entail that those who violate that obligation are not subject to punishment.

In order for this kind of objection to work, it must be the case that those in North Carolina are obligated not to sing off-key and, at the same time, that there is a norm prohibiting anyone (or anything, like the state) from punishing agents who violate that obligation. To my ear, this sounds like an impossibility. Even if there were a North Carolina law against singing off-key that was never

enforced, that does not entail that those who violate it are not subject to punishment, for the fact that those who violate it are not punished does not entail that they could not be.

But imagine there is a North Carolina law against singing off-key that state officials are prohibited from enforcing. In that case, I am inclined to say that those in North Carolina are not obligated not to sing off-key. If someone tells me that I am obligated not to sing off-key and then goes on to tell me that my singing off-key cannot be met with any sort of punishment, it would be difficult for me to make sense of his or her claim. In such a case, there does not seem to be any meaningful sense in which the I *must* not sing off-key, in which case there does not seem to be any meaningful sense in which I am *obligated* not to sing off-key. Similarly, we could imagine two parents insisting that their daughter eat her broccoli, but if there is no punishment associated with their daughter's failure to eat her broccoli, then I am hesitant to say that she is so obligated. The parents may have given their daughter a command or an imperative, but they have not put her under an obligation. This example suggests what I take to be the correct response to this objection, namely that if there is a rule or law the violation of which is not subject to punishment, then although that rule or law may be a kind of command or imperative, it does not generate an obligation.

Second, one might object to (8) by arguing that although the concept of *obligation* appears to be connected to the concept of *punishment* in the way I have described, it is more clearly connected to some other concept in a way that prevents it from being connected to the concept of *punishment* in the way I have described.

For example, one might argue that obligatory actions are the kinds of actions that others are permitted to force one to perform and that prohibited actions are the kinds of actions that others are permitted to prevent one from performing.⁶² On this view, a third party is permitted to

⁶²Eric Chwang has suggested this view in conversation.

force an agent to perform some action, ϕ , if and only if the agent is obligated to ϕ . Additionally, a third party is permitted to *prevent* an agent from performing some action, ϕ , if and only if that agent is prohibited from ϕ -ing.

This view, if true, would rule out the truth of the view for which I have been arguing. Moreover, this view seems promising, for it has the potential to explain why we think a third party would be permitted to prevent someone like the doctor in our original example from prescribing drug A, namely, because the doctor was prohibited from prescribing drug A. However, I think we should reject this view because it admits of clear counterexamples. While there are some cases in which it is permissible to force someone to do what he or she is obligated to do (or to prevent someone from doing what he or she is prohibited from doing), there are other cases in which it is impermissible to do so. For example, imagine that an agent, Joe, is legally obligated to pay ten thousand dollars in taxes to the federal government. It does not follow from this that anyone is legally permitted to steal ten thousand dollars from Joe and send it to the federal government on Joe's behalf. Nor would it be legally permissible for someone to physically force Joe to send ten thousand dollars to the government. Maybe more importantly, however, even if it were legally permissible for someone to do these things, it is conceivable that things could be otherwise. That is, the laws could be set up such that even though Joe is legally obligated to pay ten thousand dollars in taxes to the federal government, no one is legally permitted to force him to do that. What this illustrates is that there is no conceptual connection between an agent's being obligated to perform some action and its being permissible for a third party to force that agent to perform that action.

Furthermore, we do not need to adopt this view in order to explain why a third party would be permitted to prevent someone like the doctor in our original example from prescribing drug A. Instead, we could adopt the more plausible view that a third party, T , is permitted to force an agent, S , to perform some action, ϕ , if and only if T 's forcing S to ϕ maximizes expectable utility.

On this view, a third party would be permitted to stop the doctor in our original example from prescribing drug A, provided that he has good reasons to think that the doctor's prescribing drug A will harm her patient. Additionally, this view prevents a third party from stopping the doctor in our original example if he lacks good reasons to think that the doctor's prescribing drug A will harm her patient. This strikes me as the right result.

Finally, an opponent might simply deny that the concept of *obligation* is connected to the concept of *punishment* in the way I have described, or any other concept for that matter. On this view, the concept of *obligation* is primitive and the kind of argument I have given is misguided. Although one could certainly take this line, I think my view superior for at least two reasons. First, I think my view is superior simply because the concept of *obligation* does seem to be connected to the concept of *punishment* in the way I have described. Second, I think my view is superior insofar as, all else equal, it is preferable to be able to give an account of a concept in terms of other concepts rather than to take it as primitive.

In this section, I have argued that there are no good counterexamples to (8), the claim that if an agent has violated an obligation, then that agent is subject to punishment. Additionally, I have argued that there are no plausible stories to tell on which the concept of *obligation* is clearly connected to some concept in a way that prevents it from being connected to the concept of *punishment* in the way I have described. While one could simply deny that the concept of *obligation* is connected to the concept of *punishment* in the way I have described, doing so comes with significant costs.

2.2.3 An Argument for (9)

My argument for (7), the claim if an agent does not deserve punishment, then that agent has not violated a moral obligation, is valid and (8), the claim that if an agent has violated an obligation, then that agent is subject to punishment, is, as I have just argued, true. So, the last thing I need to do to establish (7), which depends on the truth of (8) and (9), is to argue for (9). Roughly, (9) says that if an agent's having violated an *obligation* entails that agent's being *subject to punishment*, then an agent's having violated a *moral* obligation entails that agent's *deserving* punishment. This is just the claim that (8) entails (10). So, below, I argue that (8) entails (10).

The first step in the argument is to note that a moral obligation is a kind of obligation, so (8), the claim that if an agent has violated an obligation, then that agent is subject to punishment, entails that if an agent has violated a *moral* obligation, then that agent is subject to punishment.

The second step in the argument is to note that although the fact that someone violates an obligation does not entail that he or she deserves punishment, the fact that someone violates a *just* obligation *does* entail that he or she deserves punishment. Why think this? Well, the only reason that one's violating an obligation does not entail one's deserving punishment is that some obligations are unjust. In particular, the obligations one must meet on account of being the child of certain parents or the obligations one must meet on account of being a citizen of a certain state are not necessarily just. Consider, for example, seatbelt laws. We can imagine someone complaining that seatbelt laws are unjust and therefore that those who violate the seatbelt laws do not deserve punishment. But now imagine that we convince this person that seatbelt laws are just. Would he continue to think that those who violate seatbelt laws do not deserve punishment? Certainly not, for to convince someone that certain laws are just is to convince that person that those who violate those laws deserve punishment. This suggests that if an agent violates a *just* obligation, then that agent deserves punishment.

Finally, the third step in the argument is to note that one's moral obligations, unlike one's filial, legal, or religious obligations, are necessarily just. So, given that all moral obligations are just, (8), the claim that if an agent has violated an obligation, then that agent is subject to punishment, entails (10):

- (10) If an agent has violated a moral obligation, then that agent *deserves* punishment.

Before moving on, I should note that just because an agent deserves punishment, it does not follow that it is appropriate for any particular person to punish that agent. This is true not only because it could be the case that no one knows that that agent deserves punishment, but also because it could be the case that punishing him or her will have bad consequences. So, even though the fact that an agent has violated a moral obligation entails that he or she deserves punishment, it does not follow from the fact that an agent has violated a moral obligation that it is appropriate to punish him or her. Imagine, for example, that an agent, Valentina, is on her deathbed when we find out that, years ago, Valentina stole thousands of dollars from Exxon Mobil. While we *could* punish Valentina in a number of ways, none of them would have good consequences. Imprisoning her would be expensive and impractical, and because she is near death, fining her would only hurt her children, who, let us stipulate, did not benefit from Valentina's thievery. In cases like this, it is inappropriate to punish one even if one deserves it.

In this section, 2.2.3, I have argued that (9) is true, that if an agent's having violated an obligation entails that agent's being subject to punishment, then an agent's having violated a moral obligation entails that agent's deserving punishment. In Section 2.2.1, I argued that (8) is true, that an agent's having violated an obligation entails that agent's being subject to punishment. If both of these arguments are successful, it follows that an agent's having violated a moral obligation

entails that agent's deserving punishment, which is (10). And if (10) is true, then its contrapositive, (7), is also true. If an agent does not deserve punishment, then that agent has not violated a moral obligation.

2.3 Objections to the Argument for (2)

Earlier, I provided an argument for (2), the claim that the doctor has not violated a moral obligation. That argument depended on (6), the claim that the doctor does not deserve punishment, and (7), the claim that if an agent does not deserve punishment, then that agent has not violated a moral obligation. Because (6) seemed obviously true, I took it for granted and spent my time arguing for (7). In the course of arguing for (7), however, I suggested that even if an agent has violated a moral obligation and therefore deserves punishment, it does not follow that it is appropriate to punish that agent because it could be the case that no one knows that that agent deserves punishment or because it could be the case that punishing that agent will have bad consequences. This may cast doubt on (6), the claim that the doctor does not deserve punishment.

In particular, an objectivist might contend that the plausibility of (6), the claim that the doctor does not deserve punishment, depends on the plausibility of the claim that it is not appropriate to punish the doctor. And, as I have suggested, that it is not appropriate to punish an agent does not entail that that agent does not deserve punishment. Thus, an objectivist might deny (6) by contending that although it is not appropriate to punish the doctor, she is deserving of punishment, nevertheless. The thought here has to be that although objectivism is true and, therefore, that the doctor has violated a moral obligation, it is not appropriate to punish her either because no one knows that she has violated a moral obligation or because everyone knows that punishing her will have bad consequences.

In response, I should note that the plausibility of the claim that it is not appropriate to punish the doctor cannot be attributed to the fact that no one knows that the doctor has violated a moral obligation, for if objectivism is true, then it is obvious that the doctor has violated a moral obligation on account of the fact that what she did did not maximize actual value.

It does, however, seem like the plausibility of the claim that it is not appropriate to punish the doctor might be attributable to the fact that everyone knows that punishing her will have bad consequences, for it is hard to see how anything good could come of punishing the doctor. Is objectivism therefore vindicated? I do not believe so. Although it is presumably true, all else being equal, that punishing the doctor will have bad consequences, that is not why we are hesitant to punish her. Instead, we are hesitant to punish her because we do not think she deserves punishment. To illustrate this, we might alter the original case such that it is obvious that punishing the doctor will have good consequences. For instance, we might imagine that the doctor is unlikeable and that everyone wants to see her punished. Even in this case, however, we would be hesitant to punish the doctor. Thus, the plausibility of the claim that it is not appropriate to punish the doctor cannot be attributed to the fact that everyone knows that punishing her will have bad consequences

In the end, then, it seems like the plausibility of the claim that it is not appropriate to punish the doctor is attributable to the claim that *the doctor does not deserve punishment*. That is (6). Thus, we should accept (6) after all.

2.4 Conclusion

In the example I described in the beginning of Section 2.1, a doctor prescribes drug A, which unexpectedly killed her patient. If objectivism is true, then the doctor violated a moral

obligation in prescribing drug A, but I have suggested that that is false. In defense of this suggestion, I argued that if an agent has violated a moral obligation, then that agent deserves punishment. So, given that the doctor does not deserve punishment, it follows that the doctor did not violate a moral obligation. Thus, objectivism must be false.

Provided that those who violate an obligation are subject to punishment, I do not see any way to deny this conclusion, and I do not see any way to deny that those who violate their obligations are subject to punishment. Thus, I believe we should abandon objectivism in favor of subjectivism or prospectivism.

CHAPTER 3

AGAINST OBJECTIVISM ABOUT MORAL OBLIGATION⁶³

3.0 Introduction

Objectivism, the view that an agent's moral obligations do not depend on her beliefs or her evidence, is the dominant one in normative ethics.⁶⁴ It has, however, come under fire of late.⁶⁵ In "In Defense of Objectivism about Moral Obligation," Peter A. Graham attempts to defend objectivism from its opponents. In this chapter, I contend that his attempt is unsuccessful.

In Sections 3.1-3.2, I discuss a common argument *against* objectivism, which I mentioned briefly in the previous chapter. Although Graham thinks he has a good objection to that argument, I show that his objection fails. In Sections 3.3-3.9, I discuss Graham's arguments *for* objectivism and argue that neither of them gives us a good reason to be objectivists (as compared to subjectivists or prospectivists). In the end, I conclude that careful reflection on Graham's paper leaves us with more reason to reject objectivism than to accept it.

⁶³ I am grateful to Peter A. Graham, Chris Heathwood, and two anonymous referees for their helpful comments on earlier drafts of this chapter.

⁶⁴ Consequentialism, Kantianism, and virtue ethics, as they have traditionally been formulated, all entail objectivism.

⁶⁵ See, e.g., Zimmerman 2006, 2008, 2014, and Mason 2013.

3.1 The Simple Argument against Objectivism

As I mentioned earlier, objectivism is compatible with a wide variety of normative theories. Objectivists could be consequentialists who hold that agents are morally obligated to maximize value, Kantians who hold that agents are morally obligated to act in accordance with the categorical imperative, virtue ethicists who hold that agents are morally obligated to act virtuously, etc. Objectivism is to be contrasted with non-objectivist views on which an agent's moral obligations depend on her beliefs or her evidence. One such view is *subjectivism*, on which an agent's moral obligations depend on her beliefs. Another such view is *prospectivism*, on which an agent's moral obligations depend on her evidence. Subjectivists and prospectivists, like objectivists, can hold a wide variety of normative theories. They could be consequentialists, Kantians, virtue ethicists, etc., in which case they would hold that agents are morally obligated to do what they believe (or what their evidence suggests) maximizes value, is in accordance with the categorical imperative, is virtuous, etc.

While objectivism is the dominant view in normative ethics, there is a common argument against it, which I will call *the simple argument against objectivism*. There are different versions of this argument, but each is based on a case in which (i) an agent believes, in accordance with his evidence, that performing a certain action is best and (ii) that agent's belief is false. The version of the simple argument that Graham discusses is based on the following case:

Syringes

Luke's patient, Nathan, has a painful, nonfatal disease. In syringe C is the cure for the disease and in syringe P is a lethal poison. Unfortunately, the labels on the syringes have, unbeknownst to Luke, been switched.⁶⁶

⁶⁶ Graham 2010: 93.

In *Syringes*, Luke believes, in accordance with his evidence, that his injecting Nathan with syringe P will cure him and, therefore, that his injecting Nathan with syringe P is best. But his injecting Nathan with syringe P will not cure Nathan; it will kill him. So, it is not best.

If objectivism is true, then neither Luke's beliefs nor his evidence is relevant to what he is morally obligated to do. So, objectivism entails that Luke's injecting Nathan with syringe C is morally obligatory. But this seems false. While it may be best for Luke to inject Nathan with syringe C, he does not seem morally obligated to do so. In fact, insofar as he believes that his injecting Nathan with syringe C will kill Nathan, it seems like his injecting Nathan with syringe C would be morally wrong. Therefore, objectivism is false.⁶⁷

3.2 Graham's Objection and a Response

Graham suggests that our judgment *that Luke's injecting Nathan with syringe C is not morally obligatory* is attributable to our intuition that Luke's failure to inject Nathan with syringe C would not be blameworthy. Then, he contends that blameworthiness and moral wrongdoing come apart and, therefore, that even if Luke's failure to inject Nathan with syringe C is not blameworthy, it might still be morally wrong. Thus, he concludes that the simple argument against objectivism fails.⁶⁸

⁶⁷ Even if this particular argument threatens only objective versions of consequentialism, we could come up with structurally analogous arguments against objective versions of Kantianism or virtue ethics. We could, for example, come up with cases in which an agent believes, in accordance with his evidence, that ϕ -ing is in accordance with the categorical imperative when, in fact, it is not, or cases in which an agent believes, in accordance with his evidence, that ϕ -ing is virtuous when, in fact, it is not. Thus, the simple argument against objectivism threatens objectivism generally.

⁶⁸ Graham 2010: 93-94.

This move goes back at least to G. E. Moore,⁶⁹ but it is too quick. It is true, *a priori*, that blameworthiness and moral wrongdoing *could* come apart, but we want to know whether they *do* come apart. In support of this claim, Graham presents two cases, a supposed case of blameless wrongdoing and a supposed case of blameworthy rightdoing. Together, they purport to show that blameworthiness is neither necessary nor sufficient for moral wrongdoing.

But Graham's cases fail to show this. Consider Graham's first case, a supposed case of blameless wrongdoing. Graham writes, "If, when leaving a party I take your jacket, mistaking it for mine, it's intuitive that I do something wrong even though I may be completely blameless."⁷⁰ If this is true, then blameworthiness is not necessary for moral wrongdoing.

In response, I want to draw attention to the familiar distinction between an agent's doing something *wrong* and an agent's doing something *morally wrong*. If I make a wrong turn on my way to a party, it does not follow that I have done anything morally wrong. I have made a mistake, but it is a morally permissible one (unless I was morally obligated not to make that wrong turn). Similarly, if you leave a party with the wrong jacket, it does not follow that you have done anything morally wrong. You have made a mistake, but it may be a morally permissible one.

There will, of course, be cases in which, in leaving a party with the wrong jacket, you do something morally wrong. Consider, for example, a case in which, while leaving a party, you take someone else's jacket, mistaking it for yours. On your way to your car, you slip your hands into its pockets and find something unfamiliar. You are surprised by your discovery, but you ignore it and drive home.

⁶⁹ Moore 1912: 192-193.

⁷⁰ Graham 2010: 94.

In this case, it is plausible that you have done something morally wrong. But it is not *your taking the wrong jacket* that is morally wrong; it is *your taking the wrong jacket even after finding something unfamiliar in its pockets* that is morally wrong. Or, more generally, it is *your taking the wrong jacket even after getting evidence that it is the wrong jacket* that is morally wrong. When we are considering cases like this one, in which you have evidence suggesting that you have taken the wrong jacket, it is plausible that your taking the wrong jacket is morally wrong. However, when we are considering cases in which you do not have evidence suggesting that you have taken the wrong jacket, it is not plausible that your taking the wrong jacket is morally wrong.

In response, one might contend that there is a morally relevant difference between the example in which I make a wrong turn on the way to a party and the case in which you take the wrong jacket while leaving a party. Making a wrong turn, one might argue, is not a moral mistake because in making a wrong turn, one does not necessarily violate any else's rights. Taking the wrong jacket, however, is a moral mistake because in taking the wrong jacket, one does necessarily violate someone else's rights. The person whose jacket one takes has a right to that jacket, a right that one violated when one left the party with it.⁷¹

The strength of this objection depends on the plausibility of an objective conception of rights on which an agent, *S*, violates another's right to something, *x*, whenever *S* deprives the other of *x*, even if *S* neither believes nor has sufficient reason to believe that he or she is depriving the other of *x*. But it is not at all clear that this conception of rights is more plausible than a subjective or prospective conception of rights on which an agent, *S*, violates another's right to something, *x*, only if *S* believes that (or *S*'s evidence suggests that) he or she is depriving the other of *x*.

⁷¹ Thank you to an anonymous reviewer for drawing my attention to this objection.

To see this, imagine a case in which, while leaving a party, you take someone else's jacket, mistaking it for yours. On your way to your car, you slip your hands into its pockets and find some unfamiliar keys. You are surprised by your discovery, so you return to the party only to find that you have taken the wrong jacket. You replace the jacket you took by mistake, grab yours, and leave the party. Have you done anything morally wrong in this case? While I do not think you have, others may disagree. To settle this disagreement, we might ask whether you owe anything to the person whose jacket you took by mistake. For if you have violated his or her rights, then you presumably owe him or her some compensation.

While it is implausible that you owe the jacket owner anything for mistakenly taking his or her jacket, the objectivist can argue that the reason for this is not that you have not done anything wrong, but rather that you have an excuse for your wrongdoing and therefore do not owe the jacket owner anything.

In the end, then, whether Graham's example gives us a good reason to think that there are cases of blameless wrongdoing seems to rest on whether we should grant the existence of moral excuses. I doubt that we should. While we do need legal excuses, the reason we need them is that we cannot expect lawmakers to write the law in such a way that it is never too demanding. If the law is to perform its action-guiding role, then we need laws against things like murder, rape, and theft. But when individuals commit these acts by reason of insanity, it would be inappropriate to punish them. That is where excuses come in. In the case of the "moral law," however, things are different. We can expect the moral law to be "written" in such a way that it is never too demanding, both because it is not written by humans and because its purpose is not to guide action. The reason we need legal excuses is that the law is not perfect. But the moral law is perfect, so we do not need moral excuses.

In discussing this example, I have tried to cast doubt on Graham's suggestion that one's taking the wrong jacket from a party could be a case of blameless wrongdoing. While more would need to be said to establish that conclusion, I do take myself to have shown that there is no more reason to accept Graham's suggestion than there is to reject it. That is enough for my purposes.

Let us move on, then, to Graham's second case, a supposed case of blameworthy rightdoing. Graham writes, "X might be obliged to chop off Y's leg—suppose X and Y both know that it's gangrenous and Y begs X to chop it off—but, if X does so not to help Y but, rather, to cause Y excruciating pain, X will be blameworthy even though she does what she knows she is obliged to do" (2010, p. 94). This, Graham suggests, entails that blameworthiness is not sufficient for moral wrongdoing.

There are at least two ways of reading the claim that blameworthiness is sufficient for moral wrongdoing, a broad reading and a narrow reading. On the broad reading, the claim that blameworthiness is sufficient for moral wrongdoing is the claim that if an agent is blameworthy at all, then that agent must have performed a morally wrong action, ϕ . On the narrow reading, the claim that blameworthiness is sufficient for moral wrongdoing is the claim that if an agent is blameworthy for performing some action, ϕ , then that agent's ϕ -ing must be morally wrong.

Graham's example is a counterexample to the broad reading, for X is blameworthy for something, but X has not performed a wrong action. Graham's example is not, however, a counterexample to the narrow reading, for X is not blameworthy for what she did, only for why she did it. To see this, ask yourself whether you would have described X as blameworthy if you had not known that she wanted to cause Y excruciating pain. Presumably, you would not have. This suggests that X is blameworthy for *wanting to cause Y excruciating pain* (i.e., for her *motive*), not for *chopping off Y's leg* (i.e., for her *action*).

If that is right, then while Graham is right to suggest that X's chopping off Y's leg is *morally permissible*, if not obligatory, he is wrong to suggest that X's chopping off Y's leg is *blameworthy*. For it is not X's *chopping off Y's leg* that is blameworthy; it is X's *wanting to cause Y excruciating pain* that is blameworthy. X is blameworthy for something, *her motives*, but she has not performed a wrong action.

In order to block the inference from the fact that Luke would be blameworthy for injecting Nathan with syringe C to the conclusion that Luke's injecting Nathan with syringe C would be morally wrong, however, Graham must provide a counterexample to the narrow reading, for the inference in question depends on the truth of the claim that if an agent would be blameworthy for performing some action, ϕ , then that agent's ϕ -ing must be morally wrong. In order to give a counterexample to this claim, Graham would have to give an example in which an agent is blameworthy for performing some action (i.e., for *what* he or she has done, not for *his or her motives*) but in which that action is not morally wrong. It is hard to imagine how one would go about constructing such a case, for how could one be blameworthy for performing an action unless that action is morally wrong? Because Graham's example does not provide us with an answer to this question, it cannot do the job he wants it to. That is, it cannot cast doubt on the narrow reading of the claim that blameworthiness is sufficient for moral wrongdoing.

In the end, if we draw a distinction between *wrongness* and *moral wrongness* and a distinction between *blameworthy motives* and *blameworthy actions*, we should not be convinced, as Graham seems to be, that blameworthiness and moral wrongdoing come apart. Beyond this, however, it is worth noting that the simple argument against objectivism does not necessarily depend on the connection between blameworthiness and moral wrongdoing. It just depends on its being true that Luke is morally obligated to inject Nathan with syringe P or that Luke is morally prohibited from injecting Nathan with syringe C. One way we might get to those conclusions is through facts about

blameworthiness, but there are other ways we might get to those conclusions. For example, we might get to those conclusions through facts about what all morally decent people would do or through facts about when we would describe someone as having violated a moral obligation. In the end, therefore, Graham's objection to the simple argument against objectivism is not decisive.

3.3 Graham's Argument from New Information

Of course, even if Graham's objection to the simple argument against objectivism fails, that is irrelevant if there are sound arguments *for* objectivism. So, in what remains, I discuss Graham's two arguments for objectivism. While each is supposed to give us a reason to prefer objectivism to its competitors, I argue that neither succeeds.

Graham's first argument comes from remarks made by W. D. Ross in *The Right and the Good*.⁷² In presenting the argument, Graham writes, "When we change our view about our moral obligations as a result of acquiring information, we do not 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."⁷³ This, Graham thinks, entails the truth of objectivism.

⁷² Ross 2002/1930: 32 writes, "Many people would be inclined to say that the right act for me is ... that which on all the evidence available to me I should think to be my duty. But suppose that from the state of partial knowledge in which I think act A to be my duty, I could pass to a state of perfect knowledge in which I saw act B to be my duty, should I not say 'act B was the right act for me to do?'" It is worth noting that Ross later abandoned objectivism in favor of subjectivism (1939).

⁷³ Graham 2010: 91.

3.4 A Response

In response, Elinor Mason writes, “This is not a strong argument.” She continues, “The way that we speak often implies this, as Graham points out. But it is not hard to think of cases where the way that we speak favors the prospective or subjective account of rightness.”⁷⁴ I think that Mason's criticism of Graham's argument is on point, so I will flesh it out quickly and move on.

Graham is obviously right that *sometimes* when we change our view about our moral obligations as a result of acquiring new information, we take ourselves to have identified what our moral obligations were all along. Usually this happens when we think that we should have been aware of that information all along. Thus, when racists, sexists, or speciesists change their view about their moral obligations as a result of acquiring new information, they often take themselves to have identified what their moral obligations were all along.

But other times, when we change our view about our moral obligations as a result of acquiring new information, we think that what we are morally obligated to do has changed. Reconsider *Syringes*. In that case, Luke believed, in accordance with his evidence, that it was better for him to inject Nathan with syringe P than with syringe C. Thus, he presumably believed that he was morally obligated to inject Nathan with syringe P. Now imagine that someone informs Luke that the syringe labels have been switched. In this case, Luke presumably comes to believe that it is better for him to inject Nathan with syringe C than with syringe P. Moreover, he presumably comes to believe that he is morally obligated to inject Nathan with syringe C. While it is possible that, in this case, Luke thinks that he has come to see what he was morally obligated to do all along, it is at least as likely, if not more likely, that Luke thinks that his moral obligation has changed. That is, it is at least as likely that Luke thinks that whereas he *was* morally obligated to inject Nathan

⁷⁴Mason 2013: 10.

with syringe P (*before* he was informed that the syringe labels had been switched), he *is now* morally obligated to inject Nathan with syringe C. Thus, Graham's (and Ross's) *Argument from New Information* does not give us any reason to think that objectivism is superior to subjectivism or prospectivism.

3.5 Graham's Argument from *Promise*

Although we have seen that Graham's *Argument from New Information* is unsound, Graham provides a second argument for the same conclusion. At the heart of that argument is the following example:

Promise

Phoebe has promised Kenneth that she will turn on a certain light at 6 p.m. All Phoebe's evidence indicates that it is not 6 p.m. Phoebe asks you what she's morally obliged to do. You know of Phoebe's promise and that it is 6 p.m., but you also know that Phoebe neither believes, nor has any evidence, that it is 6 p.m. You tell Phoebe that she's morally obliged to turn on the light.⁷⁵

Graham then uses our intuitions about *Promise* to construct the following argument for objectivism:

Graham's Argument from *Promise*

- (6) In *Promise*, you speak truly in saying Phoebe is morally obliged to turn on the light.
- (7) According to no plausible subjectivist [or prospectivist] moral theory does Phoebe have a moral obligation to turn on the light in *Promise*.⁷⁶
- (8) Therefore, objectivism is true.⁷⁷

⁷⁵ Graham 2010: 92. I have changed the names of the characters to facilitate discussion of the example.

⁷⁶ Graham leaves out prospectivism because he considers it a version of subjectivism.

⁷⁷ Graham 2010: 92. I have changed the argument in order to bring it in line with *Promise*.

We are supposed to have the intuition that (6) is true. Assuming that we do have that intuition, it is possible we have it, not because we think you speak *truly* in saying that Phoebe is morally obliged to turn on the light, but because we think it is *appropriate* for you to tell Phoebe that she is morally obliged to turn on the light. Maybe this is because, as Michael J. Zimmerman suggests, your telling Phoebe that she is morally obliged to turn on the light is a morally obligatory lie.⁷⁸ Or maybe it is because, as Mason suggests, “Someone who tells you that you ought to do Φ (when your evidence, or the evidence as you see it, does not indicate that) is saying something that is elliptical for, “The evidence from my point of view points to doing Φ .”⁷⁹ Either way, we might worry that the intuitive plausibility of (6) does not justify our believing (6).

Furthermore, it is possible we have the intuition that (6) is true, not because we think that Phoebe is morally obliged to turn on the light *before* you responded to her question, but because we think that Phoebe is morally obliged to turn on the light *after* you responded to her question. This, however, is perfectly consistent with both subjectivism and prospectivism (in which case (7) is false). For when you tell Phoebe that she is morally obliged to turn on the light, you give her evidence that it is 6 p.m., and that, in turn, gives her evidence that she is morally obliged to turn on the light. Thus, in telling Phoebe that she is morally obliged to turn on the light, you make it the case that she is morally obliged to turn on the light according to both subjectivism and prospectivism, for in telling Phoebe that she is morally obliged to turn on the light, you make it the case that she believes, in accordance with her evidence, that she is morally obliged to turn on the light.

⁷⁸ This suggestion appears in Zimmerman 2006: 344-345 and 2008: 32.

⁷⁹ Mason 2013: 11.

3.6 Graham's Argument from *Promise**

We can sidestep these worries by altering Graham's case. I will call the new version of the case *Promise**. In *Promise**, everything is the same as in *Promise*, but this time, you do not tell Phoebe that she is morally obliged to turn on the light. You simply judge that she is morally obliged to turn on the light.

In light of this change, we should alter Graham's argument accordingly.

Graham's Argument from *Promise**

- (6*) In *Promise**, your judgment that Phoebe is morally obliged to turn on the light is true.
- (7*) According to no plausible subjectivist or prospectivist moral theory does Phoebe have a moral obligation to turn on the light in *Promise**.
- (8) Therefore, objectivism is true.

Although the second premise of this new argument, (7*), is clearly true, I worry that its first premise, (6*), is not. My goal, however, is not to show that (6*) is false; it is just to show that we have no more reason to accept it than we have to reject it. If I can do that, I will have shown that Graham's second argument does not give us any reason to think that objectivism is superior to subjectivism or prospectivism.

Before discussing (6*), however, I want to alleviate the worries of those who think that, because Graham does not endorse (6*), I am arguing against a straw man. In response, note that Graham does not think that, in *Promise*, it is your telling Phoebe that she is morally obliged to turn on the light that makes her morally obliged to turn on the light. For neither subjectivists nor prospectivists would dispute that for the reasons outlined above. Instead, Graham thinks that, in *Promise*, Phoebe is morally obligated to turn on the light all along. Consequently, he needs (6*) to be true.

3.7 The Moral Obligation to Keep One's Promises

Graham does not give us an argument for (6) or (6*), so his argument for objectivism depends on the intuitive plausibility of that premise. Of course, some will not find (6*) plausible. If you are one of those people, then I trust I have already accomplished my goal. But let us say that you do find (6*) plausible. In that case, I need to explain why we have no more reason to accept (6*) than we have to reject it. To do that, I will consider three principles that would both explain (6*)'s plausibility and entail its truth. Then, I will show that none of those principles is more plausible than its denial.

First, one might think (6*) is true because one thinks that agents are always morally obligated to do what they *strictly speaking promise to do* (hereafter, simply *promise to do*). I will call this view *simple objectivism about promising*.

There are a number of problems with this view. First, it violates the 'ought' (or, more accurately, 'morally obliged') *implies 'can' principle*. Imagine, for example, that Ashley promises to meet Brad for lunch at noon but that, as it turns out, she cannot, for she will inevitably be captured by a band of evildoers intent on preventing her from keeping her promise. In a case like this, simple objectivism about promising violates the 'ought' implies 'can' principle because it entails that Ashley is morally obligated to meet Brad for lunch at noon even though she cannot.

Second, simple objectivism about promising implausibly implies that an agent's moral obligation to keep a particular promise cannot be overridden. Imagine, for example, a similar case in which Ashley, rather than being captured, comes across Chris, who is in desperate need of help. Unless Ashley helps Chris, he will die. In this case, we think that Ashley's moral obligation to keep her promise to Brad is overridden by her moral obligation to help Chris. But simple objectivism about promising does not allow for this.

In light of cases like these, objectivists should move to what I will call *sophisticated objectivism about promising*. According to this view, an agent is morally obliged to keep a promise if and only if she is neither unable to keep it nor under a weightier moral obligation to do something incompatible with her keeping it. This view avoids the problems that plague simple objectivism about promising, and it entails (6*). So, unless there are good reasons to be skeptical of sophisticated objectivism about promising, we should accept (6*).

But there *are* good reasons to be skeptical of sophisticated objectivism about promising. Consider a similar case in which Ashley promises to meet Brad at a local diner for lunch at noon. After making her promise but before meeting Brad, Ashley learns that there is an evildoer who is intent on preventing her from meeting Brad. Moreover, she believes, in accordance with her evidence, that the evildoer will be waiting for her by the front entrance of the diner. In light of this, she decides to use the diner's side entrance. Unfortunately, however, she is wrong about where the evildoer would be waiting. She is captured and, therefore, fails to keep her promise.

If sophisticated objectivism about promising is true, Ashley's failure to keep her promise is morally wrong, for she was able to keep her promise (on account of the fact that she would have kept it had she used the front door) and was not under a weightier moral obligation to do something that was incompatible with her keeping it. But it is implausible that Ashley's failure to meet Brad for lunch at noon is morally wrong. This suggests that sophisticated objectivism about promising, like simple objectivism about promising, is false.

There is a more promising view, however, which I will call *non-objectivism about promising*. According to this view, agents are morally obligated, at most, to do what they believe, or what their evidence suggests, makes it most likely that they will keep their promises (provided that they can do those things and that those things are compatible with any weightier moral obligations they have). This view can explain why it is implausible that Ashley's failure to meet Brad for lunch at

noon was morally wrong, namely, by pointing to the fact that Ashley's doing what she believed, in accordance with her evidence, made it most likely that she would keep her promise to Brad. Moreover, this view avoids the problems that plague simple objectivism about promising. Unlike objectivist views about promising, however, this view does not entail (6*).

In defense of (6*), an objectivist might argue as follows: "The non-objectivist's judgment that Ashley's failure to meet Brad for lunch at noon is not morally wrong is based on the intuition that Ashley is not blameworthy. But blameworthiness and moral wrongdoing come apart in cases like this. Ashley has done something morally wrong, but she should be excused from blame in virtue of the fact that her actions were reasonable. Thus, the example does not give us any reason to move from sophisticated objectivism to non-objectivism about promising."

The objectivist's argument is not obviously problematic. In fact, if we had good reasons to think that blameworthiness and moral wrongdoing come apart in cases like this, we would have a good reason to prefer sophisticated objectivism about promising to non-objectivism about promising convincing. But, as I pointed out earlier, we do not have good reasons to think that blameworthiness and moral wrongdoing come apart. Thus, we do not have a good reason to prefer sophisticated objectivism about promising to non-objectivism about promising.

It follows that (6*) may very well be false. For all we know, Phoebe has done everything she believes, or that her evidence suggests, makes it most likely that she will keep her promise in *Promise**. So, given that we do not have a good reason to prefer sophisticated objectivism about promising to non-objectivism about promising, we do not have a good reason to think that (6*) is any more likely to be true than it is to be false.

3.8 The Moral Obligation to Track the Time

But let us assume that Phoebe has not done everything she believes, or that her evidence suggests, makes it most likely that she will keep her promise in *Promise**. What should we say about cases like this? To answer this question, let us consider the following example.

Negligence

Having promised Kenneth that she would turn on a certain light at 6 p.m., Phoebe knows that she should use a watch to track the time. But Phoebe does not want to do that. Instead, she wants to use the sun's position to track the time. Judging by the sun's position, she estimates that it is 5 p.m. when, in fact, it is 6 p.m.

In *Negligence*, Phoebe *does not* do everything she believes, or that her evidence suggests, makes it most likely that she will keep her promise, for she decides to use the sun's position (rather than a watch) to track the time. When I think about cases like this one, as compared to cases in which Phoebe *does* do everything she believes, and that her evidence suggests, makes it most likely that she will keep her promise, I find (6*) more plausible. It just does not seem like Phoebe's epistemic irresponsibility gets her out of her moral obligation to keep her promise. If this is true, then objectivism is vindicated.

It is obvious that in *Negligence*, Phoebe has done *something* morally wrong, and a principle on which epistemic irresponsibility does not get one out of one's moral obligation to keep one's promises gets us that. In particular, this principle entails that what is morally wrong with Phoebe's behavior is *her failure to turn on the light at 6 p.m.* The problem, however, is that it is at least as plausible, if not more plausible, that what is wrong with Phoebe's behavior is *her failure to use a watch to track the time.* This is a problem for Graham because both subjectivism and prospectivism can get us that conclusion. In *Negligence*, Phoebe believes, in accordance with her evidence, that she should use a watch to track the time, and, therefore, both subjectivism and prospectivism entail that she acts wrongly when she fails to do so.

Of course, we could always supplement the principle we started with, on which epistemic irresponsibility does not get one out of one's moral obligation to keep one's promises, with another principle, on which agents who promise to do something at a certain time are thereby obligated to track the time as best they can. This would allow the objectivist to say that Phoebe's failure to use a watch to track the time was morally wrong. But this combination of principles entails that Phoebe violates two moral obligations, a moral obligation to track the time as best she can *and* a moral obligation to keep her promise to Kenneth. This seems to over-penalize Phoebe. It is true that she has violated a moral obligation, but she has not violated *two* moral obligations.

In the end, we have no more reason to accept a principle on which epistemic irresponsibility does not get one out of one's moral obligation to keep one's promises than we have to reject it, for although that principle gets us the conclusion that Phoebe has done something morally wrong, it is not clear that it correctly identifies which of her actions was morally wrong. Thus, we still do not have a good explanation for the truth of (6*).

3.9 The Moral Obligation to Play It Safe

To summarize, we cannot explain the truth of (6*) by appealing to a principle on which agents are always morally obligated to keep their promises or by appealing to a principle on which epistemic irresponsibility does not get one out of one's moral obligation to keep one's promises. But maybe we can explain the truth of (6*) by appealing to a principle on which agents who have promised to do something at a certain time, and who are not sure what time it is, are morally obligated to play it safe by doing the thing they promised to do earlier rather than later. If so, then we will have a good reason to be objectivists.

In *Promise**, we know that Phoebe has promised Kenneth that she will turn on a certain light at 6 p.m. We also know that when 6 p.m. rolls around, all of Phoebe's evidence indicates that it is not 6 p.m. But we do not know how confident Phoebe is that it is not 6 p.m. She does, after all, ask you what she is morally obliged to do, which suggests that she cannot rule out that it is 6 p.m. And if she cannot rule out that it is 6 p.m., it is possible that she is morally obliged to go ahead and turn on the light, just to be safe.

To see why this might be, we need to think about why Kenneth wants Phoebe to turn on a certain light at 6 p.m. Maybe Kenneth wants Phoebe to turn on one of his house lights to ward off thieves, or a lighthouse light to help sailors return to shore, or his Christmas lights for others to enjoy. There are, of course, many other ways to fill in the details of *Promise**, but these strike me as some of the more natural ways to do so.

Now, notice that all of these cases are such that, while it may be important that Phoebe turns on the relevant light *around* 6 p.m., it is not particularly important that she turn on the relevant light *at exactly* 6 p.m. Moreover, in each of these cases, it seems that it would better for Phoebe to err on the side of turning the lights on earlier rather than later. So, for example, it seems better for Phoebe to turn on Kenneth's house light at 5 p.m. rather than 7 p.m. if she is wanting to help him ward off thieves. The same holds true if Phoebe is supposed to be turning on a lighthouse light or Kenneth's Christmas lights. As a result, it is plausible that, in *Promise**, if Phoebe is at all unsure about what time it is, she is morally obliged to just go ahead and turn on the light.

Depending on how we fill in the details of the case, subjectivism and prospectivism can get us this result. If, for example, Phoebe believes, in accordance with her evidence, that she does not have much evidence to work with and that it would be best for her to play it safe by turning on the light sooner rather than later, then subjectivism and prospectivism both entail that (6*) is true, in which case (7*) is false.

There are, however, other ways to fill in the case such that subjectivism and prospectivism cannot get us this result. If, for example, Phoebe believes, in accordance with her evidence, that she has plenty of evidence to work with and that it would be bad for her to play it safe by turning on the light sooner rather than later, then neither subjectivism nor prospectivism entails that (6*) is true. Here is one such case.

Lighthouse

A lighthouse keeper named Kenneth is going on an overnight trip, so he asks his friend Phoebe to fill in for him. In particular, he asks Phoebe to turn on the lighthouse light around the time the sun goes down, at 6 p.m. He also stresses the importance of not turning it on too early on account of the fact that it can only run for 12 hours at a time, at which point it must turn off for three hours to recharge. Because the sun will not rise until around 6 a.m. the following morning, Phoebe's turning on the lighthouse light too early could have very bad consequences. Confident that she can do these things, Phoebe promises to fill in for Kenneth and heads for the lighthouse. She arrives early in the afternoon and finds that she has time for a nap, so she lays down to sleep. While she is sleeping, however, someone sneaks in and turns all the clocks in the lighthouse back an hour. Consequently, when Phoebe wakes up, she thinks it is 4 p.m. even though it is actually 5 p.m. Even so, Phoebe knows how important it is that she turn the lighthouse light on around the time the sun goes down, so she tries to use the sun's position to determine when to turn on the light. Unfortunately, however, storm clouds fill the sky, making that impossible. Consequently, Phoebe is forced to consult the clocks in the lighthouse, and she does so regularly, but this leads her to turn on the lighthouse light at 7 p.m., an hour too late.

When I think about a case like *Lighthouse*, I find that the intuition that Phoebe was morally obliged to play it safe by turning on the light sooner rather than later goes away because there is simply no way for Phoebe to play it safe. She just has to act in accordance with her evidence.

What this suggests is that we cannot get (6*) by appealing to a principle on which agents who have promised to do something at a certain time, and who are not sure what time it is, are morally obligated to play it safe by doing the thing they promised to do earlier rather than later. For in cases like *Lighthouse*, doing the thing one promised to do earlier rather than later could be just as bad as (if not worse than) doing it later rather than earlier.

In Sections 3.7-3.9, I have shown that Graham's Argument from *Promise** fails to give us a reason to be objectivists. Although it is initially plausible that, in *Promise**, your judgment that Phoebe is morally obliged to turn on the light is true, upon reflection, we see that there is no principle that both entails the truth of (6*) and is more plausible than its denial.

3.10 Conclusion

I began this chapter by showing that Graham's objection to the simple argument against objectivism fails. Then, I turned my attention to showing that Graham's arguments for objectivism do not give us any reasons to prefer it to subjectivism or prospectivism. In particular, I showed that while we sometimes, maybe even usually, think that new information helps us identify what our moral obligations were all along, there are other times when we think that new information changes our moral obligations. Then, I showed that while we might be tempted to think that, in *Promise*, you speak truly in telling Phoebe that she is morally obliged to turn on the light, there is no principle that both entails that Phoebe is morally obliged to turn on the light at 6 p.m. and is more plausible than its denial.

There are, of course, other arguments one could give for objectivism, and therefore other considerations that may speak in its favor. So, I do not take myself to have disproven the view on the basis of this chapter alone. I do, however, take myself to have shown that Graham's defense of objectivism fails. If there is no better defense of objectivism available, I contend that we have more reason to reject objectivism than to accept it.

CHAPTER 4

IN DEFENSE OF SUBJECTIVISM ABOUT MORAL OBLIGATION⁸⁰

4.0 Introduction

Although the leading theories in normative ethics (i.e., consequentialism, Kantianism, and virtue ethics) have traditionally been formulated as versions of objectivism, I have argued that we should reject objectivism in favor of *subjectivism*, the view that an agent's moral obligations depend on her beliefs, or *prospectivism*, the view that an agent's moral obligations depend on her evidence. While I am not the first to argue against objectivism, others who have argued against it (e.g., Elinor Mason and Michael J. Zimmerman) have tended to be prospectivists.⁸¹ In this chapter, I argue that the reasons for moving from objectivism to prospectivism also speak in favor of moving from prospectivism to subjectivism. Moreover, none of the most common objections to subjectivism are successful. Thus, subjectivism deserves more attention than it has received.

4.1 Motivating Subjectivism

Objectivism is the view that an agent is morally obligated to perform an act if and only if it is the agent's deontically best option, where an option is an agent's *deontically* best option (as

⁸⁰ I am grateful to Chris Heathwood, Elinor Mason, and two anonymous reviewers for their helpful comments on earlier drafts of this chapter. I am also grateful to Renee Bolinger and Michael J. Zimmerman for their willingness to spend time discussing this chapter's main arguments with me.

⁸¹ See Zimmerman 2006, 2008, 2014, and Mason 2013.

compared to an agent's *morally* best option) if and only if none of the agent's alternatives are as good "in terms of what matters morally, insofar as [the agent's] moral obligations are concerned."⁸² This view might sound like it privileges maximizing views of moral obligation, but it does not. For even if one has two options that are *deontically* equivalent (i.e., equally good in terms of what matters insofar as one's moral obligations are concerned), it is possible that one of those options is, nevertheless, supererogatory, for it is possible that one of those options is *morally* superior to the other in a way that is not relevant to the agent's moral obligations.⁸³

Objectivism is plausible in cases where agents *know*, or at least *can know*, what is best. It is implausible, however, in cases where agents *cannot know* what is best. Consider the following example from Frank Jackson:

The First Drug Example

Jill is a physician who has to decide on the correct treatment for her patient, John, who has a minor but not trivial skin complaint. She has three drugs to choose from: drug A, drug B, and drug C. Careful consideration of the literature has led her to the following opinions. Drug A is very likely to relieve the condition but will not completely cure it. One of drugs B and C will completely cure the skin condition; the other though will kill the patient, and there is no way that she can tell which of the two is the perfect cure and which the killer drug.⁸⁴

Assuming, with objectivists, that there is no moral value in acting in accordance with one's beliefs or one's evidence, and assuming that what is best for John is his skin condition's being cured, objectivism entails that Jill is morally obligated to prescribe whichever drug will cure John's skin condition. Let's say that that's drug B. In that case, objectivism entails that Jill is morally obligated to prescribe drug B.

⁸² Zimmerman 2014: 4.

⁸³ See Zimmerman 2014: 4 for a more thorough discussion of this issue.

⁸⁴ Jackson 1991: 462-463.

Many find this implausible.⁸⁵ Jill is morally obligated to prescribe drug A, not drug B. But why? In response to this question, Mason writes, “Intuitively, it is because she should not take the substantial risk of a really terrible outcome, the patient’s death. In other words, risk has found its way into rightness.”⁸⁶ I think that this is roughly the correct diagnosis. Unreasonably risky behavior is morally wrong. Consequently, Mason and Zimmerman reject objectivism, which does not have this implication, in favor of prospectivism, the view that agents are morally obligated to do what is prospectively best,⁸⁷ where what is prospectively best is, roughly, what an agent’s evidence suggests is her best bet given the values at stake.⁸⁸

The problem with this move is that prospectivism, like objectivism, also seems to imply that there are cases in which unreasonably risky behavior is not morally wrong. Consider the following example, which is based on an example from Holly Smith:

The Second Drug Example

Harry is a physician who has to decide on the correct treatment for his patient, Renée. Careful consideration of the literature has led Harry to believe that his doing nothing (act E) is a significantly better bet than either of his alternatives, namely, prescribing drug F or drug G. Harry’s senior colleague, however, knows that Harry has made a mistake. While Harry’s evidence *does* suggest that act E is a significantly better bet than *one* of his alternatives (i.e., his prescribing drug F or drug G), Harry’s evidence does not suggest that act E is a significantly better bet than his other alternative. In fact, Harry’s evidence suggests that his other alternative is a *slightly* better bet than act E. Harry’s senior colleague tells Harry this, and Harry justifiably believes her, but before he can ask her which of his alternatives is the slightly better bet, she is called away to deal with an emergency.⁸⁹

⁸⁵ See, e.g., Jackson 1991, Zimmerman 2008, 2014, and Mason 2013.

⁸⁶ Mason 2013: 12.

⁸⁷ See Mason 2013: 2 and Zimmerman 2014: 32, 34-37. Jackson 1991, I should note, also rejects objectivism, but he moves to a version of subjectivism rather than a version of prospectivism.

⁸⁸ Alternatively, we might say that what is prospectively best is that which maximizes *evidence-relative* expected *deontic* value. As this suggests, what is prospectively best will depend on things like what one’s evidence suggests one’s alternatives are, what one’s evidence suggests the deontic values of those alternatives are, what one’s evidence suggests the outcomes of those alternatives would be, how likely one’s evidence suggests those outcomes would be, and how deontically valuable one’s evidence suggests those outcomes are.

⁸⁹ Compare Smith 2011: 5. Smith uses her example to suggest either that there is something wrong with this kind of argument or that prospectivism is false. She does not use her example to argue for subjectivism.

Just as it would be unreasonably risky for Jill to prescribe drug B or C (because one of those drugs will kill her patient, and she does not know which one it is), it would be unreasonably risky for Harry to prescribe drug F or G (because Harry's evidence suggests that one of those drugs is a significantly worse bet than Harry's doing nothing, and he does not know which one it is). Nevertheless, prospectivism seems to require Harry to prescribe drug F or drug G, for Harry's evidence suggests that his prescribing one of those drugs is his best bet. Thus, prospectivism, like objectivism, entails that there are cases in which unreasonably risky behavior is not morally wrong.

In response to this argument, Zimmerman admits that it would be unreasonably risky for Harry to prescribe drug F or drug G but denies that prospectivism requires Harry to do so.⁹⁰ In defending prospectivism, Zimmerman distinguishes between two versions of the view. According to the first version of prospectivism, an individual's moral obligations depend on the evidence that is available to him. This is the evidence "of which the person can, in some sense, avail himself" or, in other words, "what a person would be justified in believing, if certain conditions obtained."⁹¹ Zimmerman admits that he is not sure exactly what evidence should count as being available to an agent, but this is supposedly all right because he believes we should reject this version of prospectivism in favor of a second on which an individual's moral obligations depend on the evidence he has. This is the evidence "of which the person avails himself" or, in other words, "what a person *does* justifiably believe."⁹²

Now, there is a problem with Zimmerman's suggestion that these two formulations of the second version of prospectivism are equivalent. For if a person makes a mistake in reasoning from the evidence of which he has availed himself to a belief about how to act, the first formulation will

⁹⁰ Zimmerman 2014: 72.

⁹¹ Zimmerman 2014: 73.

⁹² Zimmerman 2014: 73, emphasis in original.

entail that he is morally obligated to do whatever the evidence of which he has availed himself implies he should do. The second formulation, however, will not entail that he is morally obligated to do anything (because he does not have any justified beliefs). Given that Zimmerman would, presumably, want his view to entail that this person is morally obligated to do whatever the evidence of which he availed himself implies he should do, I assume that it is the first formulation to which Zimmerman is committed. Regardless of which interpretation one goes with, however, Zimmerman is right that this second version of prospectivism does not require Harry to prescribe drug F or drug G. Instead, it entails that Harry is morally obligated to perform act E (because he is justified in believing that act E is his best bet).

This move, from the view that agents are obligated to do what their available evidence suggests is their best bet to the view that agents are obligated to do what they are justified in believing is their best bet, is an improvement, but it does not go far enough. Imagine a third drug example in which Harry forms the belief that his doing nothing (act E) is a significantly better bet than his alternatives (viz., prescribing drug F or drug G) on the basis of the evidence of which he has availed himself. At this point, however, Harry's colleague tells him that, relative to the evidence of which he has availed himself, his belief is unjustified, and before Harry can ask her what he would be justified in believing on the basis of that evidence, she is called away to deal with an emergency.

In a case like this, in which Harry knows that his belief that act E is his best bet is unjustified, it would be unreasonably risky for Harry to do what he would, in fact, be justified in believing is his best bet. So, given that it is morally wrong to take unreasonable risks, it follows that it would be morally wrong for Harry to do what he would, in fact, be justified in believing is his best bet. Thus, Harry is obligated to do what he *believes* is his best bet, even though he knows that this belief is unjustified.

In the end, if it is morally wrong to do *what is deontically best* when one does not know what that is (on account of its being too risky), then it is also sometimes morally wrong to attempt to do *what is prospectively best* when one does not know what that is (on account of its being too risky). So, the reason we moved from objectivism to prospectivism is also a reason to move from prospectivism to subjectivism, the view that agents are morally obligated to do what they believe is deontically best (hereafter, simply *what they believe is best*).

This, of course, is consistent with subjectivism's being false, and that is the position most people take. While some in the history of philosophy have endorsed subjectivism,⁹³ few contemporary philosophers do.⁹⁴

In *Living with Uncertainty*, Zimmerman presents four common objections to subjectivism.⁹⁵ These presumably explain its unpopularity. Below, however, I show that none of them gives us a good reason to reject subjectivism.

4.2 Does "Ought" Imply "Can"?

One of Zimmerman's worries is that subjectivism entails that "ought" does not imply "can." He writes, "The Subjective View implies that some activity (whether a full-fledged act or merely a self-exertion) is obligatory if I believe it to be the best that I can do. In holding this belief I presuppose that the activity is in my control, but the presupposition may be false."⁹⁶ Thus, subjectivism seems to entail that agents could be morally obligated to do things they cannot do.

⁹³ See Prichard 2002/1932 and Ross 1939.

⁹⁴ As far as I know, Jackson 1991 is the only contemporary philosopher who endorses a version of subjectivism.

⁹⁵ Zimmerman 2008: 13-14.

⁹⁶ Zimmerman 2008: 14.

For example, imagine that a teacher believes that the best thing for her to do is to persuade her students to donate money to Oxfam. If subjectivism is true, then the teacher is morally obligated to persuade her students to donate money to Oxfam. But, of course, it may be impossible for her to persuade her students to donate money to Oxfam. So, if subjectivism is true, the teacher is morally obligated to do something she cannot do, and, therefore, if “ought” implies “can,” subjectivism is false.

In response, subjectivists can abandon a simple version of subjectivism, on which an agent is morally obligated to perform an action if and only if (i) he believes that performing it is best, in favor of a more sophisticated version of subjectivism, on which an agent is morally obligated to perform an action if and only if (i) he believes that performing it is best *and* (ii) he is able to perform it.⁹⁷

It is worth noting that a similar objection can be raised against simple versions of prospectivism,⁹⁸ for just as *one might believe* that an action is one’s best bet even though it is ultimately impossible for one to perform that act, so too *one’s evidence might suggest* that an action is one’s best bet even though it is ultimately impossible for one to perform *that* act. Thus, this objection gives us no reason to prefer prospectivism to subjectivism.

⁹⁷ This sophisticated version of subjectivism remains a version of subjectivism insofar as it remains a view on which an agent’s moral obligations depend on her beliefs.

⁹⁸ Graham 2010: 90 does exactly this by pointing out that an agent’s evidence can suggest that some action is best even if she cannot perform it.

4.3 Did Hitler Act Wrongly?

Zimmerman also worries that subjectivism implies that Hitler did nothing wrong, provided that Hitler did what he believed was best. But, Zimmerman writes, “[I]t is grotesque to think that such a perverse belief could render mass murder morally permissible.”⁹⁹ Let us assume that Hitler did, in fact, believe that his commanding mass murder was best. In that case, subjectivism *does* imply that Hitler’s commanding mass murder was morally permissible. Although this might seem grotesque, it’s not.

Consider a contemporary case, the case of the anti-vaxxer. The anti-vaxxer believes that vaccines are dangerous and therefore that parents should not vaccinate their children. What is more, he actively discourages parents from vaccinating their children. The anti-vaxxer, let us stipulate, is deeply mistaken about the expected utility of vaccinations. Contrary to what he believes, his evidence strongly suggests that vaccines are harmless. Moreover, his evidence suggests that the value of herd immunity is much greater than he would admit. Given these facts, he is not doing what is best or what his evidence suggests is best. In spite of this, I am not inclined to describe the anti-vaxxer’s behavior as immoral. It may be dangerous, but given his belief that vaccines are dangerous, it would be immoral for him *not* to discourage other parents from vaccinating their children. If the anti-vaxxer has done anything wrong, he has done something wrong in forming his beliefs, but given those beliefs, we cannot criticize him for acting in accordance with them.

This analysis of the anti-vaxxer’s behavior is consistent with subjectivism. Subjectivists cannot criticize the anti-vaxxer for acting in accordance with his beliefs, but they can criticize him for the way he has gone about forming his beliefs. So, while they cannot say that it is morally wrong for the anti-vaxxer to discourage others from vaccinating their children, they can say that the anti-

⁹⁹ Zimmerman 2008: 14.

vaxxer has done something wrong in forming his belief that discouraging others from vaccinating their children is best.¹⁰⁰ The same is true for Hitler. Subjectivists cannot say that it was morally wrong for Hitler to command mass murder (assuming he believed that his commanding mass murder was best), but they can say that Hitler did something wrong in forming his belief that his commanding mass murder would be best.¹⁰¹

Subjectivism's opponents might object that something is lost when we move from thinking that Hitler did something morally wrong in commanding mass murder to thinking that he did something morally wrong in forming his belief that his commanding mass murder was best. After all, *Hitler's believing that his commanding mass murder was best* did not result in mass murder, but *his commanding mass murder* did.

Even if it is true that Hitler's *commanding* mass murder was a more direct cause of mass murder than his *believing* that his commanding mass murder was best was, it is plausible that Hitler would not have commanded mass murder unless he had believed that it was best. When we recognize this, it is plausible that Hitler's believing that his commanding mass murder was best was not as harmless as it initially seems. Given that we are disposed to act in accordance with our beliefs, it is plausible to think that we have a very serious obligation to be careful in forming our beliefs.

One problem with this defense of subjectivism is that this supposed obligation to be careful in forming our beliefs sounds like an objective obligation, that is, an obligation that agents have regardless of what they believe. But that is not necessarily true. Notice that we would be hard-

¹⁰⁰ They can also say that it would be morally permissible for someone to stop the anti-vaxxer from spreading anti-vaxxer propaganda as long as that other person believes that doing so is best.

¹⁰¹ They can also say that it would be morally permissible for someone to stop Hitler from commanding mass murder as long as that other person believes that doing so is best.

pressed to find someone who is indifferent to whether others are careful in forming their beliefs. Furthermore, it would be surprising to find someone who thinks that although it is best that others be careful in forming their beliefs, he or she need not be. This suggests that there is widespread agreement that it is best for agents, ourselves included, to be careful in forming their beliefs.

I see no reason to think that Hitler was any different than the rest of us in this respect. Thus, on the assumption that Hitler was not careful in forming his beliefs about the value of his commanding mass murder, subjectivists are not forced to admit that Hitler did nothing wrong.

But what if Hitler was careful in forming his beliefs, or at least as careful as he thought it was appropriate to be? In that case, subjectivism entails that Hitler did nothing wrong, which is implausible.¹⁰²

There are two things subjectivists can say in response to this objection. First, subjectivists can point out that, even if Hitler thought that the best thing he could do was to command mass murder, it is hard to imagine him not recognizing that his commanding mass murder might very well be terribly immoral. And presumably, in cases like this, we think that one should hedge one's bets, for one should not risk doing something terribly immoral simply because it is slightly more likely to be better than the alternative. Thus, if Hitler thought that one should hedge one's bets in such situations but commanded mass murder anyway, then subjectivists are not forced to admit that Hitler did nothing wrong.

But if Hitler did not think that one should hedge one's bets in such situations, or if Hitler did not consider the possibility that his commanding mass murder could be terribly immoral, then subjectivists must admit that Hitler did nothing wrong. I do not find this implausible. For if Hitler was careful in forming his beliefs and acted in accordance with those beliefs, and we are willing to

¹⁰² I am grateful to an anonymous reviewer for highlighting this further worry.

grant that it is no way wrong for the epistemically responsible anti-vaxxer to discourage other parents from vaccinating their children, provided that he sincerely believes that vaccines are extremely dangerous, then it seems we should also be willing to grant that it is in no way wrong for an epistemically responsible version of Hitler, who sincerely believes that his commanding mass murder is best, to command mass murder. This does not require us to say that Hitler's commanding mass murder was good or to say that we should not criticize Hitler for commanding mass murder. It just requires us to say that if Hitler was epistemically responsible in forming his belief that his commanding mass murder was best, then he did not do anything morally wrong when he commanded mass murder.

If this still sounds implausible, subjectivists can note that cases like the one we are trying to imagine, in which an agent's evidence speaks decisively in favor of his holding some perverse moral belief, are extremely rare and, therefore, that we should not trust our intuitions about such cases. In the actual world, an agent's evidence rarely speaks decisively in favor of his holding some perverse moral belief. So, in the actual world, an agent who holds a perverse moral belief is almost always guilty of moral wrongdoing. This would explain why we are inclined to think that an agent like Hitler, who holds a perverse moral belief, must have done something morally wrong.

Because we cannot be sure whether Hitler believed that his commanding mass murder was best or, if so, how he came to believe it, I have not tried to identify exactly what subjectivists should say about Hitler's commanding mass murder. Instead, I have tried to show that even if Hitler believed that his commanding mass murder was best, subjectivists are not committed to saying that Hitler did nothing wrong. Now, if Hitler was both epistemically responsible and acted in accordance with his beliefs, then subjectivists are committed to saying that Hitler did nothing wrong, but in that case, this verdict is not implausible.

4.4 Could One Avoid Wrongdoing by Failing to Attend to One's Situation?

A third worry of Zimmerman's is that "the Subjective View implies that the failure to believe that any act is best would make it the case that one has no obligations."¹⁰³ If this is true, then one could avoid wrongdoing by simply failing to attend to one's situation (and thereby failing to form beliefs about what is best), which seems absurd. Thus, Zimmerman concludes that subjectivism is false.

My inclination is to deny that it is as easy to fail to form beliefs about what is best as Zimmerman seems to be suggesting. Typically, we cannot help but attend to our situations, and we just find ourselves having beliefs (or at least credences) about what it would be best for us to do in them. Our responsibility at that point is simply to act in accordance with those beliefs.

Nevertheless, Zimmerman is right to worry that there may be cases in which we fail to have beliefs about what would be best and that in such cases, we can sometimes resist forming beliefs about what would be best. In response to this objection, subjectivists can argue that whether we are permitted to resist forming beliefs about what would be best depends on whether we believe that it is best for us to form beliefs about what would be best. This strikes me as plausible. If we think it would be best for us to form beliefs about what would be best, then it seems that we are obligated to form such beliefs. Conversely, if we think that it would not be best for us to form beliefs about what would be best, then it seems that we are prohibited from forming such beliefs.¹⁰⁴

There may, however, be occasions on which agents not only lack beliefs about what would be best but also lack beliefs about whether it would be best for them to form beliefs about what

¹⁰³ Zimmerman 2008: 14.

¹⁰⁴ Jackson 1991: 464-465 makes a related point about how there are times when agents believe it is best for them to go ahead and do what they think is best in the situation and other times when they believe it is best for them to gather more information before doing what they currently take to be the best thing to do in the situation.

would be best. In cases like this, subjectivism entails that agents do not have any moral obligations.¹⁰⁵ Again, this does not strike me as especially implausible. If someone finds herself in a situation in which she does not have a belief about what would best or a belief about whether it would be best for her to form such a belief, then it seems plausible that she does not have any moral obligations in that situation.

4.5 Could One Be Morally Infallible?

Finally, Zimmerman writes, “[T]he Subjective View implies that all moral agents possess a certain kind of moral infallibility. For if ... we always know, whenever we have a belief about what it would be best to do, what it is that we believe, then if the Subjective View were true and we knew this ... we would always know ... what we ought to do. But this makes a mockery of the conscientious person’s inquiry into what he ought to do.”¹⁰⁶ Zimmerman’s worry is that if an agent is infallible with respect to what she believes is best, and she knows the truth of subjectivism, then she would be morally infallible. The problem with this is that no one seems to be morally infallible. We all have moments when we question what we are morally obligated to do. This seems bad for subjectivism.¹⁰⁷

In response, subjectivists can argue that some of our questions about what we are morally obligated to do are attributable to our ignorance about what we believe is best. At times our moral beliefs are dispositional rather than occurrent. Consequently, it takes some work to draw them out.

¹⁰⁵ I am grateful to an anonymous reviewer for drawing my attention to this worry.

¹⁰⁶ Zimmerman 2008: 13-14.

¹⁰⁷ Mason 2013 uses a closely-related objection to reject subjectivism, and Zimmerman’s rejection of subjectivism in his most recent book (2014) depends solely on this objection.

Moreover, subjectivists can argue that we should be more skeptical than Zimmerman is about our ability to know the truth of the correct theory of moral obligation. While we might want to admit that ethicists could know which ethical theory is the correct one, it is not plausible that ordinary individuals could *know* which theory that is. This, along with the fact that our moral beliefs are sometimes dispositional, goes much of the way to explaining why no one seems to be morally infallible.

What I have said so far, however, is consistent with there being some cases (viz., cases in which agents know both what they believe would be best and that subjectivism is true) in which agents *are* morally infallible. This might strike one as implausible. One might object, “It just does not seem like such people would stop asking for advice, even if they knew the truth of subjectivism.” I think that is right, but it is important to note that there is a difference between knowing *what one is morally obligated to do* and knowing *what would be best*. Agents who know what they believe would be best and that subjectivism is true would know what they are morally obligated to do (viz., *what they believe would be best*), but they would not know *what is actually best*. So, even if such agents would stop asking for *moral* advice, it does not follow that they would stop asking for *advice* more generally, especially if they are conscientious (i.e., they care about doing what is best).¹⁰⁸

Consider, for example, a conscientious doctor in Jill’s situation. Such an individual will ask the more knowledgeable doctor for advice about what to do, even if she knows the truth of subjectivism, not because she wants to find out what she is morally obligated to do, but because she believes that it is best for her to cure her patient, and because she believes that she will be more likely to cure her patient if she asks the more knowledgeable doctor for advice (because the more

¹⁰⁸ Such agents would presumably believe that it would be best for them to do whatever they can to form true beliefs about what would be best.

knowledgeable doctor may have information that she does not). This case shows us that even those agents who know both what they believe would be best and that subjectivism is true will still ask for advice.

But certainly agents who know both what they believe would be best and that subjectivism is true will stop asking for *moral* advice, right? I do not think so, and here is why. An agent who has a first-order belief that ϕ -ing would be best can also have a higher-order belief that her first-order belief may be false. For example, a vegetarian may believe that it would be best for her, on this particular occasion, to go ahead and eat the factory-farmed hamburger patty she has been served (because it is socially appropriate and unlikely to cause any additional animal suffering) while at the same time wondering whether she is right about that. She is pretty sure that it is best to maximize utility (and that her eating the factory-farmed hamburger patty she has been served would do that), but she cannot help but wonder whether it would be better for her to abstain. Certainly, we would not be surprised if someone going into a situation like this one were to ask for moral advice, even if that person knew both what she believed was best and that subjectivism is true.

In the end, then, subjectivism does not make a mockery of the conscientious person's inquiry into what he ought to do, for it is able to plausibly explain why even those agents who know both what they believe is best and that subjectivism is true would deliberate about what they are morally obligated to do.

4.6 Conclusion

In earlier chapters, I argued that we have good reasons for moving from objectivism to subjectivism or prospectivism. In this chapter, I have noted that the reasons we have for moving

from objectivism to prospectivism also speak in favor of moving from prospectivism to subjectivism, and I have defended subjectivism against four common objections to it. In light of that, I believe that we have most reason to be subjectivists. If nothing else, subjectivism is underexplored. If philosophers give subjectivism more attention, they may find that it is more attractive than they originally thought. Of course, they might instead hit on an objection that sinks the theory once and for all. We are sure to learn something either way.

CHAPTER 5

IS THERE BLAMELESS WRONGDOING?¹⁰⁹

5.0 Introduction

In Chapters 1-3, I argued that we should reject objectivism in favor of prospectivism or subjectivism. In Chapter 4, I argued that our reasons for rejecting objectivism in favor of prospectivism also speak in favor of our rejecting prospectivism in favor of subjectivism, and I defended subjectivism against four objections. There is, however, a further objection to subjectivism. According to that objection, we should reject subjectivism because it does not allow for the possibility of blameless wrongdoing.¹¹⁰ In response to this objection, I admit that subjectivism does not allow for the possibility of blameless wrongdoing, but I deny that it should. In short, I argue that blameless wrongdoing is impossible. If I am correct, then the fact that subjectivism does not allow for the possibility of blameless wrongdoing should not count against it. In addition, my argument has implications for the debate over the nature of excuses, for according to a plausible account of excuses, excuses just are cases of blameless wrongdoing. So, if my argument is sound, then this account of the nature of excuses is false.

¹⁰⁹ I am grateful to Chris Heathwood for his helpful comments on earlier drafts of this chapter.

¹¹⁰ This objection has been suggested by both objectivists like Peter A. Graham 2010: 94 and prospectivists like Michael J. Zimmerman 2008: 10, 171.

5.1 Blameworthiness

In this section, I argue that an agent is blameworthy for performing an action only if she believed she had a deontically better alternative. To be clear, I am not arguing that it is *morally permissible* to blame an agent for performing an action only if she believed she had a deontically better alternative. That is, I am not arguing it is *useful* (or *worthwhile*) to blame an agent for performing an action only if she believed she had a deontically better alternative. Instead, I am arguing that an agent *deserves* blame for performing an action only if she believed she had a deontically better alternative.¹¹¹ Additionally, I should point out that this is an account of when an agent deserves blame for her *actions*, not an account of when an agent deserves blame for her *character traits* or her *motives*.

To argue for this account of blameworthy action, I begin by presenting a paradigm case of blameworthy action. This helps me identify a number of features that one might think are necessary for blameworthy action. After showing that a number of those features are not actually necessary for blameworthy action, I identify one that is, namely, that the agent believed she had a deontically better alternative. After considering and then rejecting the possibility that this feature is not necessary for blameworthy action, I conclude that an agent is blameworthy for performing an action if and only if she believed she had a deontically better alternative. In the following sections, I use this account of blameworthy action to argue that blameless wrongdoing is impossible.

Consider, what I will call *The Paradigm Case* (of blameworthy action):

¹¹¹ I am therefore assuming that the word “blameworthy” is synonymous with the word “culpable.”

The Paradigm Case

A doctor, Samantha, must decide how to treat her patient, Tom's, skin condition. Samantha does not like Tom, so she gives him drug A, knowing that it will irritate his skin. Unsurprisingly, Tom develops a rash.

In this case, it is clear *that* Samantha's act (i.e., her giving Tom drug A) is blameworthy, but it is less clear *why* Samantha's act is blameworthy. Fortunately, I do not need to answer that question. I do, however, need to identify the necessary conditions for blameworthy action (in order to identify the sufficient conditions for blameless action).

Reflection on *The Paradigm Case* might lead one to suggest (a), that an agent is blameworthy for performing an action only if that act *harms* another. But an agent can be blameworthy for performing an act even if her act does not harm another. Imagine, for example, a case like *The Paradigm Case* but in which Samantha does not know that drug A will irritate Tom's skin. She just knows that it is likely to irritate Tom's skin, and she gives it to Tom in hopes that it will. As things turn out, however, drug A cures Tom. In this case, Samantha is blameworthy for giving Tom drug A even though she does not harm him. Thus, that an agent's act harms another is not necessary for that agent to be blameworthy for performing that act.

To accommodate this counterexample to (a), one might suggest (b), that an agent is blameworthy for performing an action only if that act *endangers* another. But this cannot be correct either, for an agent can be blameworthy for performing an act even if her act does not endanger anyone. Consider, for example, a case in which Samantha believes, in accordance with her evidence, that drug A is likely to irritate Tom's skin. But Samantha is wrong; drug A is sure to cure Tom. In this case, Samantha is blameworthy for giving Tom drug A even though her doing so does not endanger Tom. Thus, that an agent's act endangers another is not necessary for that agent to be blameworthy for performing that act.

Now, it is possible that, regarding this most recent example, one's intuitions conflict with mine. One might think that insofar as Samantha's giving drug A to Tom is sure to cure Tom, Samantha is not blameworthy for giving Tom drug A. Moreover, one might think that this is precisely because her doing so does not endanger Tom. How might I defend my position against this sort of objection? One thing I could do would be to present additional cases in which agents mistakenly believe they are endangering others in hopes that my opponents would have different intuitions about those cases. But this strategy is unlikely to work because it is unlikely that my opponents would have different intuitions about those other cases.

Another way for me to defend my position against this objection would be to appeal to some more general principle. For example, I could note that it is intuitively plausible that moral wrongdoing and blameworthiness are connected, argue that Samantha's giving drug A to Tom is morally wrong (given that she believes, in accordance with her evidence, that her giving drug A to Tom will irritate his skin), and finally infer that Samantha's giving drug A to Tom is blameworthy. This argument is unlikely to be convincing, however, because it only works if the credences one gives to the argument's premises are high enough to justify a credence in the argument's conclusion that is higher than the credence one gives to its denial. And it is unlikely that someone who thinks that that agents are blameworthy only if they endanger others will think that moral wrongdoing and blameworthiness are connected.

There is, however, a more promising strategy available to me, which is to argue that the reason one thinks Samantha is not blameworthy for giving Tom drug A in this most recent example, in which she mistakenly believes that drug A is likely to irritate Tom's skin, is that one thinks blaming Samantha would be useless. Then, I can point out that while it may very well be true that blaming Samantha is useless, it does not follow from this that Samantha does not deserve blame. It seems to me that that this is exactly what is going on in this most recent example; those

who think that Samantha is blameless think that because they think that blaming Samantha would be useless, not because they think that Samantha does not deserve blame.

To accommodate this counterexample to (b), one might propose (c), that an agent is blameworthy for performing an action only if her evidence suggests that it will harm another. The problem with (c), however, is that an agent can be blameworthy for performing an act even if her evidence does not suggest that it will harm another. Consider, for example, a case in which Samantha's evidence suggests that her giving Tom drug A will cure him, but for some reason or another, Samantha believes that drug A will irritate his skin. In this case, Samantha is blameworthy for giving Tom drug A even though her evidence suggests that her doing so will not harm him. Thus, that an agent's evidence suggests that an act will harm another is not necessary for that agent to be blameworthy for performing that act.

To accommodate this counterexample to (c), one might propose (d), that an agent is blameworthy for performing an action only if she believes that it will harm another. But, again, there is a problem with this suggestion, for an agent can be blameworthy for performing an act even if she does not believe that it will harm another. Consider, for example, a case in which Samantha believes that there are two drugs that will improve Tom's condition, drug A, which will relieve Tom's skin condition but will not cure it, and drug B, which *will* cure Tom's skin condition. Because Samantha does not like Tom, she prescribes drug A. In this case, Samantha is blameworthy for giving Tom drug A even though she does not believe that it will harm him. Thus, that an agent believes that an act will harm another is not necessary for that agent to be blameworthy for performing that act.

Finally, to accommodate this counterexample to (d), one might propose (e), that an agent is blameworthy for performing an action only if she believed she had a deontically better alternative. Is (e) true? I believe it is. Of course, I should admit that I did not emphasize this feature

of *The Paradigm Case* when I first described it. Consequently, on that original description, we cannot be sure that Samantha believed she had a deontically better alternative to prescribing drug A. In light of that, one might be skeptical that an agent must have believed that she had a deontically better alternative to the act she performed in order to be blameworthy for performing that act. In defense of my view, then, I would like us to consider a version of *The Paradigm Case* on which Samantha did not believe that she had a deontically better alternative to drug A. For example, imagine that an evildoer threatened to kill an innocent person unless Samantha gave Tom drug A. Samantha believed that the evildoer was telling the truth and therefore believed that her giving Tom drug A was deontically best. She gave Tom drug A, and he developed a rash. In this case, regardless of what the evildoer does or whether it was reasonable for Samantha to believe the evildoer, it is plausible that Samantha is blameless. Why? Because she did what she believed to be deontically best. This suggests that it is plausible that an agent is blameworthy for performing an act only if she believed she had a deontically better alternative.

But let us not forget that, in *The Paradigm Case*, Samantha does not like Tom. Assuming that she does not have a good reason for not liking Tom, we might think that even if Samantha believed that her giving Tom drug A was deontically best (because she believed that an evildoer would kill an innocent person unless she gave Tom drug A), she might still be blameworthy for giving him drug A if, for instance, she did so to harm Tom rather than to save the innocent person. If that is correct, then an agent's believing that she had a deontically better alternative to an act might not be necessary for her to be blameworthy for performing that act; it might be enough that she wanted to harm another.

But even if the fact that Samantha wanted to harm Tom entails that she is blameworthy, it does not follow that she is blameworthy *for giving Tom drug A*. This is because an agent can be

blameworthy for her *character traits* or, in this case, her *motives* without being blameworthy for a particular *action*.¹¹²

The view I have given, on which an agent is blameworthy for performing an action only if she believed she had a deontically better alternative is also endorsed by Michael J. Zimmerman and Gideon Rosen.¹¹³ On this view, agents are blameworthy for performing an action only if they have acted akratically. But this view has been challenged by Elinor Mason who has argued that some agents are blameworthy even if they have not acted akratically.¹¹⁴ In support of this claim, Mason points to what she calls “moral outliers,” individuals who, though they have received the typical moral education, have come to hold repellent moral views. According to Mason, moral outliers “might think it is permissible to aggressively market infant formula in the third world, or to protect oil interests in an unstable region by funding the military oppression of protest.”¹¹⁵ On her view, even if these moral outliers believe that such acts are permissible, they would still be blameworthy for performing them. Moreover, she thinks this is true even if their repellent beliefs are not the products of prior blameworthy actions.

While it is plausible that we *should* blame moral outliers for those actions they perform from their morally repellent views, I think this plausibility comes from the fact that it is plausible to think that blaming moral outliers is *useful* or *worthwhile*, for it is not plausible that such agents *deserve* blame, at least not when their repellent beliefs are not the products of prior blameworthy actions. Mason might agree, for she indicates that the kind of blame that moral outliers deserve, which she calls “objective blame,” is a different kind of blame than that deserved by those who act akratically.¹¹⁶

¹¹²I made this same point back in Section 3.2 in response to a purported case of blameworthy rightdoing given by Graham 2010: 94.

¹¹³See Zimmerman 1997b, 2008: 176, and Rosen 2003, 2004.

¹¹⁴Mason 2015.

¹¹⁵Mason 2015: 3038.

¹¹⁶Mason 2015: 3049.

If that is right, then it is not clear that Mason's view is opposed to mine. But on another way of reading Mason, her argument goes the other way. On this interpretation, Mason is arguing from the intuitive plausibility of the claim that moral outliers deserve blame to the conclusion that agents need not act akratically in order to deserve blame for performing those actions. In response to this kind of argument, I should simply note that I find *the control principle*, according to which one is not responsible for what is beyond one's control, more plausible than I find the claim that moral outliers, at least as Mason has described them, are blameworthy. So, given that the control principle entails that moral outliers are not blameworthy, I am inclined to believe the same.

In the end, however, the arguments I give below do not depend on my being right about this. For neither of the supposed cases of blameless wrongdoing that I consider below is one in which the agent in question (i) believes that she has a deontically better alternative or (ii) is a moral outlier. Thus, regardless of whether one sides with me in thinking that an agent is blameworthy for an action only if she believed she had a deontically better alternative or with Mason in thinking that an agent is blameworthy for an action only if either (i) she believed she had a deontically better alternative or (ii) she is a moral outlier who is acting from moral ignorance, you are able to accept the conclusions I arrive at below.

5.2 Blameless Wrongdoing

Objectivism, remember, is the view that an agent's moral obligations do not depend on her beliefs or her evidence. According to objectivism, or at least objectivist versions of

consequentialism,¹¹⁷ Samantha's giving Tom drug A in *The Paradigm Case* is morally wrong because her doing so will give Tom a rash (which is deontically worse than her doing nothing). This would remain true even if Samantha had believed, in accordance with her evidence, that her giving Tom drug A would cure his condition. Although objectivism is the standard view in normative ethics, it is not without its critics.¹¹⁸ In fact, cases like this one are often used to argue against objectivism. If Samantha believed, in accordance with her evidence, that her giving Tom drug A would cure his condition, then it seems like Samantha was at least morally permitted, if not morally obligated to give Tom drug A. But objectivism entails otherwise, which suggests that objectivism is false.

In one of the earliest and most influential defenses of objectivism, G. E. Moore addresses a case like this one, suggesting that the reason we think it would not be morally wrong for someone like Samantha to prescribe drug A (in a case in which she believes, in accordance with her evidence, that drug A will cure her patient) is because we think that Samantha's prescribing drug A is blameless.¹¹⁹ But that fact that it is blameless, according to Moore, does not rule out the possibility that it is morally wrong. Thus, on Moore's view, it is possible for there to be cases of blameless wrongdoing.

Ever since Moore drew this distinction between moral wrongdoing and blameworthiness, it has been used by objectivists to defend objectivism against criticisms like the one above. But this defense fails. While moral wrongdoing and blameworthiness could, in principle, come apart, there are no good reasons to think that they do.

¹¹⁷ If this example does not threaten objectivist versions of Kantianism or virtue ethics, we could alter it so that it does. Zimmerman discusses how to do this in his most recent book (2014: 39-40).

¹¹⁸ Prichard 2002/1932, Ross 1939, Jackson 1991, Zimmerman 2006, 2008, 2014, and Mason 2013.

¹¹⁹ Moore 1912: 192-193.

5.2.1 Objectivism and Blameless Wrongdoing

Although Moore was one of the first to draw a distinction between *moral* wrongdoing and blameworthiness, the distinction between *legal* wrongdoing and blameworthiness had been around for some time. There are a number of ways to commit a criminal offense but avoid punishment, but two are noteworthy for our purposes. One way for a person to commit a criminal offense but avoid punishment is for that person to show that she was legally *justified* in committing that offense. For example, imagine that Alissa is walking down the road when she hears someone calling from a nearby well. “Help me!” the person calls, “I’m stuck!” Alissa turns to run and help, but as she does, she notices a sign saying, “No Trespassing: Violators Will Be Prosecuted.” Alissa disregards the sign and runs to help the person. Although Alissa’s trespassing would ordinarily warrant punishment, it does not in this case because Alissa’s trespassing is legally *justified*.

Another way for a person to commit a criminal offense but avoid punishment is for that person to show that she should be legally *excused* for committing that offense. For example, imagine that Alissa has unknowingly been drugged with a hallucinogen and is walking down the road when *she thinks* she hears someone calling from a nearby well. Alissa turns to run and help, but as she does, she notices a sign saying, “No Trespassing: Violators Will Be Prosecuted.” Alissa disregards the sign and runs to help the person. Although Alissa’s trespassing would ordinarily warrant punishment, it does not in this case because Alissa’s trespassing is legally *excused*.

Legal justifications and excuses are two types of legal *defenses*. In a legal defense, an agent admits that one has committed an offense but denies the inference from this fact to the conclusion that she deserves punishment. In the case of legal justifications, it seems that an agent admits to having committed a criminal offense but denies that her doing so was legally wrong (i.e., illegal). In the case of legal excuses, however, it seems that an agent admits to having done something legally wrong (i.e., illegal) but denies that her doing so deserves blame. This suggests that we should

think of cases in which an agent's action is legally excused as cases of blameless *legal* wrongdoing. Moreover, we might think that if there are cases of blameless *legal* wrongdoing, then there should also be cases of blameless *moral* wrongdoing (hereafter, simply *blameless wrongdoing*). While this seems to be the strongest argument for the existence of blameless wrongdoing, I think it fails.

To see the analogy, imagine a doctor, Ben, who is tasked with treating his patient, Chrissy's, illness. Ben gives Chrissy drug B, which kills her. In this case, we can imagine Chrissy's family blaming Ben for Chrissy's death. Of course, Ben might be blameworthy, but assuming that he is not, we can imagine him defending his prescribing drug B in one of two ways. On the one hand, he might argue that his prescribing drug B was *morally justified* (hereafter, simply *justified*) by showing, for example, that had he not killed Chrissy, she would have gone on to kill numerous innocent people. On the other hand, Ben could argue that his prescribing drug B should be *morally excused* (hereafter, simply *excused*) by showing, for example, that he believed, in accordance with his evidence, that his prescribing drug B would cure Chrissy.

If Ben argues that his behavior is justified, he would not be arguing that his behavior, though morally wrong, was blameless. Instead, he would be arguing that his behavior was not morally wrong at all. If, however, Ben defends himself by drawing attention to the fact that he believed, in accordance with his evidence, that his prescribing drug B would cure Chrissy, then he would seem to be arguing that his behavior, though morally wrong, was blameless. According to objectivists, this is a paradigm case of blameless wrongdoing. For although Ben's act was not deontically best and was therefore morally wrong, Ben performed the act from ignorance and therefore does not deserve blame.

Subjectivists like myself will agree that Ben does not deserve blame. Our disagreement with objectivists, then, is over whether Ben's act was morally wrong. Why should we think, with objectivists, that Ben's act was morally wrong? Let us consider three possible reasons. First, one

might think that Ben's act was morally wrong because Ben's act was not deontically best. But this argument begs the question; it simply presupposes the truth of objectivism. Second, one might think that Ben's act was morally wrong because we can truly say that Ben should not have done what he did or that Ben prescribed the wrong drug. But as I argued in Chapter 1, the fact that we can truly say that an agent should not have done what he did, or that an agent prescribed the wrong drug, does not entail that it was *morally wrong* for that agent to do what he did. We can see this by noting that although we would say that Ben *should not* have done what he did, we would not say that Ben *violated* or even *failed to meet* his moral obligations. Similarly, while we might say that Ben prescribed the *wrong* drug, we would not say that Ben's prescribing drug B was *immoral*. So, the fact that we can truly say that Ben should not have done what he did or that Ben prescribed the wrong drug is not a good reason to think that Ben's act was morally wrong. In fact, insofar as we cannot say that Ben *violated* or even *failed to meet* his moral obligations, or that Ben's prescribing drug B was *immoral*, this suggests that Ben's giving Chrissy drug B was not morally wrong. Third, one might think that Ben's act was morally wrong because if we had known that drug B would kill Chrissy, we would have been morally permitted to stop Ben from prescribing drug B. But subjectivism can explain this fact just as easily as objectivism, for subjectivism is also consistent with the claim that if we had known that drug B would kill Chrissy, then we would have been morally permitted to stop Ben from prescribing drug B (because, in that case, we would have believed that our stopping Ben from giving Chrissy drug B was deontically best).

In the end, we do not have a good reason to think that Ben's giving Chrissy drug B was morally wrong. In fact, we have a good reason to think that Ben's giving Chrissy drug B was *not* morally wrong. If that is right, then the standard objectivist argument for blameless wrongdoing

fails.¹²⁰ This is not because objectivism gives the wrong account of blameworthiness, for objectivism does not give us an account of blameworthiness at all. Instead, this is because objectivism gives us the wrong account of moral wrongdoing.

Before concluding my discussion of objectivism, I should note that there might be other cases where it is more plausible that an agent acts morally wrongly according to objectivism and yet does not deserve blame. Imagine, for example, that evildoers threaten to kill Ben's father unless Ben kills Chrissy. Moreover, let us assume that Ben believes that his allowing the evildoers to kill his father is deontically best (because, let's say, Ben's father is curmudgeonly while Chrissy is young and full of life). Nevertheless, Ben cannot bring himself to be responsible, even if only in part, for his father's death, so he prescribes drug B, killing Chrissy. In this case, it is plausible that Ben's act is morally wrong according to both objectivism and subjectivism. Moreover, it is plausible that Ben does not deserve blame. Thus, it looks like a case of blameless wrongdoing. Upon reflection, however, we should think that Ben does, in fact, deserve blame, for the intuition that he is blameless seems to depend either on the intuition that blaming Ben is useless, or on the intuition that, had we been in Ben's situation, we would have acted the same way. But neither of these intuitions gives us a good reason to think that Ben does not deserve blame. As I have already mentioned, the fact that blaming someone would be useless does not entail that one does not deserve blame. Similarly, the fact that, had we been in Ben's situation, we would have acted the same way does not entail

¹²⁰ In Graham 2010: 94, for example, Graham describes a case in which an agent mistakenly takes the wrong coat while leaving a party, and he suggests that this is a case of blameless wrongdoing. While the subjectivist will agree with Graham that the agent in this case is blameless, he or she will disagree with Graham's assessment that this agent has done something morally wrong. Assuming, of course, that the agent's mistake was an honest one (or that the agent believed that his or her taking the wrong coat was best), subjectivism entails that there is nothing morally wrong with the agent's behavior. This seems to be the right result, for we would not say of such an agent that he or she *violated* or even *failed to meet* his or her moral obligations. The agent's action might have been *wrong*, but it was not *morally wrong*.

that Ben does not deserve blame. It might make it inappropriate for us to blame Ben, but it would not make it inappropriate for a morally better agent to blame Ben.

Although I have not canvassed all the possible objectivist arguments for the possibility of blameless wrongdoing, I have looked at the more popular ones. Given that these arguments are unsuccessful, we should be skeptical that there are better objectivist arguments for the possibility of blameless wrongdoing.

5.2.2 Prospectivism and Blameless Wrongdoing

Prospectivism, you will remember, is the view that an agent's moral obligations depend on her evidence. On one such view, an agent's act is morally wrong if and only if her available evidence suggests that it is not her best bet. For illustrative purposes, consider a standard case of ignorance:

Ignorance

A doctor, Danielle, must decide how to treat her patient, Ethan's, skin condition. Danielle believes, in accordance with her evidence, that drug C will cure Ethan. She wants the best for Ethan, so she gives Ethan drug C. Surprisingly, it kills Ethan.

According to *objectivism*, Danielle's giving Ethan drug C is morally wrong. It is, nevertheless, blameless because Danielle was ignorant of the fact that drug C would kill Ethan. Thus, according to objectivism, this is a case of blameless wrongdoing. In the previous section, I argued that objectivists are wrong about this. If there are cases of blameless wrongdoing, this is not one of them. Prospectivism gets this right, for according to prospectivism, Danielle's giving Ethan drug C was not morally wrong because her available evidence suggested that her giving Ethan drug C was her best bet.

This is one of the attractive features of prospectivism, that it is better than objectivism at capturing the intuition that moral wrongdoing and blameworthiness are connected. However, prospectivism still leaves room for the possibility of blameless wrongdoing, and this is a mistake. Consider, for example, the following case, which is based on a case described by Holly Smith in her paper “Culpable Ignorance”:

Culpable Ignorance

A doctor, Frank, must decide how to treat his patient, Gianna, an infant in respiratory distress. Although the traditional wisdom suggests that doctors give infants like Gianna high concentrations of oxygen, the latest issue of Frank’s medical journal indicates that this is actually likely to cause blindness. Furthermore, it indicates that Frank can get the exact same benefits with low concentrations of oxygen. Unfortunately, Frank has been too busy golfing to read the latest issue of his medical journal, so he mistakenly believes that it would be best for him to give Gianna high concentrations of oxygen. This is what he does, and his doing so blinds Gianna.¹²¹

Subjectivism, remember, is the view that an agent is morally obligated to perform an act if and only if he believes that it is his best bet. Because Frank believes that his giving Gianna high concentrations of oxygen is his best bet, subjectivism entails that his giving Gianna high concentrations of oxygen is morally right. But this might seem implausible, for it might seem like Frank’s giving Gianna high concentrations of oxygen is morally wrong. The version of prospectivism I described above can get us this latter result (because the evidence available to Frank, which includes the information in the latest issue of Frank’s medical journal, suggests that Frank’s giving Gianna high concentrations of oxygen is not his best bet). So, if our account of blameworthiness entails that Frank is blameless, we will have what might appear to be a case of blameless wrongdoing.

¹²¹ Cf. Smith 1983: 543.

But does our account of blameworthiness entail that Frank is blameless? Isn't the case called "*Culpable Ignorance*" for a reason, namely, because Frank is culpable (i.e., blameworthy) in it? While I think it is plausible that Frank is blameworthy in *Culpable Ignorance*, I want to contend that this is plausible, not because *Frank's giving Gianna high concentrations of oxygen* is blameworthy but because *Frank's failure to read the latest issue of his medical journal* is blameworthy. And our account of blameworthiness (on which an agent is blameworthy for an action only if she believed she had a deontically better alternative) is consistent with this, at least on the assumption that Frank believed that it would have been better for him to read the latest issue of his medical journal than it would have been for him to go golfing.

Of course, we might want to be able to say more than that Frank's failure to read the latest issue of his medical journal was blameworthy. We probably also want to say that Frank is blameworthy for Gianna's blindness, and one might worry that our account of blameworthiness prevents us from saying this (because it prevents us from saying that Frank is blameworthy for the act that directly resulted in Gianna's blindness). In response, it is true that our account of blameworthiness prevents us from saying that Frank is blameworthy for the act that directly resulted in Gianna's blindness. But it does not follow from this that our account of blameworthiness rules out the possibility that Frank is blameworthy for Gianna's blindness, for our account of blameworthiness is an account of the necessary conditions for blameworthy action, not an account of the necessary conditions for being blameworthy for a state of affairs. In light of that, our account of blameworthiness is perfectly compatible with the claim that Frank is blameworthy for Gianna's blindness. In fact, it can explain why Frank is blameworthy for Gianna's blindness if we make the following plausible assumptions: first, that an agent's failing to do what he believes is deontically best is not only necessary but also sufficient for blameworthy action; second, that Frank failed to do what he believed was deontically best when he spent his time golfing instead of reading the latest

issue of his medical journal; third, that agents are blameworthy for any actual harms that would not have occurred had they not performed some prior blameworthy action; and fourth, that Gianna's blindness would not have occurred had Frank not spent his time golfing instead of reading the latest issue of his medical journal.

That being said, we should be careful when we say that an agent is blameworthy for some harm that resulted from one of his prior blameworthy actions, for to say this can imply that the agent is more blameworthy than he actually is. For example, imagine a case in which Heather unjustifiably pushes Isaac. In this case, a number of things could happen. Isaac might be relatively unaffected, or he might trip and fall, in which case he might scrape his knee, break his wrist, or hit his head. Let us assume that Isaac hits his head and dies. In this case, is Heather to blame for Isaac's death? On the account I have just given, Heather *is* to blame for Isaac's death. But even if that is right, it does not follow that Heather is *as* blameworthy as she would have been had she intentionally killed Isaac. As long as we acknowledge this, then we should not hesitate to say that Heather is to blame for Isaac's death or that Frank is to blame for Gianna's blindness.¹²²

What this shows is that cases like *Culpable Ignorance* do not refute the account of blameworthiness I gave early on. And if that is right, then Frank is not blameworthy for giving Gianna high concentrations of oxygen. Furthermore, if Frank is not blameworthy for giving Gianna high concentrations of oxygen, that suggests that there are cases of blameless wrongdoing, for as I admitted earlier, it might seem plausible that Frank's giving Gianna high concentrations of oxygen is morally wrong. So, unless I am willing to admit that there are cases of blameless

¹²² Instead of going this route, one might say that agents are only blameworthy for the harms their available evidence suggested would result from their blameworthy actions or that agents are blameworthy for the harms they believed would result from their blameworthy actions. I appreciate Chris Heathwood's drawing my attention to these alternative views.

wrongdoing, then I must either abandon the account of blameworthiness I gave early on or argue that Frank's giving Gianna high concentrations of oxygen is not morally wrong. Because I think that the account of blameworthiness I gave early on is correct, I want to argue that Frank's giving Gianna high concentrations of oxygen is not morally wrong.

In arguing against objectivism in the last section, I drew attention to the fact that objectivism (erroneously) entails that certain acts, which we do not consider morally wrong, *are* morally wrong. In particular, I drew attention to the fact that objectivism entails that Ben's giving Chrissy drug B is morally wrong, even though Ben would not admit to having *violated* or even to having *failed to meet* his moral obligations. Whether a similar argument will work against prospectivism, however, is less clear, for it seems that Frank might admit that he violated his moral obligations in giving Gianna high concentrations of oxygen. So, if we are going to reject the prospectivist view of wrongdoing, we will need a different argument.

As I see it, that different argument is analogous to the argument that Mason gives for rejecting objectivism. According to Mason, we should reject objectivism because it fails to account for the fact that unreasonably risky behavior is morally wrong,¹²³ but the same thing can be said of prospectivism. For in cases like *Culpable Ignorance*, prospectivism fails to account for the fact that unreasonably risky behavior is morally wrong.

In *Culpable Ignorance*, it would be unreasonably risky for Frank to give Gianna low concentrations of oxygen. Remember that Gianna is an infant in respiratory distress. Frank wants to save her. Given what he believes, he is most likely to save her if he gives her high concentrations of oxygen. If he gives her low concentrations of oxygen, he believes she might not survive. Of course, we could imagine Frank's colleagues suggesting that Frank give Gianna low concentrations

¹²³ Mason 2013: 12.

of oxygen instead of high concentrations of oxygen, but unless Frank's colleagues also explain why his doing so would be superior to his giving her high concentrations of oxygen, Frank's giving her low concentrations of oxygen will, from his perspective, be unreasonably risky. But this is precisely what prospectivism would require him to do. Thus, prospectivism, like objectivism, cannot account for the fact that unreasonably risky behavior is morally wrong.

At this point, we might be left wondering why it might seem plausible that Frank's giving Gianna high concentrations of oxygen is morally wrong. As I see it, this seems plausible because it is plausible that Frank has done *something* morally wrong. But upon reflection, we should see that it is not *Frank's giving Gianna high concentrations of oxygen* that is morally wrong, it is *his failure to read the latest issue of his medical journal* that is morally wrong. Why think that it is Frank's failure to read the latest issue of his medical journal that is morally wrong? Because it is Frank's failure to read the latest issue of his medical journal that is unreasonably risky. *Culpable Ignorance*, then, is not a case of blameless wrongdoing. Although Frank's giving Gianna high concentration of oxygen is blameless, it is not morally wrong.

In this section, I have argued that the standard prospectivist argument for blameless wrongdoing, like the most popular objectivist arguments for blameless wrongdoing, fails. Again, this is not because prospectivism gives us the wrong account of blameworthiness, for prospectivism does not give us an account of blameworthiness at all. Instead, this is because prospectivism gives us the wrong account of moral wrongdoing.

5.3 Conclusion

Above, I argued that neither objectivist arguments nor prospectivist arguments for the existence of blameless wrongdoing are compelling. So, unless there are better arguments for the

existence of blameless wrongdoing, we should deny that blameless wrongdoing is possible. If this is correct, then the fact that subjectivism does not allow for the possibility of blameless wrongdoing should not count against it.

The arguments I have presented in this chapter also have implications for the debate over the nature of moral excuses, for according to a plausible view of moral excuses, an action is morally excused if and only if it is a case of blameless wrongdoing. Erin I. Kelly gives voice to this view, writing, “An excuse establishes that although the agent acted wrongly she should not be blamed, or should not fully be blamed.”¹²⁴

But if the arguments I have presented in this chapter are sound, then either this plausible view of moral excuses is false, or there aren’t any actions that are morally excused. Assuming that we want to preserve the idea that some actions are morally excused, it follows that instead of thinking of morally excused behavior as blameless *wrong*-doing, we should follow J. L. Austin in thinking of morally excused behavior as blameless *bad*-doing.¹²⁵

¹²⁴ Kelly 2013: 244.

¹²⁵ Austin 1956-1957: 2.

BIBLIOGRAPHY

- Andrić, Vuko. 2013a. "Objective Consequentialism and the Licensing Dilemma." *Philosophical Studies* 162: 547-566.
- . 2013b. "The Case of the Miners." *Journal of Ethics & Social Philosophy*, www.jesp.org, discussion note.
- Austin, J. L. 1956-1957. "A Plea for Excuses," *Proceedings of the Aristotelian Society* 57: 1-30.
- Bales, R. Eugene. 1971. "Act-Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure?" *American Philosophical Quarterly* 8: 257-265.
- Catholic Church. 1997. *Catechism of the Catholic Church: Revised in Accordance with the Official Latin Text Promulgated by Pope John Paul II*. 2nd edition. Vatican: Libreria Editrice Vaticana.
- Colorado Revised Statutes (C.R.S.)*.
- Dorsey, Dale. 2012. "Objective Morality, Subjective Morality, and the Explanatory Question." *Journal of Ethics & Social Philosophy*, www.jesp.org, vol. 6 no. 3.
- Driver, Julia. 2012a. *Consequentialism*. New York: Routledge.
- . 2012b. "What the Objective Standard is Good For." *Oxford Studies in Normative Ethics*, Volume 2. Oxford: Oxford University Press, 28-44.
- Feldman, Fred. 1986. *Doing the Best We Can: An Essay in Informal Deontic Logic*. Dordrecht, Holland: D. Reidel Publishing Company.
- . 2006. "Actual Utility, the Objection from Impracticality, and the Move to Expected Utility." *Philosophical Studies* 129: 49-79.
- Finlay, Stephen and Justin Snedegar. 2014. "One Ought Too Many." *Philosophy and Phenomenological Research* 86: 102-124.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press.
- Graham, Peter A. 2010. "In Defense of Objectivism about Moral Obligation." *Ethics* 121: 88-115.
- Hansson, Sven Ove. 2010. "Objective or Subjective 'Ought'?" *Utilitas* 22: 33-35.
- Harman, Elizabeth. 2016. "Morally Permissible Moral Mistakes." *Ethics* 126: 366-393.

- Hedden, Brian. 2012. "Options and the Subjective *Ought*." *Philosophical Studies* 158: 343-360.
- Henning, Tim. 2014. "Normative Reasons Contextualism." *Philosophy and Phenomenological Research* 88: 593-624.
- Howard-Snyder, Frances. 1997. "The Rejection of Objective Consequentialism." *Utilitas* 9: 241-248.
- Jackson, Frank. 1991. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101: 461-482.
- Kelly, Erin I. 2013. "What Is an Excuse?" *Blame: Its Nature and Norms*. Edited by D. Justin Coates and Neal A. Tognazzini. Oxford: Oxford University Press, 244-262.
- Kiesewetter, Benjamin. 2011. "'Ought' and the Perspective of the Agent." *Journal of Ethics & Social Philosophy*, www.jesp.org, vol. 5 no. 3.
- Kolodny, Niko and John MacFarlane. 2010. "Ifs and Oughts." *The Journal of Philosophy* 107: 115-143.
- Lord, Errol. 2015. "Acting for the Right Reasons, Abilities, and Obligation." *Oxford Studies in Metaethics*, Volume 10. Edited by Russ Shafer-Landau. Oxford: Oxford University Press, 26-52.
- Mason, Elinor. 2002. "Against Blameless Wrongdoing." *Ethical Theory and Moral Practice* 5: 287-303.
- . 2013. "Objectivism and Prospectivism about Rightness." *Journal of Ethics & Social Philosophy*, www.jesp.org, vol. 7 no. 2.
- . 2014. "Objectivism, Subjectivism, and Prospectivism." *The Cambridge Companion to Utilitarianism*. Edited by Ben Eggleston and Dale E. Miller. New York: Cambridge University Press, 177-199.
- . 2015. "Moral Obligation and Blameworthiness." *Philosophical Studies* 172: 3037-3057.
- Mill, John Stuart. 2001/1861. *Utilitarianism*. 2nd edition. Edited by George Sher. Indianapolis: Hackett Publishing Company.
- Moore, G. E. 1912. *Ethics*. Oxford: Oxford University Press.
- Norcross, Alastair. 1990. "Consequentialism and the Unforeseeable Future." *Analysis* 50: 253-256.
- Oddie, Graham and Peter Menzies. 1992. "An Objectivist's Guide to Subjective Value." *Ethics* 102: 512-533.
- Parfit, Derek. 2011. *On What Matters*, Volume 1. Oxford: Oxford University Press.
- Prichard, H. A. 2002/1932. "Duty and Ignorance of Fact." *Moral Writings*. Oxford: Oxford University Press, 84-101.

- Railton, Peter. 1984. "Alienation, Consequentialism, and the Demands of Morality." *Philosophy and Public Affairs* 13: 134-171.
- Regan, Donald. 1980. *Utilitarianism and Co-Operation*. Oxford: Oxford University Press.
- Rivera-López, Eduardo. 2006. "Can There Be Full Excuses for Morally Wrong Actions?" *Philosophy and Phenomenological Research* 73: 124-142.
- Rosen, Gideon. 2003. "Culpability and Ignorance." *Proceedings of the Aristotelian Society* 103: 61-84.
- . 2004. "Skepticism about Moral Responsibility." *Philosophical Perspectives* 18: 295-313.
- Ross, W. D. 2002/1930. *The Right and the Good*. Oxford: Oxford University Press.
- . 1939. *Foundations of Ethics*. Oxford: Oxford University Press.
- Sepielli, Andrew. 2012. "Subjective Normativity and Action Guidance." *Oxford Studies in Normative Ethics*, Volume 2. Edited by Mark Timmons. Oxford: Oxford University Press, 45-73.
- Sidgwick, Henry. 1981/1907. *The Methods of Ethics*. Indianapolis: Hackett Publishing Company.
- Smith, Holly. 1983. "Culpable Ignorance." *The Philosophical Review* 92: 543-571.
- . 2010. "Subjective Rightness." *Social Philosophy & Policy* 27: 64-110.
- . 2011. "The 'Prospective' View of Obligation." *Journal of Ethics & Social Philosophy*, www.jesp.org, discussion note.
- Sorensen, Ray. 1995. "Unknowable Obligations." *Utilitas* 7: 247-271.
- Thomson, Judith Jarvis. 2008. *Normativity*. Chicago: Open Court.
- Wedgwood, Ralph. 2006. "The Meaning of 'Ought.'" *Oxford Studies in Metaethics*, Volume 1. Edited by Russ Shafer-Landau. Oxford: Oxford University Press, 127-160.
- Wiland, Eric. 2005. "Monkeys, Typewriters, and Objective Consequentialism." *Ratio* 18: 352-360.
- Zimmerman, Michael J. 1997a. "A Plea for Accuses." *American Philosophical Quarterly* 34: 229-243.
- . 1997b. "Moral Responsibility and Ignorance." *Ethics* 107: 410-426.
- . 2006. "Is Moral Obligation Objective or Subjective?" *Utilitas* 18: 329-361.
- . 2008. *Living with Uncertainty: The Moral Significance of Ignorance*. New York: Cambridge University Press.
- . 2014. *Ignorance and Moral Obligation*. Oxford: Oxford University Press.