

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Culpability for Moral Ignorance and Reasonable Moral Expectations

Permalink

<https://escholarship.org/uc/item/6fw0g97k>

Author

Zinn, Carmen Michelle

Publication Date

2016

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Santa Barbara

Culpability for Moral Ignorance and Reasonable Moral Expectations

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy
in Philosophy

by

Carmen Michelle Zinn

Committee in charge:

Professor Matthew Hanser, Chair

Professor Aaron Zimmerman

Professor Christopher McMahon

December 2016

The dissertation of Carmen Michelle Zinn is approved.

Christopher McMahon

Aaron Zimmerman

Matthew Hanser, Committee Chair

September 2016

Culpability for Moral Ignorance and Reasonable Moral Expectations

Copyright © 2016

by

Carmen Michelle Zinn

ACKNOWLEDGMENTS

I began this project nearly seven years ago. Its completion has been sandwiched in between a full-time faculty position at Cuesta College and caring for my two young daughters, and has only been possible with the support of my dissertation committee, colleagues, and family members.

I owe many thanks to Professor Matthew Hanser, my dissertation advisor. He has read many drafts of these specific chapters, and endured many messier earlier drafts, as I clarified the topic and trajectory of the dissertation. His comments have been clear and direct, and have struck the perfect balance of sharpening my thinking and motivating me to improve my arguments, without being discouraging. I have stepped away from this project many times over the years, sometimes for multiple years at a time, and Professor Hanser has always welcomed me back to it, and happily continued to work on it with me. He has helped me to be a better, more careful philosopher, and modeled what it is to be a mentor. For his patience and tenacity, I will always be grateful.

I also wish to thank Professors Aaron Zimmerman, Chris McMahon, and the late Tony Brueckner. They too welcomed me back to this project after years of absence, and continued to work on it with me. Professor Zimmerman's encouragement gave me the courage and confidence to apply for reinstatement to the Philosophy program in order to complete my degree. His comments on my dissertation challenged me to defend the basic premises of the dissertation, while also clarifying its "fringe" elements. Professor Chris McMahon graciously agreed to replace the late Professor Tony Brueckner on my dissertation committee, although the topic of this dissertation falls outside his area of expertise. And, although Professor Brueckner never read the content of this dissertation, his calm presence reassured me not only of the feasibility of the project, but also of my own abilities to complete it.

My colleagues at Cuesta College, and in particular, Jane Morgan, my division chair, and Pamela Ralston, my dean, provided me time to work on the dissertation by allowing me to count the hours I spent working on it as my required 30 hours of continuing education. Pamela Ralston engaged me in a very honest discussion of the challenges women (particularly those with families) face in academia, and the ways possessing a PhD can help with them. Although brief, this discussion played a pivotal role in motivating me to finish the degree, and is one that I will never forget.

Finally, I wish to thank many members of my family. My mom, Arlys McDonald, has been my cheerleader through this process, ever reminding me that a "good dissertation is a finished dissertation." Her own accomplishments of completing her PhD in Psychology, at a time when not many women did so, of owning a thriving therapy clinic, and being a pioneer in her field have served as a sterling example of what can be accomplished with focus and discipline. I also thank my husband, Jesse Zinn, for shouldering more than his share of child care, especially in the final months of this project, for his unwavering confidence in my ability to finish it, and for completing his dissertation before me. I couldn't let him be "Dr. Zinn," while I was not! Finally, I thank my two beautiful daughters, Ivy and Opal Zinn. They have (reluctantly) sacrificed chunks of their time with me to the completion of this dissertation. Their laughter, cuddles, and delight in simple pleasures have provided me with balance and perspective, and the thought that my success will provide them all the more reason to believe in and strive for their own dreams, has spurred me on.

VITA OF CARMEN MICHELLE ZINN
AUGUST 2016

EDUCATION

University of California, Santa Barbara	Ph.D., Philosophy	2006-2016
University of California, Santa Cruz	M.A., Philosophy	2002-2005
California State University, Chico	B.A., Philosophy	1998-2002

PROFESSIONAL EMPLOYMENT

Part-time Philosophy Instructor, University of West Georgia, 2016-present
Philosophy Instructor (tenure-track), Cuesta Community College, 2012-2014
Associate (part-time) Philosophy Instructor, Mira Costa Community College, 2011
Adjunct Philosophy Instructor, Miramar College, 2011
Online Philosophy Instructor, National University, 2010-2012
Writing Program Instructor, UC Santa Barbara, 2011
Philosophy Instructor and Teaching Assistant,, 2006-2012
Philosophy Instructor and Teaching Assistant, UC Santa Cruz, 2002-2005

FIELDS OF STUDY

Studies in Moral Psychology with Professors Matthew Hanser and Aaron Zimmerman
Studies in Ethics with Professor Matthew Hanser
Studies in Epistemology and with Professors Aaron Zimmerman and Tony Brueckner

ABSTRACT

Culpability for Moral Ignorance and Reasonable Moral Expectations

By

Carmen Zinn

The focus of this dissertation is whether or not (and the conditions under which) agents are culpable for their moral ignorance and actions they perform from it. In recent years, there has been a steady interest in this topic due to its practical and theoretical import. We often find ourselves in the position of judging whether or not agents are culpable for actions they perform from moral ignorance. For example, we often hear of cases in which an agent has not only done something terrible, but also endorses his action as acceptable (thereby exhibiting moral ignorance). However, the background of these agents is sometimes such that we find ourselves wondering if the agent *could* have formed correct moral beliefs. If not, is he culpable for his moral ignorance? If he is not culpable for his moral ignorance, is he culpable for the terrible actions he performed from it? Our answers to these questions not only determine our responses to a significant subset of wrongdoers; they also have broader implications for theories of moral responsibility, and in particular theories concerning the epistemic conditions for moral responsibility.

Those who weigh in on the topic of moral ignorance tend to take one of two extreme positions concerning it. They either hold the “skeptical view,” according to which agents are almost never culpable for their moral ignorance and actions they perform from it, or they hold “culpability views,” according to which agents are almost always culpable for their moral ignorance and actions they perform from it. Proponents of the skeptical view argue that agents are culpable for their moral ignorance only if it is the result of a knowing

mismanagement of their moral beliefs. Agents very rarely knowingly mismanage their moral beliefs, and therefore are very rarely culpable for their moral ignorance or actions they perform from it. Proponents of culpability views argue that agents are culpable for their moral ignorance if it manifests insufficient moral care or objectionable moral attitudes. Moral ignorance almost always does so. Agents are therefore almost always culpable for it, and for actions they perform from it.

I argue that the skeptical view and the culpability views fail to capture the complexity of our intuitive reactions to particular cases of moral ignorance. In light of this fact, and the potential theoretical implications of these views, we ought to reexamine them. In Chapters 2 and 3, I consider the skeptical view and culpability views in-depth. I argue that evaluating these views requires an account of the general conditions under which an agent can be reasonably expected to avoid X, and how this affects her culpability for X. I offer such an account in Chapter 4. I use it to argue that both the skeptical view and culpability views are mistaken. In their place, I offer a more moderate view that vindicates what I take to be our intuitive responses to cases of moral ignorance: although agents often are culpable for their moral ignorance and actions they perform from it, there are genuine and compelling exceptions.

TABLE OF CONTENTS

Part I: An Introduction to the Debate and the Solution	1
1. Moral Ignorance and Culpability	1
1.1 Moral Ignorance, Actions Performed from a Position of Moral Ignorance, and Culpability	4
1.2 The Skeptical Argument	6
1.3 Culpability Views	7
1.4 Fairness, Reasonable Moral Expectations, and Culpability for Moral Ignorance	9
Part II: Culpability for Moral Ignorance: Existing Positions	13
2. The Skeptical Argument	13
2.1 Skepticism about Culpability for Moral Ignorance	13
2.2 The Skeptical Argument	16
2.3 Existing Objections to the Skeptical Argument: Narrowing Its Application	20
2.4 Existing Objections to the Skeptical Argument: Culpability and <i>Knowing</i> Mismanagement of Moral Beliefs	25
2.5 Moving Forward	35
3. Culpability Views	38
3.1 Unwavering Culpability for Moral Ignorance	38
3.2 Harman's Culpability View	42
3.3 Objections to Harman's Culpability View	46
3.4 Non-Volitional Culpability Views	73
3.5 Objections to Non-Volitional Culpability Views	76
3.6 Moving Forward	90
Part III: Reasonable Moral Expectations, the Moral Ignorance Debate, and Culpability for Moral Ignorance	92
4. Reasonable Moral Expectation and Culpability for Moral Ignorance	92

TABLE OF CONTENTS

4.1 Reasonable Moral Expectations and the Moral Ignorance Debate	92
4.2 Moral Expectations: A general description	95
4.3 Towards a General Account of Reasonable Moral Expectations: Common Assumptions	102
4.4 Reasonable Moral Expectations: A New Approach	117
4.5 Defending my account against objections	131
4.6 Reasonable Moral Expectations and Culpability	137
4.7 Reasonable Moral Expectations, Culpability, and the Moral Ignorance Debate	140
4.8 A Moderate and Intuitive Approach to Culpability for Moral Ignorance	168
References	170

PART I: INTRODUCTION TO THE PROBLEM AND THE SOLUTION

Chapter 1

Moral Ignorance and Culpability

There are two main (morally relevant) types of ignorance: factual ignorance and moral ignorance.¹ Non-culpable factual ignorance releases an agent from blame.

Imagine, for example, that Frank acts from non-culpable factual ignorance. His daughter has told him repeatedly that she has no desire to go to college. However, she secretly goes to college, paying her own way, and going to great lengths to conceal this from her father. Her twin brother expresses the wish to go to college, is admitted, and asks Frank to pay for his education. So, Frank pays for his son's college education, but not his daughter's. He justifiably believes that she does not want a college education, and is not pursuing one. Frank does nothing blameworthy in this scenario. Imagine instead that Frank acts from moral ignorance. He believes that it is a sin for women to seek knowledge outside their churches and homes. So, he pays for his son's college education, but not his daughter's, although he is aware that she is pursuing one. Could Frank's moral ignorance potentially release him from

¹ Some speak of "normative ignorance" instead of moral ignorance. I take normative ignorance to be a broader category than moral ignorance. Normative ignorance may include ignorance of any "should-claims," including claims about what would be prudent in a given context. Moral ignorance, on the other hand, includes ignorance about what is right or wrong, and how moral reasons or considerations ought to be weighed in our decision making processes.

blame? Can moral ignorance ever be non-culpable? If so, can it excuse an agent from knowingly, and wholeheartedly committing an action that is, in fact, wrongful?

Recently, there has been significant interest in these questions. Philosophers tend to take rather extreme positions with respect to them, with some holding that agents are almost never culpable for their moral ignorance and actions performed from it, and others holding that agents are always culpable for their moral ignorance and actions performed from it. Advocates for the “skeptical view” offer the “skeptical argument” in support of their position. They hold that moral ignorance, and actions performed from it, are culpable only if they are the result of the agent knowingly mismanaging her moral beliefs. Because the latter is very rare, the argument goes, agents are very rarely culpable for their moral ignorance, or actions performed from it. There are two prominent “culpability views”. Both hold that culpability for moral ignorance does not depend on culpability for past actions; agents can be, and always are, directly culpable for their moral ignorance. One thing that all parties to this debate share in common is the assumption that if an agent is not culpable for her moral ignorance, then she is not culpable for her actions performed from moral ignorance.

Neither the skeptical nor culpability views capture the complexity of our intuitive responses to moral ignorance. Intuitively, agents are sometimes culpable for their moral ignorance and sometimes not. E.g., a physician working in 1940 in the USA might not be culpable for believing that routine sterilization of mentally disabled individuals is morally acceptable. A presently practicing physician certainly would be culpable for this false moral belief. And, sometimes the fact that an agent acted from non-culpable moral ignorance renders blame inappropriate, while other times it does not. E.g. the physician working in 1940 arguably is not culpable for sterilizing his mentally disabled patients, insofar as he is

acting in a manner that he honestly and blamelessly believes is in their best interests. However, we sometimes judge that an agent is not culpable for holding false moral beliefs, but is culpable for acting from them. E.g. we sometimes think that very unfortunate formative circumstances render agents non-culpable for forming false moral beliefs. However, the fact that they act from this non-culpable ignorance does not seem to excuse them from seriously harming other individuals.

Questions about culpability for moral ignorance have both practical and theoretical import. The answers we give to them significantly impact our judgments about, and reactions to, a subset of wrongdoers. They also impact our understanding of the general conditions under which agents are responsible for their actions. It is generally thought that two conditions must be met in order for an agent to be responsible for her action: a control condition and an epistemic condition. Much ink has been spilt over the control condition. Philosophers are now turning their attention to the epistemic condition. While we have a fairly good grip, and consensus, on the ways in which factual ignorance might undermine an agent's responsibility, we lack this with respect to how moral ignorance might do so. As a result, we lack a clear view of the shape the epistemic condition might take.

Because our intuitions concerning culpability for moral ignorance, and actions performed from it, are not captured by current philosophical positions, and because these issues are of significant theoretical and practical importance, we ought to reexamine them. My aim in this dissertation is to do so, in an effort to bring some clarity and consensus to this debate. In what follows, I argue that both the skeptical view and culpability views are mistaken. In Chapters 2 and 3, I consider the skeptical view and culpability views in-depth. I

argue that evaluating these views requires an account of the conditions under which an agent can be reasonably expected to avoid X, and how this affects her culpability for X. I offer such an account in Chapter 4. I use it to argue that both the skeptical view and culpability views are mistaken. In their place, I offer a more moderate view that vindicates our intuitive responses to cases of moral ignorance. In particular, I argue that agents can sometimes be reasonably expected to avoid their moral ignorance and sometimes cannot. Because it is unfair to blame an agent for that which she could not be reasonably expected to avoid, agents are sometimes culpable for their moral ignorance, and sometimes are not. More surprisingly, but also in line with our intuitive judgments, I argue that agents can sometimes be reasonably expected to avoid their actions from moral ignorance, even if they could not be reasonably expected to avoid their ignorance itself. So, in some cases, agents will be non-culpable for their moral ignorance, but not for their actions performed from it.

1.1 Moral Ignorance, Actions Performed from a Position of Moral Ignorance, and Culpability

Moral ignorance differs from factual ignorance in its content. When one is factually ignorant, he lacks empirical information about particular features of his situation and/or action. For example, he might lack the information the woman he intends to marry is his mother. Or, he might lack the information that petting a bird roughly might injure it, or that he is petting the bird in a rough manner. Moral ignorance does not concern empirical facts about one's circumstances or action, but rather the moral value of the latter. A morally

ignorant agent knows, e.g., that he is marrying his mother, or is petting the bird roughly, but has false beliefs, or lacks true beliefs, about whether his action is morally acceptable.²

An agent acts from a position of ignorance, as opposed to merely in conjunction with it, when the agent would not have performed her action, had she not been ignorant.

Consider, for example, Fred, who invested in a particular company's stock because he falsely believes that the company will flourish in the future. Had Fred known that the company would actually flop, he would not have invested in it. Fred is rightly described as acting from ignorance. On the other hand, imagine that Alfred also has the false belief that the company will flourish in the near future, but does not base his decision to buy their stock on this belief. Alfred is extremely rich, and his best friend owns the company in question. Alfred invests in the company in order to support his friend, and although he hopes it will flourish, he would invest it even with the knowledge that it will not. In this case, although Alfred is in fact ignorant, he cannot be correctly said to act from ignorance, since its absence would not have affected his action.

In what follows, I understand an agent to be culpable for X, if other agents, all things equal, would be not be unjustified in blaming her for it, given her relationship to X (e.g. she performed X, X constitutes an objectionable moral attitude on her part).³ Other factors, besides the agent having the right relationship to X (e.g. whether or not the blaming agents

² In Chapter 2, I argue, along with Harman (2011) and Guerrero (2007), that all not forms of moral ignorance are potentially exculpatory. Only false moral beliefs in the permissibility of the act in question, or the lack of true beliefs accompanied by an unawareness of the unjustifiable risk of harm the action poses to others, or that the action has other objectionable qualities (e.g. it would be dishonest, unkind, greedy, and so on), can potentially exculpate.

³ I will not offer an account of the conditions under which blame is justified. It will be enough for my purposes to grant that blaming an agent for X, given her relationship to it, can be unjustified for a number of reasons, and particularly, when it is unfair to blame her for X.

are guilty of a similar transgression), can affect whether or not agents would be justified in blaming her for X.⁴ And, there may be good reasons for blaming someone who is not in fact culpable for her action. E.g. doing so might bring about some social good. I ignore these other factors that impact our justification in blaming agents, and focus exclusively on whether or not an agent could be justifiably blamed for X, based on her relation to it.

1.2 The Skeptical View

I discuss the skeptical view in Chapter 2. Its main proponents are Rosen (2004) and Zimmerman (1997, 2008). They argue that agents can be culpable for their moral ignorance only if it is the result of a culpable mismanagement of their moral beliefs. An agent would mismanage her moral beliefs if, e.g., she failed to consider contrary viewpoints, or to pay attention to counter-evidence. Agents are culpable for mismanaging their moral beliefs, according to the skeptical argument, only if they do so knowingly. Alternatively, if they mismanage their moral beliefs due to moral ignorance (e.g. they do not consider other views because they believe they should not), they are only culpable for doing so, and thereby for the original moral ignorance in question, if *this* ignorance is the result of a knowing mismanagement of their moral beliefs, and so on. Rosen and Zimmerman argue that it is very rare for agents to knowingly mismanage their moral beliefs, and therefore that agents will very rarely be culpable for their moral ignorance or actions performed from it.

Existing objections to the skeptical argument take three main shapes. Some argue against (1) the broad notion of potentially exculpatory moral ignorance employed by the skeptical argument. Some argue against (2) the claim that agents very rarely knowingly

⁴ For a detailed discussion of such factors, see A. Smith (2007).

mismanage their moral ignorance. And, some argue against (3) the claim that agents are only culpable for their moral ignorance if it is the result of a knowing mismanagement of their moral beliefs. I argue that objections to (1) and (2) are successful in narrowing the application of the skeptical argument. However, even with their success, the skeptical argument still suggests that agents are non-culpable for their moral ignorance far more often than we typically think. Avoiding this result depends on the success of objections to (3) the claim that agents are only culpable for their moral ignorance if it is the result of a knowing mismanagement of their moral beliefs. However, (3) rests on the two further claims, (a) that agents cannot be reasonably expected to avoid their moral ignorance unless it is the result of a knowing mismanagement of their moral beliefs, and (b) that agents cannot be fairly blamed for that which they could not be reasonably expected to avoid. Objections to (3) are not successful insofar as they fail to argue against these further claims. Doing so requires an account of the conditions under which an agent can be reasonably expected to avoid her moral ignorance, and how this affects her culpability for it.

1.3 Culpability Views

I consider two prominent culpability views. The first, defended by Elizabeth Harman (2011, 2014), holds that agents are always culpable for their moral ignorance insofar as it always manifests a failure to care sufficiently about things of moral significance. Harman argues that agents are always culpable for the latter. The second view is really a family of views—non-volitional culpability views—which are based on non-volitional views of blame. The latter hold that agents are culpable for X insofar as it manifests an objectionable

moral attitude/quality of will/set of values.⁵ Because moral ignorance always manifests the latter, according to non-volitional culpability views, agents are always culpable for it.

Both Harman's view and non-volitional culpability views hold that agents can be blamed for their moral ignorance, regardless of how they came to develop it. According to Harman's view, because the conditions under which an agent forms her moral ignorance do not bear on whether or not her beliefs manifest insufficient moral care, these conditions do not impact her culpability for her moral ignorance. According to the non-volitional culpability view, because the conditions under which an agent forms her moral ignorance do not impact the fact that it manifests an objectionable moral attitude/quality of will/set of values, they are not relevant to her culpability for it.

Existing objections to Harman's view and to non-volitional culpability views take two main shapes. Some argue that (1) moral ignorance does not always manifest a failure to care sufficiently about something of moral importance, or an objectionable moral attitude. Others argue that (2) moral ignorance is not always culpable, even when it manifests the former. I argue that objections along the lines of (1) are successfully applied to Harman's view and, to a lesser degree, to the non-volitional culpability views as well. In light of these objections, we can conclude that it is at most true that moral ignorance is always culpable *when* it manifests insufficient moral care or objectionable moral attitudes. I argue that existing objections along the lines of (2) are either implausible, or based on intuitions that would be rejected by proponents of culpability views.

⁵ Angela Smith (2005, 2015), Pamela Hieronymi (1994, 2006), Scanlon (1998, 2008) all endorse non-volitional views of blame, and non-volitional culpability views of moral ignorance.

However, I offer my own objection, along the lines of (2). I argue that it is plausible that given their histories (and in particular, their histories of experiencing trauma), some agents could not have been reasonably expected to bring it about that they care sufficiently about certain things of moral importance, or that they avoid holding objectionable moral attitudes. When this is true, agents are not plausibly culpable for their insufficient moral care or objectionable moral attitudes, or the moral ignorance that manifests them. Fully defending this objection requires a positive argument for these claims. This in turn requires a general account of the conditions under which a moral expectation is reasonable or unreasonable, and an account of how the reasonableness of an expectation that an agent X, bears on his culpability for X.

1.4 Fairness, Reasonable Moral Expectations, and Culpability for Moral Ignorance

In Chapter 4, I offer such an account. After arguing against standard approaches to the reasonableness of a moral expectation as tied to the difficulty of complying with the expectation, the capacities of the agent to whom the expectation is directed, or the fairness of the expectation, I argue that expectations are reasonable or unreasonable in the context of the relationships in which they are issued. The contexts of our relationships are defined by the primary characteristics of the relationship in question. The defining feature of the moral relationship (i.e. the relationship in which we stand to other moral agents, qua moral agents) is that it is a relationship of mutual regard; it is one in which all agents have equal standing, and in which the interests of all agents have equal relative importance. With this in mind, I argue that moral expectations are reasonable if they require an agent to (i) to refrain from acting in ways that treat other agents as lacking inherent or equal value and are also (iii) not

pointless, or if they require an agent (ii) to refrain from acting in ways that treat their interests as lacking or being of less relative importance compared to the interests of others, and are also (iii) not pointless.⁶ Moral expectations are unreasonable insofar as they meet neither (i) nor (ii), or insofar as they fail to meet (iii).⁷

I argue that whether or not an agent can be reasonably expected to avoid X is related to the agent's culpability for X insofar as, in most cases, a moral expectation that meets neither (i) nor (ii) requires the opposite of (i) or (ii); these expectations often require an agent to treat others and their interests as if they are of unequal value or relative importance. They are unfair insofar as they fail to treat like cases alike. When we blame an agent for X, we implicitly endorse a demand that she X. For example, if I blame Fred for violating my privacy by reading my emails without my permission, I implicitly endorse the demand that he not to do so. When we blame an agent for failing to do that which she could not have been reasonably expected to do, we often implicitly endorse an unfair expectation; we often implicitly endorse an expectation that (at the relevant time in the past) she either act in a manner that would have treated others (including herself) or their interests as having unequal relative value or importance. In doing so, *we* treat her or others as having unequal value or their interests as if they are of unequal relative importance. *We* treat her or others unfairly. It

⁶ Moral expectations are also unreasonable insofar as they require the agent in question to treat the interests of others as if they are of less weight than their own. Because this is very unusual, I focus on expectations that are unreasonable insofar as they require agents to treat the interests of others as if they are greater weight than their own.

⁷ Notice that the fact that an expectation is unfair in this sense is not a necessary or sufficient condition for its being unreasonable in other contexts. E.g. it is unreasonable to expect friends to donate their organs to us, but not because in doing so, we fail to treat like cases alike. And, it is reasonable for my children to expect me to treat their interests in getting a good education as if they are of more relative weight than the interests of other children in doing so (e.g. by expecting me to pay for them to go to a private school, but not to pay for other children to do so).

is for this reason, at least, that we are often not justified in blaming agents for that which they could not have been reasonably expected to avoid.

I then apply this account to the skeptical view and culpability views. As noted above, the success of the skeptical view (and/or objections to it) depends on (3) the claim that agents cannot be culpable for their moral ignorance unless it is the result a knowing mismanagement of their moral beliefs. The main support for this claim is the further claim that unless an agent's moral ignorance is the result of a knowing mismanagement of his moral beliefs, agents cannot be reasonably expected to avoid it, and therefore cannot be fairly blamed for it. I argue that this is mistaken. Expecting a moral agent to avoid his moral ignorance is typically reasonable insofar as it typically meets the second disjunct of (ii) and also meets (iii); it typically requires him to refrain from acting in ways that treat the interests of others as if they are of less relative importance than his own. This is true even when the agent develops his moral ignorance through an unknowing mismanagement of his moral beliefs. Given this fact, and the fact that there is no independent reason to accept the claim that agents are only culpable for their moral ignorance if it is the result of a knowing mismanagement of her moral beliefs, I argue that we ought to reject this claim, and thereby the skeptical argument.

Culpability views are also mistaken. They rest on the claim that agents can be culpable for their moral ignorance regardless of how they come to have it. However, given the circumstances under which some agents form their moral ignorance, expecting them to have avoided it is sometimes unreasonable. This expectation sometimes meets neither (i) nor (ii), and often requires the opposite of the second disjunct of (ii); it often requires the agent

to have acted in a way that would treat the interests of others as if they are of greater relative importance than her own. When this is true, the agents in question could not have been reasonably expected to avoid their moral ignorance, and we would be unjustified in blaming them for it. Culpability views must be softened to account for this.

As it turns out, a more moderate view of culpability for moral ignorance is in order. Agents can sometimes be reasonably expected to avoid their moral ignorance, and sometimes cannot. Because they are often not culpable when the latter is true, agents are sometimes culpable for their moral ignorance, and sometimes are not.

The results of this dissertation are, I hope, not only a plausible and useful account of culpability for moral ignorance and actions performed from it, but also an account of reasonable moral expectations that can be applied to other types of cases and debates within the literature.

PART II: CULPABILITY FOR MORAL IGNORANCE: EXISTING POSITIONS

Chapter 2

The Skeptical Argument

2.1 Skepticism About Culpability for Moral Ignorance

The bulk of the recent literature on culpability for moral ignorance focuses on the “skeptical argument,” developed primarily by Michael Zimmerman (1997, 2008 Ch. 4) and Gideon Rosen (2004).⁸ The skeptical argument claims that agents are very rarely culpable for their moral ignorance or actions performed from it. It does so by way of four main claims: (A) agents act from (potentially exculpatory) moral ignorance if they act while lacking an occurrent belief concerning the wrongfulness of their action, (B) agents are culpable for their actions performed from moral ignorance only if they are culpable for their moral ignorance, (C) agents are culpable for their moral ignorance only if it is the result of a knowing mismanagement of their moral beliefs, and (D) agents very rarely knowingly mismanage their moral beliefs.

The practical implications of the skeptical argument are very unintuitive; it seems to release the worst of the worst from blame. For example, Zimmerman (2002) uses it to argue that managers of Nazi concentration camps, and officers who knowingly order attacks on masses of civilians, are very likely non-culpable for their actions. These individuals often

⁸ A somewhat similar argument is offered by Vargas (2005); he argues that the “knowledge” requirement is rarely met in tracing cases, and therefore that agents are culpable for their actions far less often than we typically think.

believe they are doing the right thing, Zimmerman claims. And, it is very unlikely that they developed these horrifying beliefs by knowingly mismanaging their moral beliefs. Instead, they likely developed them by listening to and learning from individuals that they trusted, and later dismissing conflicting arguments and views that they (falsely) believed to be unworthy of consideration. Therefore, according to the skeptical argument, they are non-culpable for their moral ignorance, and their actions performed from it.

Even if we take extreme cases like these to be a misapplication of the skeptical argument (e.g. if we think that that in such cases, agents likely act from affected ignorance, if any), the suggestion that most agents who act from moral ignorance are non-culpable, is still unsettling. After all, agents acting from moral ignorance are paradigmatically culpable: they do the wrong thing while wholeheartedly endorsing their action. It would be surprising if it turns out that most agents who fit this description are instead non-culpable. E.g. if Bill simply cannot understand why others believe it is wrong for him to fail to take low-cost measures to protect his employees from serious harm, this suggests that he is more blameworthy, not less.

Nearly all participants in the debate concerning culpability for moral ignorance, and actions performed from it, accept the claim that (B) agents are culpable for their actions performed from moral ignorance only if they are culpable for their moral ignorance.⁹ However, a number of philosophers object to (A), (C) or (D). Harman (2011) objects to (A), arguing that (A) employs too broad a notion of moral ignorance; it is not a lack of true moral belief that is potentially exculpatory, but only false moral beliefs that one's action is morally

⁹ A notable exception is Harman (2011, 2014). I discuss her view in detail in Chapter 3.

permissible. Objections to (D) offer reasons to think that agents knowingly mismanage their moral beliefs more often than we think, e.g. because many instances of moral ignorance are affected ignorance. With some modifications, I take these objections to (A) and (D) to be largely successful in narrowing the scope of application of the skeptical argument. However, even with this success, the skeptical argument still suggests that agents are non-culpable for their moral ignorance far more often than we typically think. To avoid this conclusion, we must turn to objections to (C).

I consider three objections to (C). The first is offered by Elinor Mason (2015). She agrees with Rosen and Zimmerman that, in one sense of blame, agents are not culpable for their moral ignorance if it is not the result of a knowing mismanagement of their moral beliefs. However, she argues that there is another sense of blame according to which they are culpable for their moral ignorance, even if they developed it unwittingly. The second objection to (C) is offered by FitzPatrick (2008). He argues that moral ignorance is culpable if it is the result of actions that the agent could have been reasonably expected to avoid, even if the agent did not foresee that these actions might lead to his future moral ignorance, or that they were otherwise wrongful. The final objection to (C) is offered by Angela Smith (2015). She argues that culpable moral ignorance need not be traced to a past, knowing mismanagement of one's moral beliefs, insofar as culpability for moral ignorance need not be traced at all. Typically, we trace culpability for X to culpability for past actions in order to determine if X reflects an objectionable moral attitude. However, we need not trace in order to determine if moral ignorance reflects an objectionable moral attitude; it amounts to an objectionable moral attitude.

I argue that all of these objections fail to address what I take to be the main motivation for (C) the claim that agents are not culpable for their moral ignorance, unless it is the result of a knowing mismanagement of their moral beliefs. Unless an agent's moral ignorance is the result of a knowing mismanagement of her moral beliefs, one might think, agents cannot be reasonably expected to avoid their moral ignorance, and therefore cannot be fairly blamed for it. In order for objections to (C) to be successful, and therefore in order to block the disturbing conclusion of the skeptical argument, we must argue against the latter. Doing so requires an account of the conditions under which agents can be reasonably expected to avoid their moral ignorance, and how this affects their culpability for it.

2.2 The Skeptical Argument

As noted above, the skeptical argument is developed by Gideon Rosen (2004) and Michael Zimmerman (1997, 2008 ch. 4). Both philosophers begin with the claim that an agent is culpable for her action performed from moral ignorance only if she is culpable for her moral ignorance. Zimmerman (1997, 2008 Ch. 4) holds that this is so because lack of ignorance is a "root requirement" for responsibility (424, 177).¹⁰ Rosen (2004), on the other

¹⁰ Zimmerman (2008) recognizes that this is not an argument for the claim that an agent is culpable for his action performed from moral ignorance only if he is also culpable for his ignorance (177). He confesses that he's unsure how to offer an argument for it, and also points out that the claim is "deeply imbedded" in our blaming practices (with respect to factual ignorance) (178). Moreover, he argues that this cannot be a species of a more general claim, that an agent is culpable for an action flowing from mental state M, only if she is culpable M as well; if this were true, then agents would only be culpable for acting out of anger if they were also culpable for their anger, which seems implausible (177). The difference between acting from anger for which one is non-culpable, and acting from moral (or factual) ignorance for which one is non-culpable, is that in the former case, one can be aware that what one is doing is wrong (2008, 177).

hand, argues that it is unreasonable to blame an agent for acting wrongly from blameless ignorance, since this would require her to act against her (blamelessly false) all things considered judgment about what to do (306).¹¹

The remainder of the skeptical argument is focused on culpability for moral ignorance. Both Zimmerman (1997, 2008 ch. 4) and Rosen (2004) have a broad notion of potentially exculpatory moral ignorance in mind: the lack of an occurrent, true belief that what one is doing is morally wrong.¹² Both also hold that agents are only ever indirectly culpable for their moral ignorance; they are only ever culpable for it by way of being culpable for something else. Zimmerman (1997) argues that we are only ever indirectly culpable for moral ignorance because we only ever have indirect control over it (414-415).¹³ We can only bring it about that we have (or lack) a belief by way of doing something else. Because “responsibility tracks freedom,” and we only ever have indirect control over our moral beliefs, we are only ever indirectly culpable for them (414).¹⁴ Therefore, culpability for moral ignorance must be traced to some other thing, for which we are directly culpable.

¹¹ In order to judge the accuracy of the claim that agents are not culpable for their actions performed from moral ignorance unless they are also culpable for their moral ignorance, we must first determine the conditions under which the latter (if ever) is true. For this reason, I set aside the issue of culpability for actions performed from moral ignorance until Chapter 4.

¹² Zimmerman (1997) defines potentially exculpatory ignorance as failure to believe the truth at the time of acting (417), and in Ch. 4 of his (2008) book, he expressly states that in order to be non-ignorant in the sense required for direct culpability, one must have an occurrent belief that what he is doing is morally wrong (191). Rosen (2004) does not offer a defense of this view of potentially exculpatory moral ignorance. Zimmerman (2008) holds that agents are directly culpable for actions only if their beliefs play a role in their reasons for acting (191). And, he holds that only occurrent beliefs can play such a role (191). As I discuss below, Rosen (2008) adopts a slightly different view of potentially exculpatory moral ignorance.

¹³ Zimmerman (1997) argues that we only have direct control over our decisions (421).

¹⁴ Some challenge the claim that we only have indirect control over our moral beliefs. See, e.g., Montmarquet (1999). Zimmerman (2002) answers this objection directly (484-485).

Rosen (2004) uses the term “derivative culpability” to refer to the similar notion that there are some things for which we are only culpable by way of being culpable for something else (299). We are only ever derivatively culpable for moral ignorance, Rosen (2004) holds, because we are “passive” with respect to our moral beliefs; belief formation and revision are things that happen to or in us, sometimes by way of our active efforts to bring them about (302). Rosen (2004) argues that we are culpable for something with respect to which we are passive only if we are directly responsible for something else which gave rise to it (302-303).¹⁵

So, both philosophers hold that in order to be culpable for moral ignorance, there must be some past action for which the agent is culpable, of which her moral ignorance is the result. Rosen (2004) offers a much more detailed account of what shape these past actions might take. Taking a lead from negligence law, Rosen points out that agents are not culpable for their ignorance simply because they failed to take *some* measure to prevent it (301). This would cast the net of culpability too wide. Instead, they are culpable for it insofar as they failed to take reasonable measures to avoid it; i.e. the measures a reasonably prudent person in their circumstances would have taken. With respect to moral ignorance, Rosen (2004) suggests that these measures are comprised by what he refers to as our “procedural epistemic obligations” (302). He suggests that these are “impossible to codify,” but that they include obligations to do (or to refrain from doing) certain things: to ask certain questions, to take careful notes, to stop and think, to focus one’s attention in a certain

¹⁵ Rosen does not offer an argument for this claim, and recognizes that some might reject it, particularly those who believe that we can be directly culpable for our characters. However, he maintains that if an agent has done everything an agent should do in order to avoid developing a bad character trait, it is unreasonable to hold him to be culpable for it (302-303).

direction, etc.” (302). If we are culpable for failing to know any truth, Rosen claims, it is because we are culpable for having failed to fulfill some of our procedural epistemic obligations. As noted above, Zimmerman’s (1997, 2008) account is not nearly as detailed, although he seems to have something similar in mind. He suggests that culpability for moral ignorance might be traced back to an agent’s carelessness or inconsiderateness (Zimmerman, 1997, Pp. 416).

Rosen and Zimmerman argue that agents are very rarely culpable for their moral ignorance, because culpability for it can very rarely be traced to culpability for some past failure to manage their moral beliefs. Neither denies that agents fail to pay attention, ask questions, to stop and think, or to be considerate or careful. Instead, both hold that it is rare that agents do so while *knowing* they should not. If this is true, then agents are very rarely directly culpable for their failure to manage their moral beliefs well; they fail to do so from ignorance, and the argument applies all over again. Rosen and Zimmerman claim it is very rare for agents to knowingly mismanage their moral beliefs, and therefore, it is very rare that agents are culpable for their moral ignorance, and actions performed from it.¹⁶

To illustrate this, consider the following example. Rachel believes that homosexuality is morally wrong. She has listened to counter arguments, and sometimes even felt their rational pull, but she quickly banished such thoughts, which she believes were put in her mind by Satan. According to the skeptical argument, Rachel is culpable for her false moral belief (and her actions performed from it) only if it is a result of the knowing

¹⁶ Rosen (2004) adds a wrinkle to his argument. He claims that not only is it rare for agents to act akratically in this way, but also that we are unable to reliably identify akratic actions, and therefore should not be confident in our judgments that an agent acted akratically, and is thereby culpable for her action (308). Zimmerman (2008) voices a similar concern (190).

mismanagement of her moral beliefs. However, this is not clearly the case. Rachel falsely believes that she ought to dismiss arguments that conflict with her beliefs about homosexuality. So, she mismanages her moral beliefs from ignorance. She is culpable for mismanaging her moral beliefs only if she is culpable for *this* bit of ignorance. She will be culpable for it only if it is the result of a knowing mismanagement of her beliefs. This line of inquiry will either “bottom out” in such knowing mismanagement, or, as Rosen and Zimmerman suggest, will not, in which case we should conclude that she is not culpable for her moral ignorance. The latter will be true in the preponderance of cases, according to the skeptical argument.

2.3 Existing Objections to the Skeptical Argument: Narrowing Its Application

The conclusion that agents are rarely culpable for their moral ignorance rests, in part, on the claims that (A) agents act from potentially exculpatory moral ignorance if they lack occurrent beliefs that their actions are morally wrong, and (D) agents rarely knowingly mismanage their moral beliefs. (A) offers a very broad notion of potentially exculpatory moral ignorance, and therefore suggests that a host of agents act from potentially exculpatory moral ignorance including: those who simply fail to consider the wrongness of their actions; those who are undecided about the wrongness of their actions; and those who are disposed to believe their actions are wrong, but fail to advert to these beliefs at the time of acting. If agents are only culpable for their moral ignorance if they knowingly mismanage their moral beliefs, and (D) agents very rarely mismanage their moral beliefs, then the host of agents captured by (A) will turn out to be non-culpable.

Elizabeth Harman (2011) argues against (A) the claim that lack of a true belief that one's action is morally wrong constitutes potentially exculpatory moral ignorance. In particular, she argues that only false beliefs that one's action is morally permissible are potentially exculpatory. The mere absence of an occurrent belief that one's actions are wrong is not. In support of her argument, she offers examples in which agents perform a morally wrong action while being uncertain of the wrongness of their action, due to factual ignorance (448).¹⁷ She considers, e.g., a case in which an agent has .5 credence that the substance she is putting in another person's coffee is poison, and .5 credence it is sugar (448). The agent adds the substance despite her uncertainty. Harman (2011) argues that this agent is obviously directly culpable for her action; we need not determine whether or not she is culpable for her uncertainty in order to know this; she does not act from potentially exculpatory ignorance (448). Therefore, Harman argues that one should hold, at most, that not ignorance in general, but only false moral belief of the permissibility of one's action, can exculpate (448).

This conclusion needs to be both softened and broadened. It is very plausible that agents who act with the knowledge that their action poses risk of significant harm to others, while lacking a false belief that it is nevertheless justified, do not act from potentially exculpatory moral ignorance. However, it is not clear that we should conclude from this that false belief in the permissibility of one's action is the only form potentially exculpatory ignorance can take. A proponent of the skeptical argument might argue that we should instead conclude that an agent acts from potentially exculpatory moral ignorance if she lacks

¹⁷ These examples are based on cases offered by Guerrero (2007). In his paper, Guerrero (2007) argues for a number of complex claims, but his overall position is that if an agent is unsure of the moral status of an action, and the action risks significant harm to others, then the agent should not perform it.

a true belief in the wrongfulness of her action, and is not aware that her action poses risk of great, unjustified harm to others.¹⁸ There is nothing contradictory about this view, and it explains the examples offered by Harman.

At the same time, Harman's (2011) objection can be extended. After all, there are many objectionable moral features, such that if an agent is aware her action may have these features, but lacks a true belief in the wrongfulness of her action, it is not plausible she acts from potentially exculpatory moral ignorance. Included amongst these features is the fact that an action is dishonest, unkind, cruel, unfair, and so on. Imagine, for example, that a waiter finds a very generous tip on one of his tables. It has been a very busy shift, and he cannot recall if he served the patrons at this table, or if the waiter who had the shift before him served them. He is aware that keeping the money might be dishonest and unfair, but he doesn't think about the matter long enough to form the belief that it would be wrong to do so. He keeps the money. This agent would likely be culpable for his moral ignorance according to the skeptical argument, but we need not determine whether or not he is culpable for it in order to judge that he is culpable for his action; he does not act from potentially exculpatory moral ignorance.

Proponents of the skeptical argument need not accept the conclusion that the only form of potentially exculpatory moral ignorance is false belief in the permissibility of one's action. However, they must accept a significantly curtailed understanding of potentially

¹⁸In a later paper, Rosen (2008) grants that agents who perform actions while being aware that they pose risk of great harm to others, while lacking a belief that their actions are morally wrong (and, presumably, also lacking a false belief that this behavior is justified), do not act from potentially exculpatory moral ignorance. Rosen (2008) suggests that an agent performs an action from (potentially exculpatory) ignorance if, at the time of acting, she is unaware of the wrong-making features of the action (598). In cases in which an agent acts while being aware that her action risks serious harm to others, she does not act in ignorance of the wrong-making features of her action; she acts recklessly (598, fn. 14).

exculpatory moral ignorance. Moreover, this modified version of Harman's (2011) objection significantly limits the scope of the skeptical argument. Only in cases in which agents act from a false moral belief in the permissibility of their action, or lack a true belief about the wrongfulness of their action, and are not aware that their action poses risk of significant unjustified harm to others, or that it may be otherwise vicious, do we need to determine if the agent is culpable for her moral ignorance in order to determine if she is culpable for her action performed from moral ignorance. So, the skeptical argument only applies to such cases. In all other cases, we need not determine whether or not an agent is culpable for her moral ignorance in order to determine whether she is culpable for her action performed from moral ignorance.

Objections to (D) the claim that agents very rarely knowingly mismanage their moral beliefs are similarly successful in limiting the scope of application of the skeptical argument. At first blush, (D) seems plausible. After all, it seems that very few individuals *want* to have false beliefs, and most agents would therefore not engage in activities that they believe would lead to false beliefs (or lack of true ones). However, upon reflection, it is plausible that many agents have a vested interest in having false moral beliefs insofar as doing so serves their interests. Moody-Adams (1993) argues that in many cases, agents who hold widespread, false moral beliefs benefit from the status quo, and therefore avoid examining whether or not the status quo is morally permissible, revealing "an unwillingness to entertain the possibility that one might be wrong" (122).¹⁹

¹⁹ For arguments that culture can in fact render an agent non-culpably ignorant see Benson (2001), and Calhoun (1989). Benson (2001) argues that some cultural influences can affect the abilities of some agents in such a way as to render them non-responsible for their

For example, she argues that agents living in cultures that endorse slavery remain ignorant because it is in their interest to be ignorant of the wrongness of slavery. Due to the fact that an agent benefits from the practice of slavery, they are reluctant to question or revise their beliefs concerning it (Moody-Adams, 1994, p. 301-302). This argument can be applied to individuals in sub-cultures as well, and not only to those who stand to gain materially through their ignorance (e.g. slave owners), but also to those who are otherwise invested in maintaining mistaken moral beliefs. Some individuals may be reluctant to give up or adopt a moral belief out of love, loyalty, or because they would otherwise be alienated from their communities. For example, an individual who is loyal to a political party, in the way a sports fan might be loyal to a team, might ignore arguments that point to the reasonableness of opposing views, in order to maintain his or her loyalty to the party. Similarly, individuals who are invested in the truth of a particular religion might have a vested interest in ignoring or hastily denying evidence that its central tenants are mistaken.

A plausible interpretation of such agents is that they know their beliefs are, or may be mistaken, or that they ought to adopt a particular belief that they lack, but they purposefully shield themselves from evidence to this effect. As FitzPatrick points out, this is a type of *akrasia*; agents who behave in this way, act in ways that they know they ought to manage their moral beliefs differently, and fail to do so (604). If this is true, then even on Rosen's and Zimmerman's view, these agents are culpable for their moral ignorance. Insofar

beliefs and actions. Calhoun (1989) argues that cultural ignorance is sometimes not culpable due to lack of accessibility of certain advanced ideas.

as it is plausible that many instances of moral ignorance are instances of affected ignorance, the skeptical argument applies to a smaller set of cases than Rosen and Zimmerman claim.

There are perhaps other plausible objections to (A) and (D) that would further narrow the scope of application of the skeptical argument. However, on its own, this is not a satisfying response to the skeptical argument. Even if we narrow the scope of the argument so that it only applies to a small subset of morally ignorant agents, its conclusion is still likely to be repugnant. It will still suggest, e.g. that Nazi' sending Jews to their deaths under the false belief that they are complying with their moral duties are potentially non-culpable for their actions. To block this possibility, we must consider objections to (C) the claim that agents are culpable for their moral ignorance only if it is the result of a knowing mismanagement of their moral beliefs.

2.4 Existing Objections to the Skeptical Argument: Culpability and *Knowing* Mismanagement of Moral Beliefs

There are three main existing objections to the claim that (C) agents are only culpable for their moral ignorance if it is the result of *knowing* mismanagement. (1) Mason (2015) argues that agents can be culpable in a special sense for their moral ignorance even if they did not develop it knowingly. (2) FitzPatrick (2008) argues that Rosen's and Zimmerman's condition is too strong, and instead argues that agents can be culpable for their moral ignorance if it can be traced back to any act/omission that the agent could have been reasonably expected to avoid. (3) Angela Smith (2015) argues that culpability for moral ignorance need not be traced back to a knowing mismanagement of one's moral

beliefs, because culpability for moral ignorance does not need to be traced at all. We employ tracing, she argues, in order to determine whether or not an agent's action reflects an objectionable moral attitude. In the case of moral ignorance, we need not trace in order to determine this; moral ignorance is itself an objectionable moral attitude. I argue that none of these objections will be successful unless they are accompanied by an account of what can be reasonably expected of an agent, and an argument that agents can be reasonable expected to avoid, and fairly blamed for their moral ignorance, even if it was not developed through a knowing mismanagement of their moral beliefs.

(1) Mason's (2015) Objection

Mason (2015) argues that there are two types of blameworthiness: ordinary blameworthiness and objective blameworthiness. She agrees with Rosen and Zimmerman that agents are not culpable for their moral ignorance, in the sense of ordinary blameworthiness, unless the agent developed it by knowingly mismanaging her moral beliefs. However, she argues, agents can be culpable for their moral ignorance in the "objective sense," even if they developed it through an unwitting mismanagement of their moral beliefs.

Mason (2015) contrasts objective blameworthiness to ordinary blameworthiness, and to the objective attitude (outlined by Strawson (1962)). In order to be blameworthy for some action in the ordinary sense, an agent must be a member of the moral community, and must be aware, at some level, of the fact that he is doing something wrong at the time of acting (Mason, 2015, Pp. 12). Agents who are blameworthy in the ordinary sense have the right

moral goals, but do not try hard enough to attain them (Mason, 2015, Pp. 16). When we blame someone in the ordinary sense, we set aside considerations of how an agent came to be the way she is, and instead focus on her actions, and demand amends for wrongdoing (Mason, 2015, Pp. 18-19). In order to be blameworthy for an action in the objective sense, an agent must display a “morally obnoxious” attitude that is attributable to him, i.e. that “stems from and is connected in some loose (not necessarily conscious) way to the agent’s evaluative judgments” (Mason, 2015, Pp. 13). When we blame someone in the objective sense, we think of the agent as bad insofar as he is “warped, or deranged, or morally twisted”; we do not judge them by how well they conform to our values, but instead, we focus “on what they, and their attitudes are like” (Mason, 2015, Pp. 17-18).²⁰ We alienate ourselves from them; “we stand back and disapprove, despise, disavow, we do not engage” (19). This is in contrast to agents with respect to whom we take up what Strawson (1962) coined the “objective attitudes.” We adopt objective attitudes, such as pity and disgust, when we consider an agent to be entirely outside the moral community, and a proper target of management as opposed to blame. Mason (2015) suggests that we adopt the objective attitude when we judge the individual’s actions and attitudes to be a product of their circumstances, and therefore not things with respect to which they exercise agency.

²⁰ Mason (2015) distinguishes her notion of objective blame from traditional accounts of attributability insofar as on traditional accounts, attributability cannot be undermined by an agent’s unfortunate formative circumstances. E.g. on traditional accounts, Robert Harris is blameworthy for his actions, insofar as they are attributable to him, regardless of whether he became who he was through a history of abuse, or by his own will. Mason (2015) argues that objective blame is undermined by unfortunate formative circumstances insofar as such circumstances undermine our judgments that the agent is in fact an agent; objective blame is a response to agents, and is therefore undermined by unfortunate formative circumstances (18).

On Mason's (2015) view, morally ignorant agents will sometimes be culpable for their moral ignorance, and actions performed from it, even if they developed it through an unknowing mismanagement of their moral beliefs. They will not be culpable in the ordinary sense, since they do not knowingly do wrong. However, she argues that they will sometimes be culpable in the objective sense. In particular, agents whose moral ignorance is not explained by their formative (or current) circumstances will be culpable for it, and for their actions performed from it. Many morally ignorant agents will fit this description. For example, contemporary individuals with racist and sexist attitudes will likely fit this description; their bad moral attitudes are unexplained by any past or current circumstances, and therefore objectively blameworthy. So, while Rosen and Zimmerman might be right that such agents are not blameworthy for their actions in the ordinary sense, they are not entirely non-culpable for their beliefs and actions, as Rosen's and Zimmerman's arguments suggest.

There are a couple of main concerns one might have with Mason's (2015) argument. Some might argue that Mason's view lets too many agents off the hook, while others would argue that it is unjustifiably stringent. Mason suggests the following relationship between an agent's culpability and her past and/or present circumstances. Sometimes an agent's past or present circumstances render it extremely difficult for her to adopt the correct moral view (Mason, 2015, Pp. 20). When this is true, we judge that the agent is not actually an agent with respect to this subset of moral beliefs/attitudes; instead she is just a product of her environment (Mason, 2015 Pp. 19-20). Mason does not give us a more principled reason to accept that unfortunate circumstances undermine agency. While it certainly seems true that biological conditions such as disease or mental illness might do so, I think many would balk

at the idea that, e.g., widespread cultural acceptance of a practice can render individuals non-agents with respect to their beliefs and attitudes concerning the practice.²¹ At the very least, Mason (2015) owes us an explanation of why this would be so.²²

Alternatively, Rosen would likely reject the notion that agents can be culpable in any way for their moral ignorance, or actions performed from it, if they developed it through an unwitting mismanagement of their moral beliefs. Indeed Rosen (2008) argues that it is unfair to blame an agent for his objectionable moral attitudes (and actions that flow from them) if he has done nothing blameworthy (i.e. knowing) to develop them (607). There is no reason to think that Rosen would accept that it is fair to blame agents in a different way (i.e. objectively) for their moral ignorance and actions performed from it, even if they have not knowingly brought it about. Unfortunately, Mason (2015) does not argue otherwise. She suggest that we *do* blame such agents, as long as we do not view them as victims of their circumstances, but does not argue that we are justified in doing so. Rosen and Zimmerman take themselves to demonstrate that our present blaming practices are misguided, and therefore would not be surprised that we blame agents for their moral ignorance, even if they developed it unknowingly; they deny that it is fair to do so. We need an argument that it is not.

(2) FitzPatrick's (2008) Objection

²¹ Mason (2015) holds that it can, and argues that the cultural circumstances of ancient slave owners, and sexists in the USA in the 1950's, rendered them non-agents with respect to some of their moral beliefs (19).

²² Whether or not Mason (2015) is correct about this point does not impact her argument against the skeptical argument. Whether agents who face unfortunate formative circumstances are rightly understood as agents, does not impact their conclusion that agents are almost always non-culpable for their moral ignorance.

FitzPatrick (2008) also rejects the claim that an agent is culpable for his moral ignorance, and actions performed from it, only if he developed his ignorance through a knowing wrongdoing. Instead, FitzPatrick argues that agents are culpable for their moral ignorance, and actions performed from it, if they could have been reasonably expected to avoid it. He takes this to be a less demanding requirement for culpability for moral ignorance.

...we do not need to find some relevant past failure done “in full knowledge of every pertinent fact or norm”—some akratic episode in which the agent knowingly acts with “negligence or recklessness in the management of his opinion.” The true condition for culpable normative ignorance is weaker... Our question is whether he is culpable for the moral ignorance that leads to his bad actions. The natural and non-question-begging way to address this question is to ask... What, if anything, could the agent reasonably (and hence fairly) have been expected to have done in the past to avoid or to remedy that ignorance? (FitzPatrick, 2008, pg. 603)

FitzPatrick (2008) suggests that an agent can be reasonably expected to avoid his moral ignorance if he had the mental capacities required for doing so, and did not face any significant social or personal obstacles in doing so. If these conditions are met, and the agent develops moral ignorance anyway, FitzPatrick (2008) suggests that this is due to “...akrasia or due to the culpable, nonakratic exercise of vices such as overconfidence, arrogance, dismissiveness, laziness, dogmatism, incuriosity, self-indulgence, contempt, and so on” (609).

He illustrates his view, and contrasts it to Rosen’s, through an example of Mr. Potter from *It’s a Wonderful Life*. Potter is a ruthless businessman who takes advantage of the desperate situation of the poor in his town by leasing crumbling homes to them at high rents. Potter sees no problem in doing this, and throughout the film celebrates and savors his gains at the expense of others. On the skeptical view, it is likely that Mr. Potter is not culpable for

his moral ignorance of fair and acceptable business practices. Potter falsely believes that his actions are justified, and it is unlikely that he developed these false beliefs through some action which he fully understood to be a wrongful mismanagement of his moral beliefs. However, FitzPatrick (2008) argues that the skeptical view is mistaken. Potter is culpable for his moral ignorance and actions performed from it, not because he developed it through some past akratic action, but because:

Potter could reasonably have been expected, in the circumstances, to take steps that would have corrected his moral ignorance and improved his character but that he instead chose to behave in ways that merely indulged and reinforced his character defects. By doing so, he thus incurred responsibility both for continued exercises of those dispositions (as well as their becoming hardened in his character) and for the normative ignorance that was preserved or grew and eventually led to the bad business practices we were originally concerned with (FitzPatrick, 2008, pg. 608).

A number of philosophers have responded to FitzPatrick's argument. Harman (2011) argues many agents may do everything that they could be reasonably expected to do in order to avoid their moral ignorance, and yet fail to avoid it (455). In fact, she thinks this is true with respect to most cases of moral ignorance. Ethics, as Harman (2011) puts it, is hard (455). If she is right, even if we grant that agents are culpable for their moral ignorance if they could have been reasonably expected to avoid it, this condition will not be met in most cases, and FitzPatrick's argument would not seriously challenge the skeptical argument.

It is possible that this objection could be addressed by a better, more detailed account of the conditions under which an agent can be reasonably expected to avoid her moral ignorance. Such an account would, I think, suggest that agents in the cases Harman has in mind could be reasonably expected to avoid their moral ignorance. However, even if such an account were offered, FitzPatrick's (2008) general view faces a significant objection from

Rosen (2004) and Levy (2009).²³ They hold that as long as an agent has not knowingly mismanaged her moral beliefs, she could not be reasonably expected to avoid her moral ignorance. Such agents do not do anything they believe to be morally wrong when managing their moral beliefs, nor do they fail to do anything they believe they ought to do with respect to them. Instead, they manage their moral beliefs in exactly the way they believe they ought. In order to avoid developing their moral ignorance, they would either have to act against their own best judgment, or they would have to do something they see no reason to do. It is not reasonable to expect an agent to act in either manner, and therefore it is not reasonable to expect an agent to avoid her moral ignorance, unless she developed it unknowingly. If this is right, then even if FitzPatrick is right that agents are culpable for their moral ignorance as long as they could have been reasonably expected to avoid it, Rosen and Zimmerman are also right that all culpable moral ignorance must “bottom out” in an akratic action/knowing mismanagement of one’s moral beliefs.

In fact, I think this argument is mistaken; it is not unreasonable to expect an agent to do what seems irrational from her own perspective. However, in order to argue for this claim, we need a specific account of the conditions under which an agent can be reasonably expected to avoid X.

(3)Angela Smith’s (2015) Objection

²³ See Rosen (2004), pgs. 602-603. Levy (2009) argues for these claims at length, in response to FitzPatrick’s (2008) argument that agents are culpable for their moral ignorance if it can be traced to some past action which they could have been reasonably expected to avoid. Levy (2009) accepts this claim, but argues that agents cannot be reasonably expected to avoid their moral ignorance unless it was developed through a knowing mismanagement of their moral beliefs.

I'd like to consider one final objection to the claim that (C) agents are culpable for their moral ignorance only if it is the result of a past, knowing failure to manage their moral beliefs well. This objection is raised by Angela Smith (2015). Unlike the other two objections considered to this claim, Smith (2015) does not focus on the claim that agents must *knowingly* mismanage their moral beliefs in order to be culpable for their resulting moral ignorance, but on the more general claim that culpability for moral ignorance must be traced to culpability for some past action. Her target is not Rosen's and Zimmerman's argument, but the more general view that an agent's culpability for her moral attitudes must be traced to her culpability for some past voluntary act or omission. However, her objection is easily applied to the skeptical argument.

Typically, we employ tracing to explain why agents are culpable for their actions over which they lacked voluntary control. For example, we trace culpability for drunk driving to culpability for drinking recklessly. Smith (2015) argues that the justification for tracing in such cases is that actions over which an agent lacks voluntary control, do not obviously reflect her practical judgments; it is as if the movements of one's body just happens to her (124). In order for an agent to be culpable in a case in which he lacks voluntary control over his actions, it must be possible to trace in order to "find a judgment-sensitive locus of responsibility that links the agent to the thing for which he is being morally criticized" (Smith, 2015, pg. 124). E.g. in the case of a drunk driver, we hold the agent to be culpable because her action of deciding to drink without ensuring in advance that she will not harm others while drunk, reflects an objectionable attitude (i.e. a lack of concern for others). Smith (2015) argues that this justification for tracing does not apply to attitudes.

Attitudes, “*by their very nature*, are judgment-sensitive states, so there is no need to try to ‘trace back’ to some *other* judgment-sensitive locus of responsibility” (124).

The skeptical argument holds that we must trace culpability for moral ignorance to past, knowing mismanagement of one’s moral beliefs. Rosen’s and Zimmerman’s support for this claim rests on the claim that we can never be directly culpable for our moral beliefs, but only for the actions that result in them. Smith’s (2015) argument can be taken as an objection to the latter claim. On her view, our moral beliefs (or lack thereof) constitute or embody objectionable practical judgments, so there is no need to trace culpability for them to culpability for past actions. It is certainly true that false moral beliefs, and some instances of moral ambivalence, constitute objectionable moral attitudes. E.g., a false belief that it is permissible to torture animals constitutes an objectionable practical judgment, and ambivalence over whether or not it is permissible to torture animals embodies our practical judgments, particularly about the worth of non-human animals. This is less clear with respect to the lack of a true moral belief. E.g. it seems like we must trace in order to determine what the lack of a true belief that one’s discipline methods are too harsh, reflects about him.²⁴ However, even if we only grant that we need not trace in order to determine that an agent is culpable for false moral beliefs, or for some instances of moral ambivalence, Smith’s (2015) argument still presents a powerful objection to the skeptical argument. Many cases of moral ignorance are ones in which an agent acts from either a false moral belief or moral ambivalence.

²⁴ If, e.g., it just never occurred to a mother to question the appropriateness of her discipline methods, we *might* think this reflects a lack of due care for her children, but, depending on the circumstances, we may not.

A proponent of the skeptical argument might argue that this understanding of the justification of tracing is mistaken. Instead one could offer an alternative account, drawn from Rosen's (2004) and Levy's (2009) argument discussed above. One could argue that the reason we must trace culpability for moral ignorance, to culpability for a past action, is not to locate a judgment-sensitive locus of responsibility, but instead to determine whether or not an agent could have been reasonably expected to avoid her moral ignorance. The agent is culpable for her moral ignorance only if the latter is true.²⁵ If we must trace culpability for moral ignorance to a past, culpable action, then we must trace in cases of moral ignorance, and moreover, as the skeptical argument suggests, we must trace culpability to a past, knowing mismanagement of one's moral beliefs.

Defending either this view or Smith's (2015) view, again, requires that we have an account of the conditions under which agents can be reasonably expected to avoid their moral ignorance, and its relationship to their culpability for it.

2.5 Moving Forward

I have not argued in this chapter that the skeptical argument (and thereby the skeptical view) is mistaken. Instead, I have argued that we either must settle for narrowing its scope of application, or we must effectively argue against that claim that agents are culpable for their moral ignorance only if it is the result of a knowing mismanagement of their moral beliefs. Objections to the latter fail insofar as they fail to address the claims on

²⁵ As noted above, Rosen (2004) argues that we must trace culpability for moral ignorance to a knowing mismanagement of one's moral beliefs because, if the latter is impossible, then an agent could not have been reasonably expected to avoid her moral ignorance (302). In a later article, Rosen (2008) argues that it is unfair to hold agents culpable for what happens to them, in absence of previous (knowing) wrongdoing (609). So, I think this understanding of the justification of tracing is consistent with his view.

which the latter rests: agents can be reasonably expected to avoid their moral ignorance only if it is the result of a knowing mismanagement of their moral beliefs, and agents are only culpable for their moral ignorance if they could have been reasonably expected to avoid it. Evaluating these claims, I hold, requires an account of the conditions under which an agent can be reasonably expected to avoid X, and how this affects their culpability for X.

Some might wonder if such an account is in fact necessary. Some might argue that the claim that an agent could not be reasonably expected to avoid X, simply amounts to the claim that she is not culpable for not X'ing. If this is true, then it is a mistake, of course, to attempt to build an account of culpability for moral ignorance on an account of the conditions under which an agent could be reasonably expected to avoid it. However, the reasonableness of an expectation, and an agent's culpability with respect to it are distinct. If the claim that an agent could not be reasonably expected to avoid X, simply amounts to the claim that he is not culpable for Xing, it should also be the case that if an agent can be reasonably expected to avoid X, she is culpable for Xing. But, this is not true. We can be non-culpable for doing that which we could be reasonably expected to avoid. For example, we can reasonably expect a fully-abled adult to avoid eating so much that his stomach aches. However, it does not follow from this that he is culpable for eating that much, if he has no obligation not to do so, and doing so is not otherwise morally wrong. Moreover, non-moral expectations can be reasonable or unreasonable. E.g. it might be unreasonable to expect students to refrain from all spelling and grammar mistakes. But, this does not amount to saying that they are not morally culpable for failing to do so.

Alternatively, one might wonder why we need to spell out the conditions under which an agent can be reasonably expected to avoid X. One might think we have a good

intuitive grasp of what can be reasonably expected of an agent, or that we could look to the legal literature for an account of reasonable expectations. However, we are often conflicted concerning what an agent could be reasonably expected to do. E.g. it is not clear whether we can reasonably expect an agent to do something that seems irrational to her, or that we can reasonably expect agents with very unfortunate formative backgrounds to avoid false moral beliefs. The account of reasonable expectations offered from the legal perspective relies on our intuitions about what a reasonable person, in the specific circumstances of the agent in question, would do. Moreover, it is not clear that an account of what can be reasonably expected of a member of a political community is the same as what can be reasonably expected of a moral agent.

I argue in Chapter 3 that evaluating culpability views also requires an account of reasonable moral expectations and their relationship to culpability. Such an account is the proper focal point of the debate, and moving forward requires that we offer one.

Chapter 3

Culpability Views

3.1 Unwavering Culpability for Moral Ignorance

Participants in the moral ignorance debate tend to take up one of two extreme positions concerning agents' culpability for their moral ignorance. They either claim that agents are almost never culpable for their moral ignorance (skeptical view), or that agents are almost always culpable for it (culpability views). In the last chapter, I argued that the skeptical view fails to capture our intuitions concerning culpability for moral ignorance, and is inadequately defended. In this chapter, I focus on culpability views, and argue that they face similar problems.

Proponents of culpability views argue that agents are almost always directly culpable for their “pure” moral ignorance—i.e. moral ignorance not rooted in non-culpable factual ignorance. Culpability for pure moral is independent of culpability for any acts or omissions that contributed to it. This is because pure moral ignorance almost always manifests an objectionable moral attitude, and we are always directly culpable for such attitudes and the beliefs (or lack thereof) through which they are manifested.

In what follows, I consider two prominent culpability views. The first is offered by Elizabeth Harman (2011, 2014). She argues that pure moral ignorance always manifests a failure to care sufficiently about things of moral significance. We are always directly culpable for failing to care sufficiently about things of moral significance, and for beliefs (or lack thereof) which manifest this failure. We are therefore always directly culpable for our pure moral ignorance. The second culpability view, the non-volitional culpability view, is

drawn from non-volitional views of blame (defended by Scanlon (1998), Hieronymi (2004), and Angela Smith (over a series of papers)). Proponents of these views endorse one of the two following views of blame. (A) Blame is comprised of a judgment that an agent has an objectionable moral attitude (or holds beliefs that manifest an objectionable moral attitude). Or, (B) when we blame an agent for X, we demand that the agent defend the moral attitudes manifested by X. On (A), blame is justified if it is true that the agent has an objectionable moral attitude, or manifested an objectionable moral attitude in her actions or beliefs, and given our evidence, we are justified in making this judgment. On (B), blame is justified if the agent cannot defend the moral attitudes manifested by X, which will be true if they are, in fact, objectionable. Proponents of non-volitional views of blame are committed to non-volitional culpability views concerning moral ignorance. Because pure moral ignorance almost always manifests an objectionable moral attitude, agents are almost always justifiably blamed for it.

I take it that much is appealing about culpability views' position that agents can be directly culpable for their moral ignorance. In holding this, culpability views accommodate our intuitions that agents are culpable for their moral ignorance, not only because it is the result of mismanagement of their moral beliefs, but because of the nature of their moral ignorance itself. For example, we might rightfully blame an agent who holds racist beliefs in the face of obvious countervailing evidence. However, she is culpable not only because she has mismanaged her moral beliefs, but because her beliefs are morally offensive—they manifest a failure to respect other human beings. Moreover, barring any abnormal circumstances (e.g. that the agent was brainwashed and holds her beliefs as a result), it is plausible that she would be culpable for her racist beliefs even if they were not the result of

a mismanagement of her moral beliefs; she is culpable for them simply in virtue of their offensive nature.

At the same time, the claim that agents are almost always culpable for their pure moral ignorance insofar as it manifests an objectionable moral attitude, is not as appealing. Many tend to think that how an agent came to hold her false moral beliefs (or to lack true ones) is relevant to her culpability for it. In particular, given the circumstances under which an agent forms her moral ignorance (e.g. against the backdrop of very unfortunate formative circumstances), she is sometimes non-culpable for it. This is not something that proponents of culpability views can accept. The circumstances under which an agent came to be morally ignorant do not bear on the fact that pure moral ignorance manifests a failure to care sufficiently about something of moral significance. Therefore these circumstances do not bear on an agent's culpability for her moral ignorance, according to Harman's view. Similarly, the circumstances under which one came to be morally ignorant do not bear on whether or not her moral ignorance manifests objectionable moral attitudes. They therefore do not bear on the truth of our judgment that an agent's moral ignorance manifests objectionable moral attitudes, or on the ability of the agent to defend these attitudes. So too, according to non-volitional culpability views, the circumstances under which an agent came to be morally ignorant do not bear on her culpability for her moral ignorance.

In what follows, I neither defend nor object to the claim that agents can be directly culpable for their moral ignorance. Instead, I argue against the claim that agents are almost always culpable for their pure moral ignorance, insofar as it almost always manifests an objectionable moral attitude or failure to care sufficiently about things of moral significance.

In particular, I challenge the imbedded claims that (1) pure moral ignorance always manifests a lack of sufficient care or an objectionable moral attitude, and (2) that when pure moral ignorance manifests a lack of sufficient care or an objectionable moral attitude, it is thereby always culpable.

I argue that Harman is mistaken in holding that pure moral ignorance always manifests a failure to care about things of moral significance. On a sufficiently complex understanding of human psychology and the phenomenon of caring for something, agents may care sufficiently about things of moral significance, but fail to form true moral beliefs about them. At best, it is the case that pure moral ignorance very often manifests a failure to care sufficiently about things of moral significance. Proponents of non-volitional culpability views are not committed to the claim that pure moral ignorance *always* manifests an objectionable moral attitude, and I argue that they are in a good position to defend the claim that moral ignorance almost always does so. But, even they must grant that there are more exceptions to this rule than they might originally be apt to admit.

Of much greater concern is (2) the claim that when pure moral ignorance manifests an insufficient care or an objectionable moral attitude, it is thereby always culpable. Harman and proponents of non-volitional culpability views fail to fully defend (2). I argue that it is plausible that some agents cannot be reasonably expected *to* care sufficiently about things of moral significance, or to refrain from holding objectionable moral attitudes. When this is true, it is plausible that agents are non-culpable for their lack of sufficient care, or objectionable moral attitudes, and for beliefs (or lack thereof) which manifest them.

In order to defend (2) against these objections, proponents of culpability views must defend one of two further claims. They must either defend (i) the claim that agents can

always be reasonably expected to care sufficiently about things of moral significance or to avoid holding objectionable attitudes, or (ii) the claim that agents can be culpable for that which they could not be reasonably expected to avoid. Defending either claim requires an account of the conditions under which an agent can be reasonably expected to X, and the relationship between whether or not an agent can be reasonably expected to X and her culpability for X'ing (or for not X'ing). I offer such an account in Chapter 4, where I argue that (i) and (ii) are indefensible, and therefore that culpability views are also indefensible in their current formulations.

3.2 Harman's Culpability View

Elizabeth Harman (2011, 2014) develops her culpability view in response to the skeptical views of Rosen and Zimmerman, discussed in Chapter 2.²⁶ She rejects the notion that agents are only ever indirectly culpable for their moral ignorance, and instead holds that false moral beliefs can be the loci of original culpability. There are moral norms that govern not only the management of our beliefs, but also norms that govern our beliefs themselves (Harman, 2011, p. 459). The core premises of Harman's argument are (1) "Beliefs (and failures to believe) are blameworthy if they involve inadequately caring about what is morally significant [and also relevant to our actions]," and (2) pure moral ignorance always involves inadequately caring about what is morally significant.²⁷ From these two premises, Harman concludes (3) that pure moral ignorance is always culpable.

²⁶ In both articles, Harman defends a complex view concerning culpability for moral ignorance, and culpability for actions performed from moral ignorance. I discuss only the former in this chapter.

²⁷ See Harman (2011) p. 460, and Harman (2014), p. 13. Harman does not believe we are blameworthy for lacking beliefs about things that are not at all relevant to our actions. E.g., I

Premise (1) involves two separate claims: (a) beliefs (and lack thereof) sometimes involve a failure to care sufficiently about things of moral significance, and (b) agents are culpable for their beliefs when they involve such a failure. In support of (a), Harman (2011, 2014) argues that holding a belief by way of some consideration, is a way of caring about that consideration (460, 14). For example, believing that it is wrong to use animals to test the safety of cosmetics, based on the consideration that doing so causes animals undue suffering, is a way of caring about the undue suffering of animals. We form (or fail to form) moral beliefs based on our consideration of features of the moral landscape. Because of this, our moral beliefs (or lack thereof) can involve a failure to care sufficiently about these features.

It is not clear what Harman has in mind by a belief “involving” a failure to care sufficiently about something of moral significance. There are at least three things she might have in mind. She might have in mind that a belief (or lack thereof) is caused by a failure to care sufficiently about something of moral significance. E.g., she argues that Huck Finn does not care sufficiently about Jim’s humanity because his care for it does not move him to form a true moral belief about it (Harman, 2014, p. 20). It is also possible that she holds that a belief (or lack thereof) reflects or expresses a failure to care sufficiently. Alternatively, she may hold that a belief (or lack thereof) is an instantiation of a lack of sufficient care. E.g. in her example of a Mafioso who believes it is morally right to kill an innocent shop keeper (discussed directly below), she states that to hold this false moral belief is “to hold that the

have no belief concerning the morality of a senator reading every letter of every bill that comes across her desk. Because this is not at all relevant to my own actions, on Harman’s view, I am not blameworthy for lacking a moral belief with respect to it. For the remainder of the dissertation, I will assume this qualification when discussing Harman’s view, and will not explicitly reference it.

shop owner's life is cheap and can permissibly be sacrificed to his family's own goals," i.e. it *is* to care insufficiently about it (Harman, 2014, p. 14). Her view that believing something based on a consideration is "way of caring about it," also suggests a reading on which our beliefs about something are instantiations of our care about it. I refer to beliefs as "manifesting" insufficient care, in order to capture all of these senses in which a belief might "involve" a lack of sufficient care about something of moral significance.

When a belief (or lack thereof) does manifest a lack of sufficient care, Harman holds we are culpable for it insofar as we are always culpable for failing to care sufficiently about things of moral significance, and for manifestations of this lack of care. Harman does not offer a defense of this claim, but she does point out that our intuitions suggest that agents are always culpable for failing to care sufficiently about things of moral significance. E.g., she points out that we are inclined to hold psychopaths to be culpable for failing to care sufficiently about others, even if we grant that they are unable to do so, given their psychological constitution (Harman, 2014, p. 20).²⁸

Harman also does not offer a sustained defense of (2) the claim that pure moral ignorance always manifests a failure to care sufficiently about things of moral significance.²⁹ However, she offers examples in support of it. As noted above, in one of her examples, a mafia member kills a shop owner who refuses to pay a weekly extortion fee (Harman, 2014, p. 10-11). The Mafioso falsely believes he is doing the right thing insofar as he is acting to

²⁸ Harman argues that if psychopaths are non-culpable for their lack of sufficient care, this is because there may be a requirement on responsibility (i.e. for counting as a responsible agent) which they fail to meet—"the ability to be moved by any moral considerations at all" (20).

²⁹ This will of course be true if Harman holds that in order to care sufficiently about something of moral significance, we must hold true moral beliefs with respect to it. It is not completely clear that she does in fact hold this to be true. See my discussion of her analysis of Huck Finn below.

protect the financial well being of his own group. This man's beliefs manifest a failure to care sufficiently about something of moral significance, i.e. the shopkeeper's life. She believes that all pure moral ignorance will manifest such a failure. This is in contrast to moral ignorance that is based on non-culpable factual ignorance, which Harman (2014) holds will not manifest a failure to care sufficiently about something of moral significance (13-14). E.g., if it is someday established that even embryos experience pain, then the belief that it is permissible to perform very early term abortions without administering embryonic anesthesia or other pain relief, may turn out to be false. However, this false moral belief will not manifest a failure to care sufficiently about fetal pain.

If Harman is right, then whether or not an agent is culpable for his pure moral ignorance is not, as proponents of the Skeptical Argument claim, simply a matter of whether or not he has fulfilled his obligations to manage his moral beliefs well. On her view, agents are culpable for their moral ignorance (at least) because it manifests a failure to care sufficiently about something of moral significance, and this is true even if one has managed her moral beliefs well, since one can manage his beliefs well, and still arrive at the wrong conclusion. In fact, Harman (2011) argues that this happens quite often. She holds that it is true of "most (if not all) ordinary moral claims" that an agent may do everything we could reasonably expect of him in order to avoid moral ignorance (i.e. he may manage his moral beliefs well), and yet "come to deeply false moral views" (455). For example, "for many business practices that are in fact reprehensibly ruthless, we can find plenty of businesspeople ready to offer elaborate, sustained, and serious moral defenses of them" (Harman, 2011, p. 454). Moreover, these individuals often, "have thought carefully about the questions at issue and they have not violated any procedural norms" (Harman, 2011, p.

454). However, the fact that one has thought carefully about whether or not his ruthless business practices are justified does not entail that his beliefs do not manifest a failure to care sufficiently about things of moral significance.

3.3 Objections to Harman's View

Existing objections to Harman's argument target both premises of it: (2) pure moral ignorance always manifests a failure to care about things of moral significance, and (1) beliefs (and lack thereof) are culpable when they manifest a failure to care sufficiently about things of moral significance. Wieland (2015) argues against (2) the claim that pure moral ignorance always manifests a failure to care about things of moral significance. Existing objections to (1) are grounded in claims about an agent's epistemic situation. Harman (2014) points out that some might argue that sometimes false moral beliefs are justified and therefore non-culpable, even if they do reflect a failure to care sufficiently about things of moral significance. Others (e.g. Wieland (2015)) argue that not all agents are in an epistemic position that enables them to form true moral beliefs; some agents lack access to evidence that counts in favor of true moral beliefs. If this is right, some might argue that such agents are non-culpable for failing to form true moral beliefs, even if this manifests a failure to care sufficiently about things of moral significance.

I think Wieland (2015) is right to reject (2); pure moral ignorance does not always manifest a failure to care about things of moral significance. However, I do not think the example he offers to motivate this objection is convincing. Instead, I argue that it is plausible that agents who are morally ignorant due to justifiable confusion or ambivalence about what actions they ought to perform, do not thereby manifest a failure to care

sufficiently about things of moral significance. Additionally, psychological conditions (e.g. self-deception) can interfere with an agent forming a true moral belief about things for which she cares sufficiently. In light of this, I argue that at best, it is the case that moral ignorance very often manifests a failure to care sufficiently about things of moral significance.

I argue that objections to (1) that focus on the epistemic situation of the agent are largely unsuccessful. Although some agents do hold justified false moral beliefs, it is not obvious that they are non-culpable for doing so. And, I argue it is implausible that agents ever lack access to evidence that would enable them to form true moral beliefs. However, (1) is likely mistaken for a different reason. Sometimes agents cannot be reasonably expected to care sufficiently about things of moral significance, and thereby to hold moral beliefs which manifest sufficient care. When this is true, it is not plausible that agents are culpable for their pure moral ignorance, even if it does manifest a failure to care sufficiently about things of moral significance.

Objections to (2)

Wieland (2015) argues that Harman is mistaken in claiming that pure moral ignorance always manifests a failure to care sufficiently about something of moral significance. To motivate this claim, he offers the following example of an Ancient Greek slave holder.

*Case 3*** Cleo keeps two slaves and forces them to work for her without pay. Cleo is ignorant that this is wrong because she is ignorant that slavery is wrong. She has made a serious attempt to determine whether slavery is wrong, and collected all the non-moral facts about the issue (she knows that they suffer, that she could have been a slave herself if she were unlucky enough,

etc.). It is not the case that Cleo keeps slaves because she wants them to suffer. Nor is she indifferent to their suffering: she is aware of it and feels sympathy with the slaves. Still, she did not succeed in drawing the inference that slavery is wrong because of the limited social context, and concluded that it was permissible (Wieland, 2015, p. 8-9).

Wieland's example is not fine-grained enough to be a counter-example to Harman's view. In order for this example to be a counter-example to Harman's view, Cleo's failure to draw the conclusion that slavery is wrong must not be based on non-culpable factual ignorance. E.g. it must not be the case that she non-culpably believes that given the fact that slaves are barbarians, they are not fit for anything other than servitude, or that the suffering slaves experience in virtue of being slaves is relatively minor—e.g. that it is comparable to the suffering that anyone who works in a service sector might experience.³⁰ Additionally, in order to be a counter-example to Harman's view, Wieland's example must be fleshed out so that it is plausible that Cleo cares sufficiently about the suffering of her slaves. E.g. it cannot be the case that she thinks to herself that slaves are just as fit for self-determination as any free person, that slavery imposes serious suffering on slaves, but that it promotes the prosperity of Athenian Elite, and frees her from the burdens of housework, so it is justified.³¹ Were this the case, then her failure to conclude that slavery is wrong would

³⁰ It is plausible that some Greeks may have been non-culpable in holding this belief. Vlassopoulos (2010) points out that a large portion of the Ancient slave population: engaged in professional activities that took them out of the household, whether they were working with their masters or on their own. We know of slaves who worked together with their masters as potters or builders and others who worked on their own as bankers, perfume makers or shoemakers. These slaves participated in joint activities together with other free and slave persons (11).

³¹ I am assuming here that Cleo would wholehearted endorse her conclusion. This line of reasoning would not be outlandish for an Ancient Athenian. Finely (1998) and Scheidel (2005) hypothesize that economic reasons were the driving force behind slavery in Athens, and in particular that labor shortage and demand motivated the practice. Finely (1998)

suggest that she believes that the economic ease and prosperity of the Athenian Elite justifies the imposition of serious suffering upon slaves. It would then be implausible that she cares sufficiently about slaves' suffering.

Below is a rendering of Wieland's example in which Cleo's moral ignorance is not rooted in non-culpable factual ignorance, and in which it is at least initially plausible that she cares sufficiently about the suffering of slaves.

Cleo4: Cleo knows all the relevant non-moral facts about slavery. She knows that slaves are generally captured in war, and are forced to leave their homes and perform free labor. She believes that slaves are just as fit to live a life of self-determination as any free individuals, and while she believes that slavery contributes to the economic prosperity of Athens, she does not believe that it is required for it. She believes that slavery causes slaves to suffer, and she deeply sympathizes with their suffering. She wishes for their sakes that slavery did not exist. She has made a serious attempt to determine if slavery is permissible. While she believes it is morally wrong to harm slaves in ways other than forcing them to perform free labor (e.g. beating them, berating them, denying them food or water, and so on), she ultimately agrees with the prevailing view. As long as some people have power, and others lack it, the powerless will be exploited by the powerful. Being a slave is a result of bad luck, and is unfortunate for slaves, but the practice of slavery is not unjust or otherwise wrongful.

Cleo4's moral ignorance concerns whether or not it is permissible to kidnap, forcibly subjugate a group of individuals, and exploit them for their labor (in foreign lands), much to their own suffering. When put in these terms, it is difficult to accept the characterization of

argues that the abolition of debt-bondage in the early 6th century resulted in slavery being the sole source of exploitable labor (157-158). Scheidel (2005) argues that as Athenian citizens gained more civic and military responsibilities,

the combination of high time-commitments of fairly closed citizen-populations, commercial and/or expansionist opportunities, maritime access to slave markets in Asia Minor, the northern Aegean (Thrace) and the Black Sea region, and the lack of readily exploitable neighbors may well have been instrumental in precipitating the intrusion of chattel slavery into all sectors of the economy (11).

In *Politics* 1253b24–54a13, Aristotle emphasizes the importance of slavery in making leisure possible.

Cleo4 as caring sufficiently about the suffering of slaves while failing to draw the conclusion that slavery is wrong. This is particularly true given Wieland's stipulation that Cleo has thought seriously about the permissibility of slavery.

A sympathetic reading of Wieland may be that he believes that our cultural and historical settings can limit our moral imaginations in ways that cause us to form false moral beliefs (or to fail to form true ones), even when we care sufficiently about things of moral significance. But, I am not convinced that this is a plausible explanation of Cleo4's *moral* ignorance.³² The most plausible reason that Cleo4's imagination with respect to slavery

³² Whether or not cultural and historical contexts ever limit our moral imaginations is an empirical question, and not one I will explore at length here. This is mainly because I think that if cultural and historical contexts sometimes limit agents' moral imaginations in a way that renders them non-culpable for their moral ignorance, this is likely because they limit an agent's access to factual knowledge. If this is right, then moral ignorance that is the result of one's moral imagination being limited by her cultural and historical contexts will be rooted in non-culpable factual ignorance, and will therefore not pose a problem for Harman's account. E.g. given our historical context, we may be limited in our ability to imagine the possible future uses and effects of nanotechnology, and therefore may lack true moral beliefs concerning it. But, in this instance, our moral ignorance would be based in non-culpable factual ignorance and would therefore be non-culpable even on Harman's view. Alternatively, sometimes our position within a culture limits our moral imaginations in other ways that render us non-culpably factually ignorant, and this factual ignorance sometimes results in non-culpable moral ignorance. E.g. perhaps individuals of the middle, upper middle, and upper economic classes in the United States must experience the stress and insecurity of being poor in order to fully appreciate the urgency of addressing economic inequality in the U.S. However, individuals in higher economic classes in the U.S. are shielded from the realities of poverty in the U.S. due to their economic privilege. So, one might think that the evidence these individuals need in order to form true moral beliefs about the urgency of addressing poverty is inaccessible to them. If such agents are non-culpable for their moral ignorance, this is because they are non-culpable for their factual ignorance. Again, this is consistent with Harman's view. Calhoun (1989) argues that there are other types of situations in which our cultural or historical contexts may render agents non-culpably morally ignorant due to the fact that they render agents non-culpably factually ignorant. She focuses on cases in which members of disenfranchised groups in a society (e.g. women, people of color) make significant moral advances (based on factual revelations—e.g. that science has a male bias) that take time to reach the broader population.

might be limited by her cultural and historical context, is explained by Vlassopoulos (2010). He argues that, despite their impressive imaginations, many Ancient Greeks could not imagine a world without slavery (2). He argues that this was because

Douleia [a main Greek term for the relationship between slaves and masters] for them was not a relationship of property which could be abolished by legislative *fiat*; rather it defined a situation in which an individual or a community was under the power of another individual or community. *Douleia* is the pragmatic result of the fact that there exists inequality of power and wealth among individuals and communities. Some people are *douloi*, because there are others who have the wealth and power to force people to execute their orders or to afford not to do things on their own but to have other people do them on their behalf. The Greeks understood clearly that as long as there were people who, because they had more wealth and power, were able to make others to obey their orders, *douleia* could not be extinguished (Vlassopoulos, 2010, p. 6).

However, even if we stipulate that Cleo4 finds it unimaginable that slavery be abolished, this doesn't explain her moral ignorance. One need not be able to imagine the abolition of a practice in order to judge the practice to be wrong. And, it is still very difficult to accept that Cleo4 cares sufficiently about the suffering of slaves if, despite thinking very hard about the matter, she concludes it is permissible to kidnap them and force them to perform free labor in a foreign land.

I'm not convinced that Wieland can succeed in offering an example in which an agent cares sufficiently about the suffering of slaves, has all the non-moral facts straight, has thought hard about the morality of the practice of slavery, and has still failed to form a true moral belief about it. However, regardless of whether or not this true, his general point that an agent can care sufficiently about something of moral significance while remaining (purely) morally ignorant with respect to it, is quite plausible. This may happen when an agent is justifiably confused about what she ought to do in a given scenario, or when an

agent has psychological conditions (e.g. self-deception, fear) that interfere with her forming true moral beliefs about things for which she cares sufficiently. Consider the following examples.

Confused Environmentalist: Sarah is deeply concerned about the environment. She donates regularly to conservation groups, stays informed about global warming, and shares her knowledge and concern about the environment with others. She also believes that if every human being consumed as many natural resources as she does, this would be unsustainable. However, she is unsure whether or not she ought to substantially alter her lifestyle. This would require significant, on-going sacrifices on her part, and would have negligible effects, given the rate of consumption of other human beings. Let's stipulate that it's true that Sarah ought to change her lifestyle. However, she fails to form this belief. Sarah's failure to form an accurate moral belief is due to genuine confusion about her moral obligations, and not to self-interest. I.e. she does not resist this conclusion because she does not want to have to substantially change her lifestyle; she would gladly do so if she believed it was the right thing to do.

Cleo5: Cleo thinks to herself that slaves are just as fit for self-determination as any free person, and that slavery imposes serious suffering on slaves. It also promotes the prosperity of the Athenian Elite, and frees her from the burdens of housework. She avoids thinking very hard about the permissibility of slavery (the truth about her actions would be too much to bear). However, she feels extreme guilt for owning slaves (though she would not describe her feelings in these terms), and has recurring nightmares about their experiences, is depressed due to keeping slaves, keeps a diary about the suffering she imposes on her slaves and the ways they have been wronged, acts so as to ease their burden whenever she can, dreams of one day releasing them from slavery, and so on.

Sarah fails to form a true moral belief that she ought to substantially change her lifestyle, but it is not plausible (given my description of the case) that this failure manifests a failure to care sufficiently about the environment or her duty to protect it. She isn't sure whether or not she ought to make large sacrifices that have little impact. I take it that this is a genuinely difficult moral issue—i.e. one about which professional ethicists might be ambivalent. That Sarah is wrestling with this issue, and that she considers making significant, on-going sacrifices for the good of the environment, suggests how seriously she takes her responsibility to do her part, and how much she does in fact care about the

environment, despite failing to form true moral beliefs with respect to it. It is plausible that many morally ignorant agents fit this description—they may care sufficiently about things of moral significance (as revealed by other aspects of their psychology, behavior, and actions), but fail to form true moral beliefs because it is very difficult to do so insofar as this would require them to form true moral beliefs about genuinely difficult moral issues.

It is plausible that Cleo5 cares sufficiently about the suffering of slaves, despite failing to form true beliefs about the permissibility of slavery. The psychological strife Cleo5 experiences due to the suffering she imposes on slaves, her attempts to ameliorate their suffering, and her dreams of freeing them, all manifest her care for the suffering of her slaves.³³ However, some psychological conditions prevent her from reaching true moral beliefs about their suffering. E.g. she might unconsciously avoid reaching the conclusion that slavery is wrong because doing so would result in greater guilt on her part, result in her thinking ill of others whom she holds in high esteem (and who think slavery is unproblematic), or require her to act in ways that will have serious negative consequences for her.³⁴

If I am right about these cases, then Harman's claim that pure moral ignorance always manifests a failure to care sufficiently about things of moral significance ought to be significantly curtailed. Many cases of moral ignorance are likely to be ones in which the agent fails to form true moral beliefs despite caring sufficiently about things of moral

³³ Of course, one might insist that Cleo *does* believe that slavery is wrong and simply does not admit this to herself. We may have good evidence for believing this to be the case. If this were the case, Cleo5 poses no problem for Harman's view since Cleo5 would not be morally ignorant. But, Harman would owe some explanation concerning how we can determine what someone's *real* belief is, if not by her own testimony.

³⁴ Note that I am not here claiming that Cleo5 is non-culpable for her moral ignorance, but only that it is plausible that she cares sufficiently about the suffering of slaves, despite holding false moral beliefs about the permissibility of slavery.

significance (as revealed by her other actions, behaviors, emotional reactions, and so on), either because she is grappling with genuinely difficult moral issues, or because psychological conditions interfere with her forming true moral beliefs about the objects in question. In light of such cases, it is more plausible that moral ignorance very often manifests a failure to care about things of moral significance, than that it always does so.

Harman would likely resist my analysis of these cases. (A) She might resist my analysis of both cases by arguing that an agent cares sufficiently about things of moral significance only if she forms true moral beliefs with respect to them. (B) She might reject my analysis of Cleo5 by resisting the claim that psychological conditions can interfere with an agent forming true moral beliefs about something for which he cares sufficiently. Neither response is convincing.

(A)

Harman may resist my analysis of Cleo5 and Sarah (and agents like them) by arguing that if they cared sufficiently about the environment and the suffering of slaves, they would hold true moral beliefs with respect to these things. This response is suggested by Harman's discussion of Huck Finn.

Huck believes that he is morally required to turn Jim in, but at a crucial moment when he could easily do so, Huck does not. He "resolves to be bad" instead. One version of the case is this: Huck does genuinely believe that it is morally required to turn Jim in, but despite this he is moved by Jim's humanity, and this is why he refrains from turning Jim in. It has been asked: is Huck praiseworthy for acting? I have not offered an account of praiseworthiness, but the following sits nicely with the view I have proposed:

An agent is praiseworthy for a morally good action just in case the agent's action resulted from caring about the features of the situation that make the action a morally good action.

On this view, Huck *is* praiseworthy for refraining from turning Jim in. This is a conclusion that I can happily embrace.

But on my view, Huck is also blameworthy for something: his moral belief that he should turn Jim in. Huck's psychology, in my view, involves his both caring about Jim's humanity – it moves him to refrain from turning Jim in – yet also not adequately caring about Jim's humanity – it does not move him to *believe* that Jim deserves to not be a slave; it does not prevent his false moral belief. There are two ways we could understand Huck's psychology, in both of which he is somewhat blameworthy on my view. One possibility is that Huck does care about Jim's humanity, but not *fully*. The other possibility is that Huck simultaneously has two conflicting attitudes, two conflicting levels of care toward Jim's humanity: he cares about it fully, but he also cares very little about it. (Compare the way that a person might have two conflicting beliefs: I believe I will be off campus on Tuesday, and I also believe I'll have lunch in the cafeteria with Adam on Tuesday) (Harman 2014, p. 14-15).

Harman's view, as revealed in this analysis of Huck Finn, seems to be that an agent can be said to care sufficiently about something of moral significance only if she is moved to form a correct moral belief with respect to it. If this is her position, it would allow her to reject the claim that Sarah and Cleo5 care sufficiently about the environment or the suffering of slaves, despite failing to hold true moral beliefs about them. However, the claim that agents care sufficiently about things of moral significance only if they are moved to form true moral beliefs about it seems to me to be a mistake.³⁵

Whether or not we ought to accept this principle depends on the conception of care that we adopt. Harman does not offer an account of care, and I do not intend to do so here. However, I think it would at least be very challenging for Harman to offer an account of care, according to which an agent cares sufficiently about something of moral significance

³⁵ Wieland (2015) also raises this concern, but does not defend it (8-9). If Harman does not mean to suggest here that, as a general principle, an agent cares sufficiently about something of moral significance only if she is moved to form a true moral belief with respect to it, then it is not clear how she could avoid describing agents like Sarah and Cleo5 as caring sufficiently despite being morally ignorant. So, whether she holds this to be a general principle or not, my point stands that Harman is likely mistaken in holding that all pure moral ignorance manifests a failure to care sufficiently about things of moral significance.

only if she forms a true moral belief with respect to it. Any plausible account of what it is to care about something must grant that caring involves having or being disposed to have certain feelings, desires, beliefs, and patterns of deliberation and attentiveness with respect to the object of care. For example, because I care about my daughter, I am disposed to feel distress when she is hurt, and joy when she succeeds or is happy. I am disposed to desire her wellbeing and to be near her. I believe that I ought to sometimes sacrifice my own interests for her sake. I often take her interests to be reasons for acting, and I attend to her general behavior and well being (e.g. does her behavior suggest that she is happy at school? Is she learning at music class? Is she adjusting well to a friend moving away?). Forming certain beliefs are only *one* element of the complex phenomenon of caring. It is plausible that if the other elements are in place, but the belief is missing with a good explanation (e.g. the agent has strong unconscious motivation for not holding the belief), an agent can be said to care about the object in question.

Perhaps caring about “things of moral significance”, e.g. the suffering of sentient beings, dignity, autonomy, rights, and so on, is more like caring about an end in which one is not personally invested, than it is like caring about individual persons. However, I am not convinced that these are distinct phenomena, or that they are sufficiently distinct to ground the claim that an agent cares sufficiently about something of moral significance only if she forms a true moral belief with respect to it.³⁶ E.g., I care about freedom of speech. My care

³⁶ Shoemaker (2003) makes a similar distinction between “caring for” and “caring about” (94). He points out that some might think that:

caring for certain others, as many feminist thinkers have shown, and as my dog example reveals, is a rather thick concept, involving as it does a personal concern for the object of care for its own sake, whereas caring about things (like peace on earth, say) is a more general, or much thinner, concept and does not seem to involve such personal

for it bears much resemblance to the care I have for my daughter, though the emotions involved in my care for freedom of speech are comparatively very mild. I am disposed to feel outraged when this freedom is violated. I am disposed, when it is in my power, to take some actions to prevent such violations and to restore freedom of speech when it has been denied. I am disposed to take the fact that something would violate someone's freedom of speech as a reason for acting (or not). Because I care about freedom of speech, I pay attention to things I otherwise would not (e.g. presidential candidates whose remarks suggest they would favor legislation that limits freedom of religion). Again, the phenomenon of caring is complex, and belief is only one aspect of it. It is plausible that one could care about an end in which she is not personally invested, even if she lacks true moral beliefs with respect to it (with good explanation).

Another possibility is that by "caring sufficiently about things of moral significance" Harman has in mind that we fully appreciate the moral value of the thing in question. But, this cannot be right. To fully appreciate a thing's value is primarily to correctly judge its value.³⁷ Understanding pure moral ignorance to be culpable insofar as it manifests a failure to make a correct judgment with respect to a thing's moral value would be circular.

(B)

concern (94).

Like me, he concludes that a distinction along these lines is "overdrawn" since both types of caring involve a certain investment on the part of the agent who cares (95).

³⁷ We might also aim to protect and preserve something if we fully appreciate its value. However, if we do not also correctly judge its value, these actions suggest that the thing is valuable to us, but not that we fully appreciate its value.

Alternatively, Harman might only reject my analysis of Cleo5 by arguing that psychological conditions cannot interfere with agents forming true moral beliefs with respect to things for which they care sufficiently. However, it would be somewhat puzzling for Harman to make such a claim. She grants that psychological conditions such as depression may interfere with an agent's actions manifesting what she really cares about (Harman, 2014, p. 19). It is not clear why she would not also grant that this can happen with respect to beliefs. Moreover, there is good reason to believe that psychological conditions do interfere with agents forming true moral beliefs about things for which they care sufficiently. Psychological conditions such as anxiety and trauma prevent agents from forming true factual beliefs.³⁸ There is good reason to think that psychological conditions can also interfere with an agent forming true moral beliefs.³⁹ If psychological conditions can interfere with an agent forming true moral beliefs, at least some of the time, this may cause her to form false beliefs, or to avoid forming true moral beliefs, that do not manifest what she really cares about.

There are numerous first-hand accounts of this happening. For example, Ishmael Beah (2007) describes how, shortly after he and fellow child soldiers were rescued in Sierra Leon, and detoxed from drugs, they “now had time to think; the fastened mantle of our war memories slowly began to open. We resorted to more violence to avoid summoning

³⁸ E.g. according to James and Gilliland (2012), soldiers with PTSD often form false factual beliefs such as that they are invincible, or that they are in ever-present danger (at home) (159).

³⁹ Tsang (2002) explains that in one form of this phenomenon—“moral rationalization”—agents distort the morality of an action they have committed or are about to commit in order to “perceive that their actions are consistent with their valued moral standards” (11). In a highly influential article, Skykes and Matza (1957) describe ways in which some criminals similarly distort the moral nature of their crimes in order to be able to commit them and live with having done so.

thoughts of our recent lives.” Beah (2007) also explains that he and his fellow child soldiers were angered by relief workers’ assurances that none of it was their fault. They did not view themselves as helpless victims, but as ferocious warriors who had done what they had to do. Reflecting on his reckoning with his silence as Jews and ethnic minorities in Danzig were deported, Grass (2007) describes the difficulty he faced in arriving at an objective moral assessment of his own actions:

...because many kept silent, the temptation is great to discount one’s own silence, or to compensate for it by invoking the general guilt, or to speak about oneself all but abstractly, in the third person: he was, saw, had, said, he kept silent...and what’s more, silent within, where there is plenty of room for hide and seek.

As soon as I summon up the boy I was at the age of thirteen, subject him to the third degree, and feel tempted to judge him as I would a stranger to whose needs I am indifferent, to condemn him, I see a kid of average height in shorts and kneesocks constantly grimacing, running to his mother and crying, “I was just a child, just a kid...” (28-29).

Even as Grass (2007) confesses his complicity with the Nazi regime, he reports his struggle with coming to terms with his actions:

They want accurate information about other things, about what else has been encapsulated [in his novels], about what has been swallowed in shame, about secrets in varying disguise, about nits nesting in sackhair. Eloquently avoided words. Slivers of thought. Things that hurt. Even now...(62).

Both Beah and Grass, at one point in time, avoided forming true moral beliefs about their past actions.⁴⁰ However, it is implausible that even at the time that they did so, this

⁴⁰ Of course, neither Beah nor Grass is fully culpable for his actions. Beah is arguably not at all culpable. He was a desperate child, alone in a war zone. He was brainwashed, drugged, and traumatized. Grass too was a child when he joined Hitler’s Youth (and as a testament to his youthful mindset) describes himself as being seduced by the glamour of being part of the movement. However, it is also not the case that their actions were morally

manifested of a lack of care about things of moral significance on their part. Their words do not reveal that they failed to care sufficiently, but that they cared so deeply that they had difficulty reckoning with the moral reality of their actions. Their future actions testify to this care. Beah has written two novels about his experience as a child soldier, and is a human rights activist, fighting to end the exploitation of children in violent conflicts. Grass has been called “the conscience of Germany” for a literary and political career dedicated to calling out the German public and politicians for their collective forgetfulness of the atrocities committed the Nazi regime, and their compliance with it.⁴¹

Still, one might argue that agents like Cleo5, Grass, and Beah are not accurately described as caring sufficiently about things of moral significance while failing to hold true moral beliefs with respect to them. Instead, one might argue that they have true moral beliefs, but intentionally fail to acknowledge or focus on them. Or, one might grant that they were unintentionally self-deceived, but deny that this entails that they lacked true moral beliefs about their actions at some point, if one understands self-deception to be a situation in which one believes both P and not-P.⁴²

I grant that these are possibilities. It may very well be that I have not accurately characterized the psychology of Grass and Beah, or that my analysis of Cleo5 is implausible. It may also be the case that some self-deceived agents do not lack true moral beliefs, but instead bury or ignore them. However, psychologists grant and observe cases of self-deception in which agents do not believe P and not-P. von Hippel and Trivers (2001) explain that in some forms of self-deception,

unproblematic. I take it that both, initially, resisted the truth about the moral gravity of their actions.

⁴¹ See Buruma (2006).

⁴² Lynch (2016) reports that this understanding of self-deception has waned (513-514).

the individual need not have two representations of reality to self-deceive. Rather, people can self-deceive in the same way that they deceive others, by avoiding critical information and thereby not telling (themselves) the whole truth (2).

While Beah, Grass, and Cleo5 may not meet this description, their cases hopefully motivate the claim that human psychology (and life) is sufficiently complex that some moral agents likely meet this description, and as a result, fail to form true moral beliefs about things for which they care sufficiently.

Through the above arguments and examples, I have tried to motivate the claim that pure moral ignorance does not always manifest a failure to care sufficiently about things of moral significance, and to argue against what I take to be Harman's potential responses to this move: (A) forming a true moral belief about something is a necessary condition for caring sufficiently about it, and (B) psychological conditions cannot interfere with one forming a true moral belief with respect to something for which she cares sufficiently. If what I have said is convincing, I think it is a best the case that pure moral ignorance often manifests a failure to care sufficiently about things of moral significance.

However, even if Harman weakens her claim that pure moral ignorance always manifests a failure to care sufficiently about things of moral significance, if she is right that when it does so, it is always culpable, her view still amounts to a very strong position concerning culpability for moral ignorance. I.e. it still amounts to the view that pure moral ignorance very often culpable, simply in virtue of the fact that it manifests insufficient care. There are a number of possible objections to this more constrained view as well.

Objections to (1)

As noted above, objections to (1) the claim that moral ignorance is always culpable when it manifests a failure to care about things of moral significance, mainly focus on agents' epistemic situations. As Harman (2014) points out, some might reject (1) on the grounds that sometimes agents are justified in holding false moral beliefs, and that agents are not culpable for holding justified false moral beliefs, even when these beliefs manifest a failure to care sufficiently about things of moral significance.⁴³ Alternatively, Wieland (2015) argues that some agents are not in an epistemic position to form true moral beliefs due to the fact that they lack evidence for doing so. When this is true, one might think that agents are not culpable for their resulting moral ignorance. I argue that while it is plausible that some agents are justified in holding false moral beliefs, Harman's position that they are nevertheless culpable for doing so is not as unintuitive as it may at first seem. I argue against Wieland's (2015) claim that some agents are not in an epistemic position to hold justified true moral beliefs, though it may be the case that, given their epistemic position, some agents cannot be reasonably expected to do so. Finally, I argue that there is another, non-epistemically based reason for doubting the claim that moral ignorance is always culpable when it manifests a failure to care about things of moral significance; sometimes it is not reasonable to expect agents to care sufficiently about things of moral significance.

⁴³ Harman also considers the objection that some agents form their moral ignorance as a result of bad luck—they are unfortunate to suffer poor formative circumstances or constitution. Some may think that agents cannot be culpable for something that is merely a result of luck. Harman (2014) dismisses this worry, explaining that she believes we can be and are culpable for things that are the result of bad luck (16). Because the existence of moral luck is controversial, it is unproductive basis for a culpability view, or an objection to one. For this reason, I do not discuss this objection above.

Justified False Moral Beliefs

Harman offers a number of responses to the potential objection that some agents are epistemically justified in holding false moral beliefs, and are thereby non-culpable for these beliefs, even if they manifest a failure to care sufficiently about things of moral significance. First, she holds that it is plausible that moral ignorance is never epistemically justified insofar as experience provides agents with ample evidence for accurate moral beliefs, and against false ones (Harman, 2011, p. 461; Harman, 2014, p. 22). Second, Harman (2014) argues that even if agents are epistemically justified in holding false moral beliefs, they are also in a position to have epistemically justified, true moral beliefs (21). Finally, Harman (2014) suggests that it is possible for epistemically justified, false beliefs to be blameworthy (21). She points out that it is plausible, e.g., that we owe it to our friends to give them the benefit of the doubt, even when evidence supports thinking ill of them (Harman, 2014, P. 21). If we instead (justifiably) believe ill of them, we would arguably be culpable for doing so.

I'll consider possible objections to Harman's first and third responses first. I take Harman's claim that false moral beliefs are never epistemically justified to be the weakest of her three responses to the justified false belief objection. Consider the following example.⁴⁴

Debating Doctors:

Two physicians are treating a 14 year old girl for stomach pain. She has internal bleeding, and is also pregnant. The girl has not told her mother about her pregnancy, and requests that the doctors not tell her either. The doctors disagree about the best course of action. Both doctors believe that they will violate their patient's right to privacy if they disclose her pregnancy to her mother. Both doctors believe it is in the patient's best interests for her mother to know about the pregnancy; the child needs her mother's support, and the mother needs to know that her daughter is pregnant in order to make an informed decision about how to treat her daughter's internal

⁴⁴ This example is based on episode 17 ("I Wear the Face") of season 12 of *Grey's Anatomy*.

bleeding. Moreover, the mother will eventually learn about the pregnancy if they do not tell her. Both doctors also believe that the patient is foolish for not telling her mother about her pregnancy.⁴⁵

Doctor A believes that they ought to reveal the patient's pregnancy to her mother. Acting in the patient's best interests is more important than respecting her right to privacy. Doctor B believes that they should not reveal the patient's pregnancy to her mother; it is more important to respect her right to privacy than it is to act in her best interests.

The doctors agree on the non-moral facts of this case; their disagreement is purely moral.

It is not that the evidence (e.g. their past experiences, moral reasoning, judgments derived from empathy, and so on) justifies both doctors in holding their respective beliefs. One of the doctors is mistaken about the comparative value of respecting the patient's right to privacy and acting in her best interests. This doctor holds a justified false (purely) moral belief. Given this example, and many others we could think of, it is highly plausible that Harman is mistaken in holding that pure moral ignorance is never epistemically justified.

Harman allows for the possibility that she may be wrong, and that some false moral beliefs may be justified. She argues that they are nevertheless culpable. One might balk at this, insisting that an agent cannot be culpable for holding justified false beliefs. There are two possible responses to ease the counter-intuitiveness of Harman's suggestion. First, agents will often hold justified false beliefs concerning genuinely difficult moral issues. As I argued earlier, it is plausible that false beliefs about such issues do not manifest a failure to care about things of moral significance, in which case, on Harman's own view, they may be non-culpable. So, Harman may be able to grant that many justified false moral beliefs are non-culpable. Additionally, the claim that we can be culpable for justified false moral beliefs

⁴⁵ I am assuming there that there are no federal or state laws that either require the doctors to inform the mother, or prohibit them from doing so.

(within the context of Harman's view) is not as problematic as it may at first seem. Harman grants that culpability comes in degrees, and is proportionate to the lack of care manifested by an agent.⁴⁶ In many cases, an agent who holds justified false moral beliefs will not manifest a significant failure to care about things of moral significance, and will therefore only be slightly culpable for her false belief. Moreover, agents who hold justified false moral beliefs are not culpable for mismanaging their moral beliefs (because they have not done so). E.g. in the above case, the doctor who holds the false moral belief cares very much about her patient's wellbeing and her right to privacy, but misjudges their relative importance. If we grant that her false belief manifests a very slight deficient care for things of moral significance, Harman would hold that she is only very slightly blameworthy for her false belief, and not at all culpable for mismanaging her beliefs. In other cases, where an agent holds a more egregiously false justified moral belief, this will manifest a more serious failure to care about things of moral significance on her part. She will therefore be more culpable on Harman's view. But, this does not strike me as being as counter-intuitive as the claim that agents can be culpable for holding justified false moral beliefs initially seems.

Accessibility

Still, one might object to Harman's claim that agents are culpable for their pure moral ignorance, even when it constitutes a justified false moral belief, insofar as agents are always in a position to form justified true moral beliefs. Harman (2014) does not say much to elaborate this point, but does offer an example to illustrate it:

My college friend Moon and I studied math together. She

⁴⁶ See again Harman's (2014) discussion of Huck Finn, who she holds is "somewhat" blameworthy (20).

likes to talk to me about math, though it has been years since I studied it. She tells me that a certain mathematical claim is true; I believe her. In fact, that claim is false, and if I thought about it I might figure that out – I remember enough math to figure it out. In this case, I am epistemically justified in believing the claim though I am also capable of becoming epistemically justified in believing it is false (22).

Harman's thought seems to be that if we are in a position to hold justified true beliefs, and yet hold justified false beliefs, this manifests a lack of care about the object in question. In the example above, the fact that Harman does not access available evidence to form a justified true belief manifests a lack of care about the truth of the mathematical claim in question; she doesn't care enough about its truth to think seriously about it.⁴⁷ Although Harman would not be culpable, in this example, for failing to care sufficiently about the truth of the math claim, an agent would be culpable for manifesting a similar lack of care with respect to the truth of a moral claim.⁴⁸

There are a couple of possible objections one could raise to this suggestion. First, one might argue that agents are not always in an epistemic position to hold justified true moral beliefs. Alternatively, one could argue that the fact that they are in such a position

⁴⁷ Alternatively, one might think that if one is in a good position to hold justified true beliefs, but instead holds justified false beliefs, this is because she has mismanaged her beliefs. She is thereby culpable for the false beliefs she holds. I do not think that this is Harman's point since she rejects the notion that culpability for false moral beliefs depends on culpability for mismanaging them.

⁴⁸ Mason (2015) objects to Harman's claim that agents are culpable for holding justified false moral beliefs because they are in a position to hold justified true moral beliefs. She argues that this claim commits Harman to the more general, and highly implausible principle that the mere fact that an agent could have avoided a moral wrong implies that she is culpable for it (Mason, 2015, p.8). Instead, in order to be blameworthy for X, Mason (2015) holds, the agent's having X'ed must reflect something about the quality of the agent's will, and in particular, it must show a lack of moral concern (Mason, 2015, p. 8-9). However, I think this is a misinterpretation of Harman's claim. As noted above, I do not think that Harman's point is that agents are culpable for holding justified false moral beliefs *simply* because it was possible for them to avoid doing so. Instead, the fact that they hold justified false moral beliefs when they are in a position to hold justified true moral beliefs manifests a failure to care sufficiently about things of moral significance.

does not entail that they are thereby culpable for failing to hold true moral beliefs, insofar as, given their epistemic situation, some agents cannot be reasonably expected to do so.

Wieland (2015) argues that some agents are not in a position to hold epistemically justified true moral beliefs insofar as agents do not always have access to evidence that would enable them to form justified true moral beliefs.⁴⁹ He argues that there are three potential sources of evidence supporting true moral beliefs: testimony, mere reflection, and direct experience (Wieland, 2015, p. 9). It is not always the case that agents have access to evidence via these sources. Testimony does not always provide one with evidence that tells against the morality of a practice. E.g., it may not do so in cases in which an immoral practice is widely accepted (Wieland, 2015, p. 10). Some agents simply lack the capacity to conclude, through mere reflection, that a practice is wrong (10). Wieland (2015) holds, e.g. that not everyone has the capacity to see “by mere reflection that the suffering of slaves outweighs the benefits of slavery for a given society (such as Athens)” (11). This is because not everyone “is in a position to make such rather philosophical inferences” (Wieland, 2015, p. 11). Finally, not everyone will have access to direct experience that provides evidence for true moral beliefs (e.g. not all Athenians will have direct contact with slaves).

While I grant that some agents may lack access to evidence in favor of true moral belief via testimony and direct experience, I am not convinced that (fully-abled, adult) agents (of average intelligence) ever lack access to evidence in favor of true moral beliefs via mere reflection. Consider Wieland’s (2015) own example of an agent who purportedly lacks access to evidence supporting the true moral belief that slavery is impermissible.

⁴⁹ Wieland (2015) does not make the further claim that agents are non-culpable for holding false moral beliefs when they lack access to evidence that would support true moral beliefs.

Wieland (2015) suggests that this agent is unable to make the philosophical inference, via mere reflection, that “the suffering of slaves outweighs the benefits of slavery for a given society (such as Athens)” (11). Let’s think about what evidence one might access through mere reflection, in favor of this belief. One might reflect that it would be truly terrible to be a slave. She might think about how, were she a slave, the fact that she was contributing to the economic well being of Athens would do little to console her. More abstractly, she might consider that the fact that it’s generally unacceptable to harm others for one’s own economic gain (e.g. it’s wrong to steal, cheat, manipulate, and so on). Or, she might think about whether or not she wants to live under a government that is willing to exploit the few in order to promote the welfare of the many. Even more abstractly, she might consider the value of money and economic stability in comparison to the value of freedom from suffering and the exercise of autonomy. I see no reason to assume that fully-abled, adult agents of average intelligence could not engage in reflection of this sort, and thereby access evidence that supports true moral beliefs about the permissibility of slavery via mere reflection.

Of course, some people might be better at accessing evidence for moral beliefs via mere reflection than others. Moral reflection arguably involves capacities that some might have in greater abundance, and skills that are improved by practice. Some agents might simply be more empathetic or imaginative, and therefore be better at arriving at the right conclusions through moral reflection. Alternatively, some agents might be better at basic reasoning (e.g. recognizing X as a member of some larger set, Y) due to natural ability or practice. But, the mere fact that some agents are better at accessing evidence for true moral beliefs, and thereby in a better position to form justified true moral beliefs via mere reflection, does not suggest that some agents are not able to do so at all.

Still, one might argue that even if an agent could form epistemically justified true moral beliefs on the basis of evidence that is available to them, and even if their pure moral ignorance manifests a failure to care sufficiently about things of moral significance, they are nevertheless non-culpable for it insofar as they cannot be reasonably expected to form true moral beliefs, given their epistemic situation. This might seem to be true, e.g., if through no fault of one's own, one is not very good at accessing evidence in favor of true moral beliefs through mere reflection, and the agent has evidence for false moral beliefs in the form of testimony or direct experience. Consider the following fictionalized example.⁵⁰

Mr. S is from a rural pocket of Malawi. He falsely believes that women and girls are of less moral value than men and boys. His belief is similar to the rather widely held belief that non-human animals have less moral value in virtue of being members of a different species; Mr. S believes that women and girls have less moral value in virtue of being a different sex. This is the prevailing belief in Mr. S's community, where women and girls are widely accepted to be similar in value to livestock. Some men in Mr. S's community treat their wives and daughters with great love and respect, but Mr. S and other men in his community do not take this to be evidence that they are equals. Instead they understand these men to have a particularly strong emotional attachment to their wives and daughters, similar to the attachment that some individuals have to pets. Mr. S isn't very empathetic or imaginative, and therefore is not very good at arriving at moral truths via mere reflection.

One might argue that Mr. S could not be reasonably expected to avoid holding his false moral belief that women and girls are of less moral value than men and boys. Testimony confirms this belief. Direct experience arguably does not provide evidence for or

⁵⁰ This example is based on an actual agent, Mr. Simbeye, discussed by LaFrainere (2005). Mr. Simbeye offered his twelve year old daughter as a sexual and domestic servant to his much older neighbor in order to repay a debt. He was genuinely surprised to learn that this practice is illegal in Malawi, and considered by many outside his geographic area to be abusive. It is a common practice in his village, and in neighboring villages. In offering an explanation for his behavior, he explained that he did not know he was abusing his daughter, and also that he had underestimated her monetary worth.

against it. Just as there is not much evidence via direct experience that counts for or against a view that moral value depends on one's species membership, there is not much that counts for or against a view on which moral value depends on one's sex. If Mr. S is poor at arriving at moral truths via mere reflection, due to his diminished capacities, one might argue that Mr. S cannot be reasonably expected to avoid his false moral belief about the moral value of women and girls.

I think it is very doubtful that many agents cannot be reasonably expected to form true moral beliefs, given their epistemic situation, and that they are thereby non-culpable for holding false moral beliefs that manifest a failure to care sufficiently about things of moral significance. However, arguing for this claim requires an account of the conditions under which an agent can be reasonably expected to X. Therefore, an argument for this will have to wait until chapter 4.

Caring Sufficiently

There is a remaining objection to (1) the claim that pure moral ignorance is always culpable when it manifests a failure to care sufficiently about things of moral significance. If some agents cannot be reasonably expected to care sufficiently about things of moral significance, then it is plausible that they cannot be culpable for failing to do so, or for holding beliefs which manifest a failure to do so. I think it *is* quite plausible that some agents cannot be reasonably expected to care sufficiently about things of moral significance.

Some agents who have suffered traumatic events may not be reasonably expected to care sufficiently about things of moral significance. As a response to traumatic events, some agents engage in "emotional numbing"; they suppress the normal emotions they might feel

about actions or events. This can become habitual, in which case, it will sometimes manifest a failure to care sufficiently about things of moral significance. For example, in a recent study of delinquent youth, Kerig et. al. (2012) found that youth who experience a traumatic event perpetrated by someone they trusted (e.g. sexual abuse by a trusted adult) were more likely to engage in “habitual emotional numbing,” which contributed to “withdrawal of empathy and the cultivation of callousness toward others” (277). This is echoed in self-reports and reports of agents who have experience trauma. Consider the following self-report from a veteran, Billie Mac, cited in Gilliland and James (2013).

Looking back on it, I cannot believe how callous I have become. SOP [standard operating procedure] was “It don’t mean nothin’, screw it, drive on.” This would be right after a B-40 round had blown your buddy’s brains all over you. You had to put it behind you to survive. A guy fell off the construction site I was working on last fall and splattered himself all over the pavement. I sat on a steel beam about 30 feet above the guy and just kept eating my lunch. No big deal (160)!

Consider also the following report of a young woman, charged with assault, who has a history of trauma.

Sadhanna is a 22-year-old woman mandated to outpatient mental health and substance abuse treatment as the alternative to incarceration. She was arrested and charged with assault after arguing and fighting with another woman on the street. At intake, Sadhanna reported a 7-year history of alcohol abuse and one depressive episode at age 18. She was surprised that she got into a fight but admitted that she was drinking at the time of the incident. She also reported severe physical abuse at the hands of her mother’s boyfriend between ages 4 and 15. Of particular note to the intake worker was Sadhanna’s matter-of-fact way of presenting the abuse history. During the interview, she clearly indicated that she did not want to attend group therapy and hear other people talk about their feelings, saying, “I learned long ago not to wear emotions on my sleeve.”

Sadhanna reported dropping out of 10th grade, saying she never liked school. She didn't expect much from life. In Sadhanna's first weeks in treatment, she reported feeling disconnected from other group members and questioned the purpose of the group. When asked about her own history, she denied that she had any difficulties and did not understand why she was mandated to treatment. She further denied having feelings about her abuse and did not believe that it affected her life now. Group members often commented that she did not show much empathy and maintained a flat affect, even when group discussions were emotionally charged (Substance Abuse and Mental Health Services Administration, 2014, p. 62).

Two things are highly plausible about the above cases. (i) The agents in question fail to care sufficiently about things of moral significance. (ii) The agents are not culpable for this failure. Billie Mac fails to care sufficiently that another human being fell to his death before his eyes. Sadhanna fails to care sufficiently about the suffering experienced by the members of her therapy group. Both lack the types of beliefs that would be indicative of care. Billie Mac reports that the death of another human being was "no big deal." Sadhanna reports that she doesn't see the point of discussing the suffering of the members of her therapy group. Both also lack the affective and behavioral responses indicative of care. Billie Mac calmly eats his lunch above the splattered remains of his coworker. Sadhanna shows a lack of empathy and "flat affect." Billie Mac does nothing to help, nor does Sadhanna. Perhaps at one point, these individual would have cared about such things, and perhaps through therapy they will come to care about these things again. But it is implausible that, at the relevant point in time, they cared sufficiently about them.

It is equally implausible that they are culpable for failing to care sufficiently about these things, or for any pure moral ignorance that may manifest this failure. This is because it is highly plausible that neither Billie Mac nor Sadhanna could be reasonably expected to

avoid their lack of care. It is arguably the result of them habitually tuning out their emotional responses, of them learning to not care. This is an activity from which we arguably could not reasonably expect them to refrain. As Billie Mac states, engaging in it was a matter of survival.

If agents sometimes cannot be reasonably expected to care sufficiently about things of moral significance, and therefore are not culpable for failing to do so or for forming beliefs that manifest this lack of care, the Harman is mistaken in claiming that moral ignorance is always culpable when it manifests a failure to care sufficiently about things of moral significance.

3.4 Non-volitional Culpability Views

Non-volitional culpability views reach the same conclusion as Harman's view by a different route. They are based on non-volitional views of blame, according to which voluntary control is not necessary for culpability. Instead, proponents of non-volitional views of blame argue that the proper target of blame is our objectionable moral attitudes. According to some non-volitional views, blame consists of a judgment that an agent has or has displayed an objectionable moral attitude, and as is the case with any judgment, it is justified as long as this judgment is accurate and we are justified in reaching it, given our evidence.⁵¹ E.g. when we blame someone for driving recklessly, we judge him to have

⁵¹ For example, Scanlon (1998) argues that blame is a judgment that an agent has failed to govern herself according to principles which no one can reasonably reject (1998 p. 268-269). Blame is justified in case the aforementioned judgment is true. Hieronymi (2004)

displayed a lack of respect for the lives and safety of others. This judgment is justified if he in fact displayed a lack of respect and we have sufficient evidence for judging him to have done so. According to other non-volitional views of blame, when we blame someone for X, we implicitly demand that she justify the moral attitudes reflected by X, and if she cannot, we blame her for them.⁵² For example, when I blame Sam for cheating on an exam in my class, I demand that Sam justify this action. To do so, he must defend his evaluation of the reasons for and against cheating. Perhaps he can do so. Perhaps, e.g., he had good reason to believe his life depended on passing this exam, and rightfully judged his life to be more valuable than being honest in this particular scenario. If so, I ought to rescind my blame. But, if he cannot do so—if he knows he behaved badly, or he falsely judges that getting good grades is more valuable than being honest—my blame is justified.⁵³

argues that when we blame an agent, we simply acknowledge that the agent has exhibited an attitude of disrespect towards us or a fellow agent. Because blame simply consists of the observation that an agent has exhibited an attitude of disrespect, it will be justified just in case this observation is true and we have adequate reasons for believing it to be true (Hieronymi, 2004, p. 128-129).

⁵² This is Angela Smith's view, which she develops over a series of papers. See, e.g. Smith (2005), Smith (2008), Smith (2012), and Smith (2015).

⁵³ The description of non-volitional views of blame offered in this paragraph captures the most prominent non-volitional views, but does not capture all non-volitional views. E.g. Sher (2001) also argues for a non-volitional account of culpability. His view is that we're culpable for our character traits even though he believes that we often cannot be reasonably expected to avoid them, or reform them. He also does not rely on the fact that these traits were developed through our voluntary actions to support his view. He instead suggests that the connection between an agent's ability to avoid X (i.e. the agent's control over X) is not as intimately linked to warranted blame as we typically think. Robert Adams (1985) argues for a similar view as well, suggesting that we can be rightfully blamed for that over which we lack any type of control. His reasoning for this is two-fold. First, he holds that morality is not only action guiding, but also correctional and therefore moral criticism applies even when the thing for which we are criticized is something over which we lacked power. Second, he thinks that it is important that agents accept responsibility for things (mental and emotional states in particular), even though they lacked control over them, insofar as doing so is necessary to repair moral relations amongst agents after wrongdoing.

On either non-volitional view of blame, we are directly culpable for our objectionable moral attitudes. Our culpability for them need not be traced to culpability for past actions which led to them, or to a failure to perform actions that would have eliminated or mitigated them and their effects. In fact, how we came to have the objectionable attitudes in question is largely irrelevant to our culpability for them. Facts about the formation and maintenance of our attitudes do not affect whether or not our judgments that agents in question hold objectionable attitudes are true and justified, or whether or not their attitudes are defensible. E.g. the fact that Carol grew up in a very abusive household is not relevant to the accuracy or justification of my judgment that she is cruel. If blame simply amounts to such a judgment, and is justified if the judgment is true and we are justified in making it, then the fact that Carol grew up in a very abusive household is not relevant to whether or not she is blameworthy for being cruel and exhibiting her cruelty. Similarly, her background is not relevant to whether or not her cruelty is defensible. What if she defends this cruelty by asserting that the feelings of those she is harming are not important? If the parties in question (or their surrogate) were to object to this claim by asserting that their feelings are indeed important, it would not be appropriate for her to cite the abuse she endured in response.

Although proponents of non-volitional views do not directly address or focus on the issue of culpability for moral ignorance, their views are nevertheless a very important part of this discussion insofar as they entail some kind of culpability view with respect to moral ignorance. The latter is a fact that has been observed by a number of participants in the

moral ignorance debate.⁵⁴ Pure moral ignorance in the form of a false moral belief, almost always manifests an objectionable moral attitude. And ignorance of a moral truth often does so as well⁵⁵ According to non-volitional views of blame, then, moral ignorance is almost always culpable. Moreover, the conditions under which an agent formed her moral ignorance are not relevant to her culpability for it. The objectionable nature of ignorance entails that an agent is culpable for it, and these conditions do not do not bear on whether or not the content of the moral ignorance is objectionable or defensible.

3.5 Objections to the Non-volitional Culpability View

One could object to the non-volitional culpability view by arguing that voluntary control is, in fact, required for culpability. This, of course, would be a large project that would take us far a field from the issue of culpability for moral ignorance. Fortunately we

⁵⁴ E.g. see Rosen (2003), pg. 73-74 (he refers to non-volitional views as “quality of will accounts”); Mason (2015) (who refers to non-volitional views as “attributionist” views); Levy (2009), pg 740; and Wieland (2015). Björnsson (forthcoming) brings non-volitional accounts directly into the discussion concerning moral ignorance by arguing that they imply that there are no epistemic conditions for culpability, and therefore that the skeptical argument rests on a mistake.

⁵⁵Imagine, e.g., that Bethany loves fur coats and owns many of them. She loves how soft and pretty they are. She knows that they are made from the fur of animals that would not be killed if there were no demand for fur coats. She is also aware that animals suffer in order for fur coats to be made. But, Bethany lacks beliefs about whether or not it is permissible to kill animals to make fur coats. It has never occurred to consider the matter because it has never occurred to her that the fact that animals must suffer and be killed in order for fur coats to be made, possibly counts against the permissibility of manufacturing fur coats. Bethany’s lack of a true moral belief manifests an objectionable moral attitude toward the value of the lives of non-human animals, and their suffering. This is only the case with respect to issues that directly intersect with our lives E.g. I might lack true moral beliefs about who is responsible for paying for the demolition of dangerous buildings in Nigeria insofar as I have never considered the matter. But, this does not reflect objectionable moral attitudes on my part insofar as this issue does not intersect with my life in any way.

need not take on this issue in order to argue that non-volitional culpability views are flawed. I argue that even if we grant that agents can be directly culpable for their moral ignorance insofar as it reflects an objectionable moral attitude, it does not follow that agents are therefore always culpable for their ignorance. I argue for this claim by first examining three existing objections to the non-volitional culpability view. (1) Some argue that moral ignorance does not always reflect an objectionable moral attitude. (2) Some argue against the claim that agents are always culpable for their objectionable moral attitudes by arguing that treating the objectionable nature of an attitude as sufficient for its blameworthiness, fails to recognize the difference between *being* bad and *deserving* blame. (3) Others argue against the claim that agents are always culpable for their objectionable moral attitudes by arguing that this view implausibly suggests that agents are culpable for their moral ignorance, even when they have done everything they could be reasonably expected to do in order to avoid it.

I argue that (1) is true only with respect to a restricted set of cases of moral ignorance, and therefore does not pose a significant problem for non-volitional culpability views. Both (2) and (3) are based on intuitive responses to particular cases, which proponents of the non-volitional culpability view are happy to reject. But, we need not accept that there is simply a clash of intuitions here. Instead, I argue that these objections can be supported by two further claims. (A) Agents are not culpable for that which they could not be reasonably expected to avoid, and (B) sometimes agents cannot be reasonably expected to avoid their objectionable moral attitudes, or the moral beliefs which manifest them, due to the circumstances under which they were shaped and carried into the present. If this is right, the non-volitional culpability view is mistaken. The mere fact that moral

ignorance constitutes or reflects an objectionable moral attitude does not itself render the ignorance culpable. The conditions under which moral ignorance is formed do affect the agent's culpability for it.

(1) Moral ignorance and objectionable moral attitudes

Wieland (2015) argues against the claim that moral ignorance almost always manifests objectionable attitudes, and in particular, rejects the view that ignorance almost always manifests ill will. He does so by considering a couple of ways of understanding ill will, and arguing that according to those understandings, it is not true that moral ignorance almost always reflects it. First, he suggests that moral ignorance may reflect ill will if it is the result of a failure to make a serious attempt to discover if X is morally wrong (Wieland, 2015, p. 6). However, on this interpretation of ill will, many instances of moral ignorance will not reflect it. Many agents think very hard about moral questions, and still arrive at very objectionable conclusions regarding them. Alternatively, moral ignorance about X may manifest ill will if it manifests, “an indifference for the fact that makes X wrong” (7). On this view, e.g., moral ignorance about slavery might manifest ill will if it manifests indifference for the fact that slaves suffer from being enslaved. However, Wieland (2015) argues that not all cases of moral ignorance will display ill will on this view, either. In many cases, agents are not indifferent towards the features that make X wrong, but still fail to draw the conclusion that X is wrong (recall, e.g., Cleo).

It is not clear that we should accept either of these characterizations of ill will. One might instead suggest that ill will amounts to a failure to adequately respect moral agents or sentient beings. Even on this understanding of ill will, it may be that there are a some of

cases in which pure moral ignorance does not manifest it. However there is no reason why proponents of the non-volitional culpability view could not accept this, and acknowledge that in such cases, agents are not culpable for their moral ignorance.⁵⁶ They would likely need to adopt this strategy, e.g., with respect to some cases in which agents make a sincere effort to come to the right conclusion about a genuinely difficult moral issue, but arrive at the wrong conclusion. E.g., recall Sarah the confused environmentalist, and the debating doctors. Sarah is uncertain if she ought to make significant sacrifices that will fail to have any significant impact. Just as this uncertainty does not plausibly manifest a failure to care sufficiently about the environment, it also does not plausibly manifest other objectionable

⁵⁶ Relatedly, one might worry about cases in which an agent's moral ignorance constitutes objectionable moral attitudes, but where these attitudes are only slightly offensive, and therefore the agent does not seem to be culpable for them. For example, many of us have friends who are chronically late. Some may feel very bad about this tendency, while some are shocked when others are annoyed by it. Let's stipulate that the latter hold that it is no problem to be 10-15 minutes late to a meeting, and thereby fail to adequately value other people's time. However, they do not grossly fail in this respect. They simply think that avoiding the stress of being rushed, or of better organizing their activities, is more important than the minor inconvenience being late causes those with whom they associate. Because this is such a minor failing, we might be inclined to simply tolerate this about them, and not blame them for it. One might wonder if proponents of the non-volitional view can satisfactorily accommodate cases like this. I think they can. It is open to them to argue that blame can come in degrees (reflected, perhaps, by the strength of the attitudes we adopt on account of our blame, or the degree to which we judge the agent to have shown a lack of regard), and that the appropriate degree of blame is proportional to the offensiveness of the attitude in question. So, they might argue, we *do* blame our chronically late friend, but to such a small degree that it registers as tolerance rather than blame. There may be other ways to handle cases like this. E.g. Holly Smith (2011) argues that an objectionable attitude can constitute a moral fault without being blameworthy if the attitude fails to capture an agent's overall evaluative assessment of a moral situation (e.g. of the comparative value of other people's time and their own interests) (140). This might apply here. It is possible that as long as the agent above shows that he cares about other people's time in other ways, his chronic tardiness does not reflect his full assessment of the value of other people's time, and is therefore a moral fault rather than something for which he is culpable.

moral attitudes (as I have described the case).⁵⁷ However, proponents of non-volitional views would be likely to accept that such agents are an exception to the rule that pure moral ignorance always manifests an objectionable moral attitude, and that they are therefore non-culpable (at least in this respect) for their moral ignorance. To suggest otherwise would be to insist we're blameworthy unless we hold perfectly accurate moral beliefs, even with respect to genuinely difficult moral issues. But, again, it does not seem problematic for proponents of non-volitional culpability views to grant that in some cases, pure moral ignorance does not manifest an objectionable moral attitude, and is (in this respect), therefore non-culpable.⁵⁸

Of much greater concern for non-volitional culpability views are objections to the claim that an agent is always culpable for her moral ignorance if it manifests an objectionable moral attitude. Levy (2005) and Rosen (2002) raise objections to the latter by pointing to cases in which X clearly reflects objectionable moral attitudes but the agent of X is not obviously culpable for X on account of this. Levy (2005) argues that some agents are bad, and not also blameworthy, in virtue of the fact that it is not their fault that they have their objectionable attitudes. Rosen (2002) argues that some agents are not culpable for their moral ignorance, even though it manifests objectionable moral attitudes, insofar as they have

⁵⁷ Again, it is possible to interpret Sarah's confusion as a form of affected ignorance (i.e. she does not conclude that she should make significant sacrifices because she does not want to make such sacrifices). However, as I am imagining her, she would happily make such sacrifices if she judged this to be the best course of action.

⁵⁸ Non-volitional culpability views will have to grant fewer exceptions to their view than Harman. This is because it is less plausible that many instances of pure moral ignorance fail to manifest an objectionable moral attitude than it is that many fail to manifest a failure to care sufficiently about things of moral significance. E.g. agents like Cleo5 and Grass are likely exceptions to the former, but not the latter. It is plausible that these agents care sufficiently about things of moral significance, but not that their failure to hold true moral beliefs does not manifest objectionable moral attitudes.

done everything they could have been reasonably expected to do in order to avoid their ignorance and/or the attitudes in question.

Levy's (2005) objection is not directed against non-volitional stances towards moral ignorance in particular, but to a non-volitional stance toward culpability for attitudes more generally. He argues that, contrary to non-volitional views, the fact that X reflects/constitutes an objectionable moral attitude does not imply that an agent is culpable for it, and that how an agent came to have his attitudes is relevant to his culpability for them. In rejecting the non-volitional stance, Levy argues that it is committed to saying that agents whose objectionable moral attitudes are the result of traumatic brain injury, or poor formative circumstances, are nevertheless culpable for them. This, he claims, is highly unintuitive. In support of this, he discusses the well known cases of Phineas Gage and Robert Alton Harris (Levy, 2005, p. 5-9). According to Levy (2005), before the accident in which a tamping iron punctured his skull and damaged his brain, Gage was "sober and hardworking" (8). After this accident, Gage was "dissolute," "anti-social" and had a diminished "ability to control his impulses" (Levy, 2005, p. 8). These changes were permanent.⁵⁹ Assuming that Gage's new attitudes (as Levy describes them) can be properly

⁵⁹ Macmillian (2000) argues that common accounts of Gage's injury and the psychological changes he suffered as a result depart largely from the only detailed first-hand accounts we have of Gage (from the physician who treated him and from Gage's mother). Of particular interest, Macmillian (2000) suggests that changes in Gage's personality may not have been permanent. According to Gage's physician and Gage's mother (as reported by Macmillian), once Gage regained his strength after his injury, he held was employed on a horse farm in Chile for eight years (and returned due to illness), and was reported to have a special fondness for, and gentleness with his nieces and nephews, and animals (66). However, it is not important for Levy's purpose that his description of Gage is accurate. It isn't even important that some agents do, in fact, suffer significant, permanent negative changes in character as a result of traumatic brain injury or illness. Levy could simply point out that *if* an agent did suffer such changes as the result of a traumatic injury or illness, non-volitionists would claim that they are culpable for their new objectionable attitudes. It is,

described as his own, proponents of non-volitional views will insist he is culpable for them; they constitute objectionable moral attitudes, and how he came to have them does not alter this fact.⁶⁰ However, this is highly implausible. Robert Alton Harris suffered a horrible, emotionally and physically abusive upbringing. As an adult, he unapologetically murdered two teenagers in cold blood, calmly ate their lunch, and then bragged about the murder to others. Levy (2005) suggests that it is plausible that “Harris may not have had a genuine chance to become a better person; nevertheless the attitudes he expresses in his awful actions are genuinely his” (5). Both Gage and Harris, according to Levy (2005), are not blameworthy for their objectionable attitudes because it is not their fault that they have them; they are simply bad in virtue of having them (5).

Levy’s (2005) argument is directly applicable to non-volitional culpability views about moral ignorance. Many of the cases in which we are inclined to excuse an agent for her moral ignorance are ones in which their moral ignorance seems to be the result of some

however, plausible that some agents do suffer permanent negative changes in their character as a result of traumatic injury and illness. E.g. Burns et. al. (2003) describe a patient who acquired pedophilia and a inability to suppress sexual urges as the result of an orbitofrontal brain tumor. The patient’s tumor was noticed after he was sent to the hospital (from a treatment center) complaining of a headache. After his tumor was resected, his symptoms resolved. However, had he not been treated, his condition would have been persisted.

⁶⁰ As noted above, Levy (2005) likely mischaracterizes Gage. However, if we grant that some agent, Gage*, could exist and meet the description offered by Levy, it is likely that proponents of the non-volitional view would grant that Gage*’s new attitudes, at least at some point in time, are properly described as his own. Smith (2005) argues that the notion of an attitude being an agent’s own should be “spelled out in terms of the very network of beliefs and attitudes which I am suggesting ground our attributions of responsibility. . . . A reasonable account of the conditions of responsibility should preserve our sense of the rational interrelations among our attitudes . . . (262). Assuming Gage*’s shift in attitudes was fairly global (i.e. most of his new attitudes reflected a lack of regard), it seems like this condition is met. Scanlon (1998) suggests in order for an attitude to be properly described as one’s own, it must bear “the right kind of stable and coherent connections between what one says, does, and how things seem to one at one time, and what one says, does, and how things seem to one at later times. . .” (278). At some point in time, this will arguably be true of Gage*’s new attitudes.

unfortunate circumstance(s). So, if Levy is right about Gage and Harris, this would suggest that in relevantly similar cases of moral ignorance, agents are not culpable for their moral ignorance even if it reflects/constitutes objectionable moral attitudes. Moreover, it would suggest that the conditions under which an agent formed her moral ignorance are relevant to her culpability for it.

Of proponents of non-volitional accounts of blame, Angela Smith has been the most outspoken with respect to cases like Gage and Harris, and more generally with respect to cases in which it seems that an agent is not culpable for his objectionable attitudes because it not his fault that he has those attitudes in the first place. She argues that in such cases, either the “agents” are not really agents at all, or they are, in fact, culpable for their attitudes. Smith (2005) argues that few of us are fully responsible for acquiring our particular values, but that this does not impact whether or not we are culpable for these values, for retaining them, and for displaying them (268-269).⁶¹ She cautions us against ever holding that agents are non-culpable for their attitudes; doing so suggests that they are the passive victim of their own judgments, which in turn suggests that they are not agents at all (Smith, 2008, P. 390-391). The latter is true only if the agent’s upbringing or circumstances have “damaged their very capacity to critically evaluate and respond to reasons” which in turn inform the attitudes they adopt (Smith, 2015, P. 125). Otherwise, the agent’s attitudes accurately reflect her own moral values, and if they are objectionable, she is culpable for them. Smith (2005, 2008)

⁶¹ Miller (2014) argues that Smith ought to accept a historical condition on culpability: “(HB’) An agent S is blameworthy for an attitude (or value) A only if it is not the case that (a) A was acquired in a way that bypassed S’s rational capacities completely and that (b) A is unsheddable” (Miller, 2014, P. 484). Such a condition would help Smith explain why agents like Gage* are not culpable for their attitudes, since his new attitudes likely met (a) and (b). It won’t help explain our reactions to cases like Harris’ which are more relevant to the topic of culpability for moral ignorance.

does allow that the conditions that contributed to an agent's attitudes can affect the degree to which we blame her (268, 390). If, e.g., an objectionable attitude has been strongly reinforced by an agent's culture, it might be very difficult for the agent to revise it, in which case we may blame her less harshly for it (Smith, 2005, p. 268).

Aside from this last response, none of Smith's comments are likely to be particularly convincing to those who, like Levy, find it to be very odd to claim that an agent might be culpable for her objectionable moral attitudes even if they are the result of a traumatic brain injury or abuse that no human being should have to endure. Smith's first two responses are reiterations of her view that agents are culpable for their objectionable moral attitudes (unless they are not agents at all). Levy rejects this. Her suggestion that how an agent came to have her attitudes (and the difficulty of shedding them) can lessen the degree of blame that is justified, might go some way towards lessening the counter-intuitiveness of holding that how an agent came to have an attitude is not relevant to his culpability for it. However, it is not clear how a non-volitional view of blame could support this claim. If blame constitutes a judgment that an agent has objectionable moral attitudes, or a demand that he defend those attitudes, it is not clear why facts about how he came to have those attitudes should affect the degree to which we blame him. Only facts about the degree to which an attitude is objectionable or defensible should do so.

So, we are faced with a clash of intuitions. Proponents of non-volitional views will insist that an agent is culpable for her objectionable moral attitudes (including moral

ignorance) regardless of how she came to have them, while Levy (2005) insists she is not.⁶² However, it is possible to offer a different defense of Levy's claims. If the conditions under which an agent forms his attitudes, or the conditions that explain why he persists in his moral ignorance, make it unreasonable to expect him to have avoided these deficiencies, it is plausible that he is not culpable for them. This is the objection Rosen raises to non-volitional culpability views.

Rosen (2003) argues that in some cases, agents are morally ignorant, and while their moral ignorance reflects or constitutes an objectionable moral attitude, they are nevertheless not culpable for these things insofar as they have done everything they could have been reasonably expected to do in order to avoid them. In support of this, he offers the example of an Ancient Hittite slave owner (Rosen, 2003, p. 65-66, 72-73). The slave owner buys and sells slaves, and separates families for his own material gain. He believes he is entitled to do so. Like others during this time period, he does not believe that the slaves are inferior. Instead, he thinks the fact that they are slaves is just bad luck. But though he acknowledges that slavery is a misfortune that could befall anyone, this does not lead him to question the warrant with which he buys and sells slaves, despite the misery he observes in the wake of these behaviors. Because slavery is socially accepted by all (or almost all) of those whose opinions he care about, he doesn't question the permissibility of the practice. Rosen

⁶² Levy (2005) does rely on the intuitive pull of his examples. However, he thinks that they point to the truth of a volitional view of blame, for which one could offer a more principled defense. As I noted above, I am not interested in taking up the issue of whether or not voluntary control is require for culpability. Moreover, I am not convinced that examples like that of Harris, at least, support volitional views. Our reluctance in blaming Harris is not due to the belief that no voluntary actions or omissions of his led him to hold his attitudes, but due to the belief that, given his past, he could not have been reasonably expected to avoid his attitudes.

(2003) argues that there is nothing the Hittite slave owner could have been reasonably expected to do in order to avoid his moral ignorance regarding slavery. Rosen (2003) argues that “one is normally under no obligation to rethink the uncontroversial normative principles that form the framework for social life,” but even if we stipulate that the Hittite man did reflect, “Given the intellectual and cultural resources available to a second millennium Hittite lord, it would have taken a moral genius to see through to the wrongness of chattel slavery” (65-66). So, Rosen (2003) argues that even if we grant that the Hittite’s moral ignorance reflects objectionable moral attitudes, due to what it would have taken for him to avoid this ignorance (i.e. that it would have required him to behave in ways we could not reasonably expect), he is not culpable for it (73).

Rosen’s claim that we are under no obligation to question uncontroversial moral principles, even if such principles endorse the intentional imposition of suffering upon one group of individuals for the benefit of another group, is certainly not uncontroversial. And, one might certainly reject Rosen’s description of the Hittite’s epistemic situation. As I argued previously, it is implausible that such agents lack access to evidence via mere reflection that would count in favor of the conclusion that the practice of slavery is morally wrong. However, Rosen’s claims may be more plausible with respect to other cases. They may be plausible, e.g. with respect to agents like Harris, who form deformed moral beliefs as a result of very unfortunate formative circumstances. Like Billie Mac and Sadhanna, discussed earlier in the chapter, it is possible that Harris (and others with similar backgrounds) responded to the unrelenting abuse he faced as a child and adolescent by “numbing” himself so that he did not experience the emotions one might typically experience when subject to such abuse. As discussed earlier, habitual numbing like this can

sometimes result in callous, uncaring attitudes towards others (which certainly describe Harris). Expecting him to avoid these attitudes, and thereby to avoid moral ignorance that manifests them, would require him to take action (i.e. to refrain from engaging in emotional numbing) that we could not reasonably expect him to take (since emotional numbing may have been the only resource Harris had to endure the abuse he suffered). It is therefore not clear that we could reasonably expect Harris, and others like him (e.g. Billie Mac, Sadhanna) to avoid his objectionable attitudes or resulting moral ignorance.

Because it is plausible that it is unfair to blame agents for that which they could not be reasonably expected to avoid, and it is plausible that some agents cannot be reasonably expected to avoid their moral ignorance, it is plausible that even if moral ignorance reflects or constitutes an objectionable moral attitude, agents are not always thereby culpable for it. It is also plausible that the conditions that explain why an agent is ignorant are indeed not irrelevant to her culpability for that ignorance.

Proponents of non-volitional views are not without resources to respond to this objection. Hieronymi (2004) argues that the claim that it is unfair to blame agents in such cases rests on the view that blame carries a particular force and that it is only fair to subject an agent to this force if her action/attitude meets certain historical conditions. This is a claim that she rejects. Many hold that the characteristic force of blame rests in the negative consequences it imposes upon the wrongdoer, and that it is fair to impose negative consequences upon an agent on account of X only if the agent had a reasonable opportunity

to avoid X.⁶³ However, Hieronymi argues that the characteristic force of blame does not lie in any negative consequences we impose as a result of blaming an agent, but instead lies in the judgment of blame itself; i.e. in the judgment that the agent has show a disregard or ill will. “That judgment—even if incorrect—makes it the case that you no longer stand in relations in which your good will is recognized on all sides” (Hieronymi, 2004, P. 124). Standing in relations of mutual regard is very important to us, and so blame carries a painful force. Because the force of blame is explained by the content of the judgment and its implication, blame cannot be rendered unfair by the fact that an agent could not have been reasonably expected to avoid that for which she is blamed. Judgments are only unfair if they are unjustified, and the historical conditions under which one came to have a certain attitude, or which led one to perform a particular action never render the judgment that an agent showed disregard or ill will unjustified all on their own (Hieronymi, 2004, p. 124).

I do not wish to take issue with Hieronymi’s argument that blame cannot be rendered unfair by its force; I take this argument to be highly plausible. However, it does not resolve the worry that blame might be rendered unfair in some other manner, and that for this reason, it might be unfair to blame agents for things we could not reasonably expect them to avoid. It is plausible that when we blame an agent for failing to show proper regard, we implicitly endorse the demand that she show proper regard.⁶⁴ But, if the demand is not itself

⁶³ For example, Wallace (1996) argues that the reactive attitudes associated with blame are painful for the recipient, Watson (1996) argues that to blame someone involves taking a stance that certain negative responses to them would be appropriate, and Wolf (1990) argues that to blame someone involves assessing something that is a deep aspect of their agency, and therefore to assess them in a particularly deep and painful manner. For Hieronymi’s (2004) discussion, and rejection of each particular, view see p. 119-121, 120-122, and 122-124.

⁶⁴ I have in mind demands in the sense of a claim one or more agents makes against another. We often do demand that agents fulfill their moral obligations, but obligations are

fair, we are treating the agent under judgment unfairly by endorsing it. For example, Susan is a stock broker. She works for another broker, Stan. If the market is open, either Stan or Susan must be in the office. Stan takes all major holidays off, and has done so for the past 10 years while he has worked with Susan. On his way out of the office, he informs Susan that he will, once again, be taking Christmas Eve off. Susan refuses this arrangement. Stan reports this to his wife, who blames Susan for refusing. In doing so, Stan's wife endorses the demand that Susan work on Christmas Eve. Because the demand that Susan work (yet again) on Christmas Eve is unfair, Stan's wife treats Susan unfairly in endorsing it (i.e. in blaming Susan for refusing the arrangement).

Something similar may be thought to occur when we blame agents for things they could not be reasonably expected to avoid. Consider an example offered by Hieronymi (2004).

Suppose that a coworker is both extremely competitive and extremely insensitive, due to untoward formative circumstances beyond her control. As a result, she is constantly putting others down and manipulating people and circumstances so as to gain the best of limited goods. As presently constituted, she is unable to respond to others in ways that show proper regard for them (Hieronymi, 2004, P. 126).

Hieronymi (2004) argues persuasively that it is not unfair to blame this agent due to the force of the judgment of blame. However, even if the judgment that this agent has shown ill will is accurate, it is not clear that it is fair to blame her if we take seriously the claim that she could not avoid her actions (e.g. manipulating others, putting others down), due to her constitution, and that she could not be reasonably expected to avoid her constitution due to

not themselves demands. E.g. I might demand that my husband fulfill his obligation to assume equal responsibility for our children, but his obligation to do so is not itself a demand. Therefore, on the view I sketch above, a demand that an agent fulfill his obligation might be unfair, while the obligation itself is not. I am not convinced that obligations can be unfair, but holding others to them certainly can be.

her poor formative circumstances. In blaming her, we endorse the demand that she show proper regard to others. If she cannot be reasonably expected to do so, then it is arguable that this demand is unfair, as is endorsing it.⁶⁵

If this is right, then Rosen's (2003) objection to non-volitional culpability views is a powerful one. It is highly plausible that some agents cannot be reasonably expected to avoid their moral ignorance, and it is highly plausible that it is therefore unfair to blame them for it. This is true even if we grant that it is not unfair to impose the force of blame upon an agent on account of X, despite the fact that she could not have been reasonably expected to avoid X. If an agent cannot be reasonably expected to avoid her moral ignorance, it is arguably unfair to demand that she do so, and thereby unfair to blame her for failing to do so. This suggests that the non-volitional culpability view is mistaken in holding that the mere fact that an agent's moral ignorance manifests an objectionable moral attitude implies that she is culpable for it, and in holding that the conditions under which an agent formed her moral are irrelevant to her culpability for it

3.6 Moving Forward

I have argued that Harman's culpability view and the non-volitional view face the same main problem. It is plausible that agents sometimes are not culpable for their moral ignorance because they could not be reasonably expected to avoid it. This is true even when an agent's moral ignorance manifests an objectionable moral attitude, or a failure to care sufficiently about things of moral significance. As noted above, it is open to proponents of culpability views to argue either (i) that agents can always be reasonably expected to avoid

⁶⁵ Hieronymi (2004) grants that "issuing a demand to someone unable to meet it seems not only pointless but also unfair" (118).

their moral ignorance, or (ii) that agents can be culpable for that which they could not be reasonably expected to avoid (e.g. that it is not, in fact, unfair to blame agents for that which they could not be reasonably expected to avoid). As also noted above, arguing for either claim requires an account of the conditions under which an agent can be reasonably expected to avoid her moral ignorance, and the relationship between whether or not an agent can be reasonably expected to avoid X and her culpability for Xing (or not X'ing). I offer such an account in the following chapter, where I argue that (i) and (ii) are implausible.

However, I also argue that we should not abandon culpability views altogether on these grounds. Proponents of culpability views could accept that agents are not culpable for their moral ignorance if they could not have been reasonably expected to avoid it, and therefore give up the claim that the conditions under which an agent formed her moral ignorance are irrelevant to her culpability for it, without relinquishing the view that, when agents are culpable for their moral ignorance, they are directly culpable for it. In doing so, they would retain the core elements of their view while adopting a more moderate, and intuitive stance towards culpability for moral ignorance.

PART III: REASONABLE MORAL EXPECTATIONS, THE MORAL IGNORANCE DEBATE, AND CULPABILITY FOR MORAL IGNORANCE

Chapter 4

Reasonable Moral Expectations and Culpability for Moral Ignorance

4.1 Reasonable Moral Expectations and the Moral Ignorance Debate

In Chapters 2 and 3, I argued that skeptical arguments against blameworthiness, and strict culpability views of moral ignorance, depend for their success on a general account of the conditions under which an agent can be reasonably expected to X, and the relationship between what can be reasonably expected of an agent, and her culpability. So far, no account like this has been offered. In this chapter, I develop one, and apply it to the moral ignorance debate. I then argue, on the basis of this discussion, that the skeptical argument against blame ought to be rejected entirely, and that strict culpability views of moral ignorance must be significantly revised.

After considering and rejecting common assumptions concerning the conditions under which a moral expectation is reasonable in section 4.3, I argue for an account of reasonable moral expectations (in section 4.4) comprised of the following sufficient and necessary conditions. A moral expectation is reasonable if it either requires an agent (i) to refrain from acting in ways that treat other agents as lacking inherent or equal value and (iii) is not pointless, or it if requires an agent (ii) to refrain from acting in ways that treat their

interests as lacking or being of less relative importance compared to the interests of others and (iii) is not pointless. A moral expectation is unreasonable if it meets neither (i) nor (ii), or if fails to meet (iii).

After considering and answering a number of objections to this account of reasonable moral expectations in section 4.5, in section 4.6, I argue that on this account of reasonable moral expectations, blaming an agent for failing to do that which she could not have been reasonably expected to do, is often unjustified. Moral expectations that are unreasonable on my view, it turns out, often also fail to treat like cases alike. When we blame an agent for failing to meet some expectation, we implicitly endorse the expectation. If the expectation itself is unfair, because it fails to treat like cases alike, in implicitly endorsing it, we fail to treat like cases alike. So, on my view, in blaming agents for failing to meet a demand they could not be reasonably expected to meet, we treat them or others unfairly.

In section 4.7, I apply this view of reasonable moral expectations and their relation to culpability to the moral ignorance debate. I argue that, given my account of reasonable moral expectations, proponents of skeptical arguments cannot defend a key premise of their argument: moral ignorance is culpable only if it is the result of a knowing mismanagement of one's moral beliefs. In Chapter 2 I argued that this premise is only plausible if it is also true that agents can be reasonably expected to avoid their moral ignorance only if it is the result of a knowing mismanagement of their moral beliefs. Using the account of reasonable expectations that I defend in section 4.4, I argue that this further claim is implausible. If this is right, proponents of the skeptical argument lack the resources to defend one of their

central premises. Given this, I argue, we ought to reject the skeptical argument, and thereby the view that agents are almost never culpable for their moral ignorance.

In Chapter 3 I argued that it is plausible that given their epistemic or cultural circumstances, some agents could not be reasonably expected to avoid their moral ignorance, and that some agents could not be reasonably expected to avoid the insufficient moral care or objectionable moral attitudes manifested by their moral ignorance. In this chapter, I offer a positive argument for both claims. In light of them, I argue that proponents of culpability views should not hold that agents are always culpable for their moral ignorance insofar as it manifests insufficient moral care or objectionable moral attitudes. Instead, they ought to hold that agents are culpable for their moral ignorance if it manifests these objectionable states *and they could have been reasonably expected to avoid both their moral ignorance and the states in question*. While this constitutes a significant alteration of culpability views, it does not require their proponents to abandon their main aims and claims. They can still maintain that agents are very often culpable for their moral ignorance.

The result of this discussion is what I hope will be a highly useful account of reasonable moral expectations, and a highly intuitive approach towards agents' culpability for their moral ignorance. Although I do not take a stance on when agents are culpable for their moral ignorance, I argue that, whatever view one adopts concerning the latter, one must accept the condition that agents are only culpable for their moral ignorance, if they could have been reasonably expected to avoid it. This condition prevents participants in the moral ignorance debate from taking up the extreme positions that have characterized it up until this point. When one accepts this condition, one arrives at a more moderate view: although agents are culpable for their moral ignorance more often than not, there are real exceptions.

4.2 Moral Expectations: A general description

Before turning to my main arguments, it will be helpful to outline the general notion of moral expectation with which I will be working. I understand having an expectation of another agent in terms of holding the agent to a demand.⁶⁶ I'll unpack this view in a moment, but first notice that expecting something of an agent, in the sense of holding her to a demand, is different from holding a predictive expectation about how she will act, or judging that she ought to act in a particular way. I can hold an agent to a demand without also predicting that she will act so as to fulfill it. E.g. I might hold all of my students to the demand that they complete the reading for the course; my course might be intentionally designed such that students who read tend to do well in many ways that students who do not read cannot (e.g. on reading quizzes). This can be true even if I predict that many of my students will not complete the reading. Similarly, I might hold drivers in the Atlanta metro to the demand that they not look at their cell phones while driving. I may give them disapproving looks when I see them doing this, or I may toot my horn and wag my finger at them, or even report their activity to the police. And, again, this might be true even if I

⁶⁶ This is Wallace's (2008) understanding of expectations. However, Wallace understands holding an agent to a demand to be captured by being susceptible to the reactive attitudes if the demand is flouted, or at least being disposed to believe these attitudes to be appropriate if the demand is flouted. I part ways with him here. My own view of what is involved in holding an agent to a demand is heavily influenced by MacNamara's (2011) account of "holding demands." A holding demand "second-personally imposes normative burdens" on another agent... "with the aim of inducing in her first-personal practical uptake of the ought that binds her" (MacNamara, 2011, p. 92). These demands are, "the forward-looking analogue of punishments and reproofs" (MacNamara, 2011, p. 92). Tognazzini (2015) briefly discusses interactions aimed at "enforcement of a standard" which he characterizes as attempts to provide forward-looking motivation to act as one ought (40). He suggests these interactions are something like holding demands, but that they do not only involve moral obligations (as MacNamara suggests). My view also bears similarities to this idea.

predict that the majority of drivers in the Atlanta metro *will* look at their phones while they drive. Moreover, when we hold an agent to a demand, we do something besides judging that they ought to comply with the demand. For example, I might tell my husband, for the millionth time, that our daughter really should practice piano more often. He might reply “well, *do* something about it. Hold her to it.” If holding her to the demand that she practice the piano more often involved nothing more than judging that she ought to comply with the demand, his suggestion that I hold her to it would not make sense.⁶⁷

When we have an expectation of an agent in the sense in which I am concerned, we hold her to a demand. When we hold an agent to a demand, we guard against her future infractions of it. This might involve taking proactive measures to guard against her infractions, or responding to her infractions of the demand with measures intended to guard against her future infractions of it.⁶⁸ These “measures” might take the shape of threats of negative consequences, negative reinforcement, penalties, and reproof or rebuke, all *aimed at the prevention of future infractions*. Suppose, for example, that I demand that my daughter practice piano more often. I might guard against her (future) infraction of this demand by asserting a rule: that she cannot play with her friends until she has practiced piano, and demanding that she follow this rule on pain of suffering some negative

⁶⁷ I borrow this type of example from MacNamara (2011), who makes a similar point (about holding an agent to a demand being non-reducible to moral appraisal) with a similar example (90).

⁶⁸ I have no position about how to judge whether, in a particular instance, we in fact hold an individual to a demand. It seems that in some cases (e.g. cases in which we recognize a high probability that the individual in question will violate the demand), in order for it to be plausible that we actually hold the individual to the demand, we must take proactive actions to guard against infraction. In other cases, we plausibly hold an agent to a demand simply by reacting to its infraction in ways intended to guard against future infraction. It is sufficient for my purposes here that we can be said to hold an agent to a demand if we are ready to guard against its infraction, while acknowledging this requires different things in different circumstances.

consequence. If enforcement proves necessary, I might withhold her allowance if she does not practice. Or I might reprove her if she does not practice, “It’s only 20 minutes! You can spare 20 minutes to learn a skill that you will value for the rest of your life.”⁶⁹

To hold a moral expectation of an agent is to hold the agent to a moral demand, i.e. a demand targeted to an agent qua moral agent (as opposed to as a friend, boss, husband, and so on) and one the issuer takes to be primarily supported by moral reasons. E.g., a shopkeeper might expect others not to steal from his store. The demand that they not do so is moral insofar as it is addressed to moral agents qua moral agents, and the shopkeeper (as I am imagining him) takes it to be primarily supported by moral reasons, e.g. that stealing is dishonest, unjust, disrespectful, and so on.⁷⁰ The means through which we guard against the infraction of moral demands (by those to whom our expectations are directed), are similar to the means through which we guard against the infraction of non-moral demands. We guard

⁶⁹ As this paragraph suggests, I think the activity of holding an agent to a demand involves imposing (or being ready to impose) a negative consequence if the demand is infringed. We do often encourage desired behavior through positive means. E.g. I would likely be more successful in getting my daughter to practice piano if I praised her lavishly for doing so, offered her rewards for doing so, or figured out a way to make practicing really fun for her. However, while employing these forms of positive reinforcement are likely to bring about increased practicing on her part, I do not think that in employing them, I could be said to be holding my daughter to the demand that she practice piano more often. Offerings of positive reinforcement invite and encourage a behavior. Demands insist on a behavior. So, I am not convinced that we can hold an agent to a demand through positive reinforcement. This line of thought is echoed in the work of other philosophers. E.g. Watson (1996) argues that holding agent responsible involves holding an agent to a requirement by responding to her violation of this requirement with “sanctions” and “unwanted or adverse treatment” (237).

⁷⁰ I do not mean to suggest that all moral demands are sufficiently supported by moral reasons, but that whatever reasons the issuer takes to support them, are primarily moral reasons. E.g. an anti-abortion activist might demand that fellow activists make women feel terrible about their decision to obtain an abortion, in the hopes that these women will change their minds. I take it that this is a moral demand, even if it is not sufficiently supported by moral reasons. The reasons that the activist takes to support this demand are mostly moral, e.g. that this will protect a fetus’ right to life, or prevent innocent children from being harmed.

against infractions of moral demands by threatening negative consequences for those who violate them, and by responding to infractions by imposing these negative consequences.⁷¹ Imagine, e.g., that Augie is in her 90's and though she still gets around, she uses a walker, and is obviously a bit weak and unsteady on her feet. She boards a bus on which all the seats are taken. She scans the passengers and notices that with the exception of a healthy young woman, who is focused on a book she is reading, all of the passengers in the front seats are also elderly or disabled. Augie politely asks the young woman for her seat. The woman briefly looks up, indicating that she has heard Augie's request, but ignores her and continues to read. I take the demand that the young woman give her seat to Augie to be a moral demand. It is issued to the young woman qua moral agent, and Augie (or someone else) issuing it would likely take it to be supported by moral reasons; to give up her seat involves no significant sacrifice on the part of the woman, and would prevent risk of significant harm to Augie (from standing during the ride). Augie might hold this young woman to this moral demand in ways similar to those we saw above. She might try to persuade the bus driver to give the woman the option of giving up her seat or getting off the bus. Or, Augie might verbally rebuke the woman, informing her that "Common decency calls for you to move. You are being disrespectful and inconsiderate. Get up."

⁷¹ We are often not in a position to threaten or impose penalties upon moral agents, qua moral agents. E.g. I am in no position to penalize another agent for taking my parking spot, or for failing to help as I am scrambling to pick up my spilled belongings. For this reason, I think we most often guard against infractions of moral demands through verbal rebuke and reproofs. My discussion will focus on these forms of enforcement. One might worry that verbal rebuke and reproofs that are responses to anticipated or actual violations of moral demands are simply expressions of the negative reactive attitudes, and are therefore properly characterized as forms of blame rather than attempts to guard against future infractions of demands. While verbal rebukes and reproofs aimed at preventing future infractions of a moral demand might "air" our negative reactive attitudes, they do not aim at doing so; their aim is forward-looking, not backward looking. I therefore do not think they are properly construed as forms of blame.

On the view I have sketched, expecting an agent to X—holding her to the demand that she X—is a forward looking activity. When we wonder whether or not an expectation is reasonable, on my view, we are questioning whether this forward looking activity, of holding an agent to a demand, is justified. Is it appropriate to hold this agent (or set of agents) to this demand; is it appropriate to guard against their present and future infractions of it? Our answer to this question will depend on at least three things: (1) the appropriateness of the demand itself, (2) the relationship between the demand and the recipient (e.g. what is required for *this* agent to comply with the demand), and (3) the relationship between the agent holding to the demand, and the agent being held to it. Notice that we might also wonder whether the means through which we hold an agent to a demand are justified. While there surely are norms that govern the appropriateness of the means through which we hold agents to demands, it is the question of whether or not it is appropriate to hold an agent to a particular demand at all that is relevant to an agent's culpability. And, this is the question on which I will exclusively focus.

It might at first seem puzzling that the reasonableness of holding an agent to a demand (a forward looking activity) could impact the legitimacy of blame (a backward looking activity or judgment). Along similar lines, one might wonder if moral expectations as I have described them above, are the same expectations that philosophers typically have in mind when they claim that agents cannot be culpable for that which they could not be reasonably expected to avoid. I can offer a few initial words of reassurance.

When we blame an agent, we often do so insofar as she has failed to live up to some a demand we have placed upon her. E.g. we blame agents for failing to respect other agents, or for failing to be fair or honest. If it was inappropriate to issue a demand, this will

undermine the propriety of blaming the targeted party for failing to meet that demand. It is unreasonable of me to demand that my neighbor pay my rent. And it would be similarly inappropriate for me to hold her to this demand by imposing or threatening negative consequences (e.g. reproof). It is plausible that this has implications for whether or not I am justified in blaming her for failing to pay my rent.

However, when trying to determine whether or not an agent is culpable for X, by way of considering whether or not he could be reasonably expected to avoid X, we are typically faced with a situation that is much more complex. Instead of wondering if we would be justified in blaming an agent for an action that she might perform or fail to perform (e.g. not paying my rent), we are in the position of wondering if she is culpable for an action or omission that has already taken place. In even more complicated cases, we are considering whether or not an agent is culpable for something she did in the distant past (perhaps before we even existed), or whether she is culpable for something that cannot be done differently in the future. Consider the following examples. (1) We wonder whether Luis is culpable for lying to his wife last night about the nature of one of his former relationships. (2) We wonder whether or not Luis is culpable for lying to his wife, years ago, about the nature of one of his former relationships. (3) We wonder whether or not Luis is culpable for chopping down his wife's beloved apricot tree without her consent. In the first case, one might rightfully wonder what the appropriateness of guarding against Luis's future infractions of the demand that he not lie to his wife about his former relationships, has to do with whether or not he is culpable for having lied last night. This may seem even more puzzling in the second case, when more time has elapsed between the past misdeed and the

present expectation. In the third case, there is no relevant present demand to which to hold Luis. The deed is done; the apricot tree is gone.

In fact, I do not think these cases are as complicated as they seem. In some cases in which we are wondering whether or not an agent is culpable for something she has already done, the fact that we could (now) reasonably hold her to the demand that she not X, is plausibly relevant to her culpability for Xing. This is true when nothing significant has changed in an agent or her situation, between the time she performed the act in question, and the present. For example, if nothing has changed about Luis or his situation since last night, then whether or not we can now reasonably hold him to the demand that he be honest with his wife about his past relationships is relevant whether or not we can justifiably blame him for violating this demand last night. If it would be inappropriate to reprove Luis for lying to his wife, with the aim of preventing his future infractions of the demand that he not do so, then plausibly bears on whether or not he is blameworthy for having lied to her. Cases in which a lot of time has elapsed, or there is no relevant demand to which we could now hold the agent are, admittedly, a bit more complicated. In these cases, the fact that we could (or could not) now reasonably hold the agent to a demand that she X may not bear on her culpability for \sim Xing. But, the fact that we (or others) could have (or could not have) reasonably done so at the time that she \sim Xed, is relevant to her culpability. For example, it is possible that much about Luis' situation when he lied to his wife years ago is different from his present situation. Perhaps he had just met his wife, and had no idea he would marry her. Perhaps he thought she was a one-night-stand, and did not see the point of honestly discussing his past with someone he would never see again. Under these circumstances, whether or not we could now reasonably hold Luis to the demand that he not lie to his wife,

is irrelevant to whether or not he is culpable for lying to her all those years ago. However, whether or not we (or someone else, perhaps his wife) could have *then* held him to this demand, is plausibly relevant to his culpability for violating it.⁷² In cases in which there is no relevant demand to which we could now hold the agent, the question of whether or not we could reasonably do so is obviously not relevant to the agent's culpability for her past action. For example, we cannot hold Luis to the demand that he not chop down his wife's favorite apricot tree without her consent; the tree is gone. But, it is either the case that we (or others) could or could not have reasonably held Luis to the demand that he not chop it down without her consent, at the time that he did. And if, for some reason, it would have been inappropriate to guard against Luis' infractions of this demand, at the time he chopped down the tree, surely this is relevant to his culpability for violating this demand.

4.3 Towards a General Account of Reasonable Moral Expectations: Common Assumptions

To my knowledge, although philosophers often appeal to reasonable expectations, and assume that agents cannot be blamed for that which they could not be reasonably expected to avoid, no one has attempted to offer a general account of the conditions under which a moral expectation is reasonable. However, there are common themes in the remarks and arguments that philosophers make with respect to what can be reasonably expected of an agent. Philosophers tend to hold that whether or not an expectation is reasonable, depends on the capacity of the agent to whom the expectation is communicated, the difficulty an

⁷² It may, of course, be difficult for us to determine whether we or others could have reasonably held an agent to a demand at some time in the past. But, this does not impact the claim that the fact that we (or others) could or could not have done so is plausibly relevant to whether or not the agent is culpable for his past violation of the demand.

agent faces in meeting it, and the fairness or unfairness of expecting her to act in the manner demanded of her.⁷³ In this section, I consider whether any of the latter are conditions on the reasonableness of moral expectations, and argue that none of them are. To arrive at a general account of reasonable moral expectations, by identifying the conditions under which moral expectations are reasonable or unreasonable, we must abandon these traditional approaches.

Capacity and Reasonable Expectations

A number of philosophers hold that, in order for a moral expectation to be reasonable, (i) the agent to whom it is communicated must have the capacity to meet it.⁷⁴ Some also hold (ii) that the agent must have a fair opportunity to exercise this capacity.⁷⁵ I argue that in its current formulation, (i) is implausible, but does point to an important related condition on the reasonableness of an expectation: in order for an expectation to be reasonable, it must not be pointless for us to hold an agent to the demand in question. I argue (ii) is better understood as a claim about the difficulty or fairness of an expectation.

⁷³ E.g., Wolf (1986) argues that it is unreasonable to expect agents with extremely deprived backgrounds to see anything wrong with their immoral actions insofar as they lack the capacity to do so (232). Levy (2009) argues that agents can be reasonably expected to conform their behavior to normative standards “only if they have the capacity to behave accordingly” (735). Ciurria (2014) defends an account of reasonable expectations according to which an agent can be reasonably expected to avoid X only if she has the capacity to do so, and a fair opportunity to exercise this capacity. Wallace (1996) argues that agents cannot be reasonably expected to avoid their actions if they lack “powers of reflective self-control,” or the powers to grasp moral reasons and to control one’s behavior in light of them (215-217). Benson (2001) argues that we sometimes cannot reasonably expect an agent to do something if it “would be so extremely difficult that we could not fairly hold her to the demand that she do so” (614). Fricker (2016) argues that moral expectations are only reasonable if they are not excessively demanding (168).

⁷⁴ E.g. Wolf (1986), Ciurria (2014), and Wallace (1996)

⁷⁵ E.g. Ciurria (2014)

There are two ways to interpret (i). One might understand it to amount to the claim that an expectation is reasonable only if the agent does not lack the very general capacities required to comply with it, such as, e.g., the capacities to form beliefs in response to reasons, alter those beliefs in response to new evidence, form intentions, act on those intentions, and so on. Alternatively, one might interpret it as the claim that an expectation is reasonable only if the agent does not lack the *personal* capacity to comply with it, determined by the particular psychology of the agent, at the relevant time. Imagine, e.g., that Carol witnesses a brutal attack. She is paralyzed by fear—she is frozen, unable to think or move—and does nothing to aid the victim. This reaction is a non-voluntary, physiological reaction that is the result of a malfunction of her cerebellum.⁷⁶ Carol has lived a sheltered life, and has never truly believed herself, or someone nearby, to be in serious danger. She was therefore not aware that she would have this reaction to perceived danger. While Carol has the general capacities required to aid the victim (e.g. the ability form beliefs and intentions, to act on those intentions, and so on), she lacks the personal capacity to do so. Given her psychology and physiology at the time of witnessing the attack, she cannot help her.

If anything, the reasonableness of an expectation depends on the personal capacities of the agent to whom it is addressed. While an agent must have the general capacities of moral agency in order for it to make any sense to hold her to a moral demand, it is not plausible that this is a condition on the reasonableness of a particular moral expectation. The fact that an agent lacks the general capacities to comply with a moral expectation indicates that no moral expectations (and not just this one in particular) ought to be addressed to her.

⁷⁶ For a discussion of the physiological responses (and in particular, the neural pathways) involved in the phenomenon of being frozen with fear, see Koutsikou et. al. (2014).

So, the question is whether we should accept the personal capacity claim: in order for a moral expectation to be reasonable, an agent must not lack the personal capacity to fulfill it.⁷⁷ Recall that in having a moral expectation of another agent, we hold her to a moral demand; we guard against her future infractions of the demand. So, our question is whether or not we should accept the claim that if an agent lacks the personal capacity to comply with a demand, this entails that it is unreasonable to guard against her infraction of it, by threatening and imposing negative consequences aimed at preventing such infractions. I argue that there is insufficient support for this claim.

One might argue that it is pointless to hold an agent to a demand if she lacks the personal capacity to comply with it. However, I am not convinced that this is true, except in a very small minority of cases. In many instances, an agent *temporarily* lacks a personal capacity to comply with a demand. In such cases, it is not pointless to nevertheless hold the agent to the demand. For example, drunk drivers temporarily lack the capacity to comply with the demand that they not drive recklessly. In holding a drunk driver to the demand that he not drive recklessly, we might verbally rebuke him for doing so. “How could you have been so stupid? Don’t you realize that you could have killed someone or yourself? You could have killed a child! You cannot take risks like that.” Such a rebuke, aimed at preventing the agent’s future infractions of the demand that he not drive recklessly, is a way of holding him to this demand. While this might be unsuccessful in preventing future infractions on his part, it certainly isn’t pointless. It makes good sense to try, with whatever

⁷⁷ This is what most philosophers have in mind when they suggest that an agent’s incapacity renders an expectation unreasonable. Most tend to focus on the particular, personal capacities of the agent (e.g. his ability to feel empathy with others).

resources we have, to prevent future instances of reckless (drunk) driving on behalf of this agent.

Even when an agent *permanently* lacks the personal capacity to comply with a demand, I am not convinced that it follows that it is pointless to hold her to it. Though we may sometimes permanently lack the personal capacity to comply with a demand, we may often be able to compensate for this lack of capacity in ways that allow us (eventually, with the help of others) to bring it about that we not violate the demand. Consider Carol again. Imagine that Carol's brain malfunction cannot be treated through therapy, surgery, or drugs. She will always, unavoidably respond to the perception that she or others are in serious danger by freezing in terror. We might nevertheless demand that Carol do something in response to perceived danger. For example, we might (gently) reprove her. "Look Carol, I know it's not your fault that you freeze in these types of situations. But this can't go on. It's not safe for you or others. You need to be able to do something to help yourself or other victims in an emergency. You can't just stand there and watch it happen!" I am not convinced that this is pointless. While Carol might not be able to change the fact that she freezes in the presence of perceived danger, she might be able to bring it about in other ways that she can do something to secure her own safety and the safety of others (e.g. she might be able to wear an alert bracelet that she only needs to push in order to get assistance). It is at least possible that our reproof will prevent Carol's future infraction of the demand that she do something to help when she or others are in danger. It is therefore not pointless to hold her to this demand by taking measures to guard against her future infractions of it.

There may be some cases in which it is pointless to hold an agent to a demand, insofar as the agent in question entirely lacks the capacity to bring it about (even at some

distant point in the future, even with assistance from others) that she comply with the demand. Such cases would be very rare. However, they do exist. E.g. we ordinarily hold agents to the demand that they recall important facts about family members. It would be pointless to hold an agent with advanced Alzheimer's disease to this demand if he permanently lacks the capacity to do so, and there is nothing he could do, even with the assistance of others, to bring it about that he do so. But, again, such cases are rare, and they are not typically what philosophers have in mind when arguing that an expectation is unreasonable insofar as an agent lacks the capacity to fulfill it.⁷⁸

Alternatively, one might think there is something unfair about having an expectation of an agent that they cannot meet, and that for this reason expectations are reasonable only if the agents to whom they are directed have the personal capacity to meet them. Having an expectation of an agent amounts to holding her to a demand, which in turn involves threatening and imposing negative consequences upon a person aimed at preventing her from violating the demand. One might think it is unfair to threaten and impose negative consequences upon an individual for something she lacks the capacity to avoid. However, I think this depends on the purpose of imposing the consequences. It is plausible that it is unfair to impose negative consequences, for retributive purposes, upon an agent for violating a demand with which she lacked the capacity to comply. Negative consequences that are imposed for retributive purposes are justified, in part, by the fault of the agent in question. If

⁷⁸ Most philosophers who endorse the capacity claim do not have in mind agents who lack the capacity to bring it about (sometime in the future, with the help of others) that they comply with the demand. Instead, they have in mind agents who face significant challenges in doing so. E.g. Wolf (1986) argues that agents with very poor formative circumstances might lack the capacity to comply with a demand (232). Ciurria (2014) argues along similar lines, that given his terrible childhood, and the fact that he may have been suffering from mental illness, Robert Alton Harris lacked the capacity to comply with the moral demand that he not murder others (9-10).

there is no fault, these consequences are not justified. Negative consequences imposed for the purpose of preventing infractions of a demand are justified by the worthiness of this aim. If we assume this aim is justified, then in most cases, the fact that an agent lacked the capacity to comply with the demand in question, does not render it unfair to nevertheless hold her to the demand. Again, if it is impossible for the agent, even with time and the assistance of others, to ever comply with the demand, then holding him to it would be pointless. In such cases, we would be unjustified in holding the agent to the demand insofar as we are not justified in imposing pointless negative consequences upon others. But, again, in most cases the fact that an agent lacked the capacity to comply with a demand does not entail that she cannot bring it about that she eventually does not violate the demand. In these cases, it is not unfair, at least for this reason, to impose negative consequences upon an agent in order to prevent future infractions of the demand.

Philosophers who endorse (ii) typically understand an agent to have a fair opportunity to exercise her capacities to fulfill an expectation as long as she does not face significant obstacles in doing so. E.g. FitzPatrick (2008) suggests that agents have a fair opportunity to exercise their capacities (with respect to avoiding moral ignorance) as long as they do not face personal or societal limitations that would make doing so exceptionally difficult (605).⁷⁹ Ciurria (2014) argues that agents who suffer poor formative or cultural circumstances lack the fair opportunity to exercise their capacities to fulfill relevant moral

⁷⁹ FitzPatrick (2008) argues that Potter (from *It's a Wonderful Life*) faced no "relevant limitations in his social context or in his capabilities that should have made the necessary broader reflection and information gathering impossible or unreasonably difficult for him," he could have been reasonably expected to avoid his moral ignorance (605).

expectations.⁸⁰ For this reason, I think (ii) is best understood either as the claim that in order for an expectation to be reasonable, it must not be too difficult for an agent to fulfill it, or it must not be unfair, as opposed to understanding it to be a claim that is focused on the capacities of the agent in question.

Difficulty and Reasonable Expectations

Let's turn to the claim that the reasonableness of an expectation depends on the difficulty one faces in fulfilling it.⁸¹ One might understand this claim to amount to the claim that an expectation is reasonable only if one would not face significant difficulty in complying with it. I argue that this claim is implausible.

The fact that it would be significantly difficult for an agent to meet an expectation does not, by itself, entail that the expectation is unreasonable. It may be significantly difficult for an agent to fulfill an expectation given the physical or psychological effort he would have to exert in order to fulfill it (e.g. withstanding torture, resisting an addictive urge), or given the costs he would endure in order to do so (e.g. by sacrificing relationships or giving up a lot of valuable time). For example, it may be very difficult (psychologically and physically) for a doctor to be alert and careful after a 48 hour shift, but it is plausible that her patient and others could still reasonably expect her to be. It would not be inappropriate for the patient, or others, to impose negative consequences upon the doctor in

⁸⁰ Cuirria (2014) argues that a sexist man living in the 1950s in the USA could be understood as lacking the fair opportunity to exercise his capacities in order to avoid sexist attitudes, insofar as he lacks "epistemic alternatives" in his sociohistorical setting (7). She also argues that extenuating circumstances such as "duress, severe childhood abuse, and the inaccessibility of requisite ethical resources," can deprive an agent of the relevant capacities or the fair opportunity to exercise his capacities to meet an expectation (9).

⁸¹ Benson (2001), FitzPatrick (2008), and Fricker (2016) can all be interpreted as endorsing this view.

order to prevent her from violating this demand. E.g., it would certainly be appropriate for another physician to reprove the doctor if she is being careless, “I know you’re exhausted, but you *must* wake up and pay attention to what you are doing!” It is also plausible that others could reasonably expect a physician to be honest about a mistake she made, even if she might lose her job as a result. E.g. upon hearing her plan to cover up her mistake, her friend might justifiably rebuke her, “You can’t lie! Who could trust you again if you do? Your integrity is more important than your job, and the board and your patient are entitled to the truth. You are better than this.” Similarly, it is plausible that an agent who is seriously injured on the side of the road could reasonably expect me to help him, even if doing so would require me to forgo an exceptional and rare opportunity. Imagine I tell the agent that I am not sure I can help. It would require me to miss an interview that I worked very hard to get, and that I will not likely be offered again. It would not plausibly be inappropriate for the injured agent to reprove me, “What do you mean you cannot help? You might miss a *job interview*? I am badly hurt! Look at my leg. I need your help!”

Given these examples, one might think that it is the difficulty of fulfilling an expectation, balanced against the degree to which others need us to fulfill it, that bears on its reasonableness. Perhaps an expectation is unreasonable if an agent would face great difficulty in fulfilling it, and the need for her to fulfill it does not outweigh this difficulty. However, this too is implausible. Imagine, e.g., that Samantha suffers from kleptomania.⁸² She experiences intense urges to steal items that she could easily afford, and for which she has no need or particular desire. She often throws the stolen objects in the trash, and isn’t

⁸² This description is based on the clinical characteristics of kleptomania, and reports of case studies in Grant and Kim (2002). Kleptomania is no longer considered to be an individual disorder, but is instead considered to be a type of impulse control disorder.

sure why she steals them in the first place, except that it is the only way to relieve the tension she feels just prior to stealing. She is ashamed of her behavior, and has tried many times to resist her urges to steal. She has even considered suicide to stop the urges. But, whenever she tries to resist the urges, they get stronger. It is obviously very difficult for Samantha to comply with the demand that she not steal. However, it is at least not obvious that a shopkeeper could not reasonably expect her not to steal small items from his store, even if this would have only a minor negative impact on him and his business. There is at least nothing obviously inappropriate about him holding her to the demand that she not steal small items from his store. He might tell her sternly, “I have my eye on you.” Or, he might respond to her theft through verbal rebukes, or by calling the police. Even if he is well aware of her disorder, it is not plausible that it would be unreasonable for him to hold her to the demand that she not steal from his store. In response to her theft, it would still be appropriate for him, e.g., to reprove her, “It is not okay to steal, no matter how ill you are!” It is implausible that the shopkeeper just has to tolerate that Samantha is going to steal from his store, and give up holding her to the demand that she not do so.

It is highly plausible that the reasonableness of an expectation is sometimes sensitive to the difficulty an agent faces in meeting it. E.g. if it is very easy for me to get medicine that would help my sick husband who is too ill to fetch it himself, it is plausible that he could reasonable expect me to do this. If, on the other hand, I would have to travel the Amazon rainforest in search of a rare plant, he could not reasonably expect me to do it (under normal circumstances). What I have said above does not deny that the reasonableness of our expectations is sometimes sensitive to the difficulty an agent would face in meeting them, but rather that the relationship between this difficulty and the

reasonableness of an expectation cannot be spelled out in terms of the former being a condition on the latter. In section the following section (4.4) I argue that the relationship between the difficulty an agent faces in meeting an expectation and the reasonableness of the expectation is indirect; sometimes the former impacts whether or not an expectation meets a different condition for reasonableness.

Reasonable Expectations and Fairness

Finally, one might think that the fairness of an expectation is a necessary and/or sufficient condition for its reasonableness.⁸³ We sometimes do, in fact, use the terms interchangeably. We might, for example, describe the prices at a store as “reasonable” or “fair”. We might describe a professor’s expectations of her students, or parents’ expectations of their children as “reasonable” or “fair.” A law might be called “reasonable” or “fair,” and so on. There are three main ways a moral expectation might be thought to be unfair: (1) agents may not have an equal opportunity to fulfill it, (2) it may impose a disproportionate burden upon one agent as compared to others, or (3) it may not treat like cases alike. I argue that (1) suggests only that a state of affairs is unfair, and not that an expectation is unfair. With respect to (2), I argue that some expectations may involve imposing a disproportionate burden upon an agent, but it is not clear that this, alone, entails that the expectation is unfair. Expectations would be unfair if (3) they fail to treat like cases alike, but it isn’t readily apparent how an expectation might do so.

⁸³ I am not convinced that anyone holds this view, although some hold the reverse, that the reasonableness of an expectation is a necessary and/or sufficient condition for its fairness. E.g. FitzPatrick (2008) argues that whether or not agents are culpable for their moral ignorance turns on what, if anything, the agent could “reasonably (and hence fairly) have been expected to have done in the past to avoid or to remedy that ignorance” (603).

One might think that an expectation is unfair if the agents to whom it is directed do not have an equal opportunity to fulfill it. However, while it may be unfair that agents lack equal opportunity to fulfill an expectation, this consideration (alone) does not suggest that the expectation itself is unfair. For example, an algebra teacher's expectation that students be able to solve basic equations is not rendered unfair by the fact that his students lack equal opportunity to fulfill it. It may be unfair that some of his students lack parental support, financial resources, or lots of free time to spend on their homework, while other students have ample support, resources for outside tutors, and no obligations that compete with their school work. But, this does not render the expectation that students learn basic equations in an algebra class, unfair. E.g. it does not make it the case that it is unfair for the teacher to test their algebra skills, and to base their grades in part on this, or to reprove them for being lackadaisical about learning algebra, with the purpose of preventing them from violating the demand that they learn it.

A closely related claim is that expectations can be unfair, and thereby unreasonable, if they impose a disproportionate burden upon some agents in comparison to others. There are two ways having an expectation of an agent might be thought to impose a burden on that agent. When we have an expectation of an agent, we hold her to a demand by guarding against her infraction of it. We do the latter by threatening and imposing negative consequences upon an agent for infractions, with the aim of preventing them (or future ones). This practice might be thought to impose a burden upon the agent insofar as, in order to avoid negative consequences, the agent must comply with the demand. Alternatively, this practice might be thought to impose a burden upon an agent in the sense that it involves imposing negative consequences for infractions of the demand—these negative

consequences may themselves be thought to be burdens. I'll focus on the first sense in which an expectation imposes a burden upon an agent.⁸⁴

On this suggestion, the burden imposed by an expectation is arguably comprised by the effort the agent must exert in order to comply with the demand (and thereby avoid the negative consequences). If this is right, then even if having an expectation of an agent is not unreasonable or unjustified in other ways (e.g. the agent issuing the demand does not fail to treat like cases alike, holding the agent to the demand is not pointless, the demand itself is justified), the expectation may impose a disproportionate burden upon the agent who is held to the demand. However, it is not plausible that the expectation is unfair for this reason; it is implausible that an expectation is fair only if it imposes equal burdens upon all agents who are subject to it.

If an expectation is otherwise justified, it is not rendered unfair by the fact that it imposes disproportionate burdens upon an agent, as long as this imposition is not intended by the issuer of the expectation.⁸⁵ To see this, consider the following example. An employer

⁸⁴ Expectations may impose a disproportionate burden upon an agent in the second sense if the individual issuing the demand imposes harsher consequences upon one agent for failing to comply with a demand than she does with respect to others who fail to do so. In some cases, this might be unfair. But, this concerns the fairness of the means through which an agent holds an agent to a demand. And, as noted above, I am concerned with the reasonableness of holding an agent to the demand at all.

⁸⁵ There are many complicated cases in which agents are disproportionately burdened by an expectation, not intentionally, but due to undetected racism or sexism. In such cases, the expectation is arguably unjustified. I do not discuss such cases above for two main reasons. I think in these cases, if the expectation is unjustified, this will often not be due to the fact that it imposes a disproportionate burden upon some agents, but because it holds them to an unjustifiable demand. E.g. an employer might have the expectation that employees must ask for a raise in order to get one. This expectation arguably places a disproportionate burden upon women, who tend to find asking for a raise to be intimidating and scary [see Ludden's (2011) discussion of Linda Babcock's research on this matter]. This result may be unintended by the employer, but may nevertheless seem to render the expectation unfair. I agree that the expectation is unjustified, but I am not convinced this is because it

expects all of their employees to arrive at work on time. She runs a stock brokerage firm, and because the market is only open during set hours, it is important that her employees are present at work during these hours. She holds her employees to this demand by having a policy that anyone who is frequently late will be subject to termination. It is much less difficult for some agents to meet this demand than it is for others. For example, agents who have no children, and have reliable transportation face far fewer obstacles in arriving at work on time than individuals who have to get their children ready for school or childcare, and then take public transportation to work. However, if the fact that some employees are disproportionately burdened by this expectation is not an outcome intended by the employer, but instead an accidental result of the agents' circumstances, it is not clear why it alone would render the expectation that agents arrive at work on time, unfair.⁸⁶

(unintentionally) places a disproportionate burden upon women, but because the demand that employees must ask for a raise in order to get one, is arguably unjustified. Raises ought to be based on performance, and not on the act of asking for one. Similarly, one might think that holding students to the demand that they do well on the SAT is unjustified insofar as it (hopefully unintentionally) places a disproportionate burden upon underprivileged black students. Santelices and Wilson (2010) argue that some questions on the verbal section of the SAT favor white students insofar as they involve words commonly used in middle class white communities that are not commonly used in underprivileged black communities. These words may have contextual meaning that is well understood by middle class white students, but not by underprivileged black students. If this is right, then the expectation that they do well on the SAT arguably imposes a disproportionate burden upon underprivileged black students who must learn the contextual meaning of words as they are used in privileged white communities in order to meet this demand. I am sympathetic to the thought that holding students to the demand that they do well on the SAT (e.g. by making scholarships and college admissions contingent upon this) is unjustified. However, again I am not convinced that this is because it places a disproportionate burden upon some students, but instead because there is something unjustified about the demand itself (i.e. that all students do well on an exam that tests the contextual meaning of words as they are used in privileged white communities).

⁸⁶If instead, the employer hopes to “encourage” poor, working parents to leave its workforce by imposing greater burdens upon them than those faced by their richer, non-parent counter-parts, this would be unfair. I am not convinced that we should characterize such actions as instances of holding a group of agents to a demand, since negative

One might worry that if an expectation places an agent under a very significant burden while not doing so with respect to other agents who are subject to it, it is unfair. However, I am also not convinced that this, alone, entails that an expectation is unfair. Imagine, e.g., that in order to arrive at work on time at 9am, Gloria must get up at 4 am each weekday, and pay for early morning childcare. In the morning, she wakes up, gets herself ready, and then wakes her toddler, who is very cranky due to having to wake up so early. She gets him ready, feeds him breakfast, and then they take the bus to his childcare provider. Her toddler cries as she drops him off, and Gloria feels very guilty. This feeling sticks with her as she rides three different buses to work, which takes two hours total. George, on the other hand, has no children, and lives close to the office. He obviously does not have to pay for supplemental childcare, and he gets up each day at 7am, gets ready, and then makes the 10 minute walk to work. The expectation that employees arrive at work on time imposes a *much* greater burden upon Gloria than it does upon George. But, it is not clear that it is thereby unfair, for this reason alone, especially since the employer has very good reasons for expecting her employees to be at work on time—her the success of her brokerage firm depends on this.⁸⁷ One might worry that if the burden is great enough, the expectation will seem unjustified. I think this is a version of the view that an expectation is unreasonable if it would be significantly difficult for an agent to meet it. As noted above, I am sympathetic to

consequences imposed for violations of the demand are not intended to prevent future violations of it, but instead have more insidious purposes. E.g. in the case under consideration, the employer imposes penalties for infractions of demands, but they do not do so with the intention of preventing future in fractions, but with the intention of forcing or “encouraging” some individuals to find different employment.

⁸⁷ This is not to suggest that society at large should not make an effort to make the daily lives of single working-parents easier (e.g. by providing high-quality, subsidized, local childcare, or by making public transportation more efficient). It is also not to suggest that it would not be good or kind of Gloria’s employer to be lenient with her if she is late.

the idea that the reasonableness of an expectation is sensitive to the difficulty agents would face in meeting it, though I do not think the difficulty of meeting an expectation is a condition on its reasonableness.

Finally, one might think that an expectation is unfair if it fails to treat like cases alike. It is not readily apparent how an expectation can be unfair for this reason, though it may create an unfair state of affairs. For example, I would treat my ten-year-old daughter unfairly by holding her to the demand that she clean her room if I do not also hold her twin sister to this demand, assuming I have no good reason for treating the children differently. However, while this is unfair, it is not clear that this is because it is unfair to hold my daughter to the demand that she clean her room. We can see this by the fact that I could remedy the unfairness of this situation by holding both girls to this demand. As will become apparent in the following section, I think expectations can fail to treat like cases alike, and thereby become unfair. Moreover, if a moral expectation is unreasonable, it will also be unfair in this sense. However, it is not the unfairness of the moral expectation that explains its unreasonableness.⁸⁸

4.4 Reasonable Moral Expectations: A New Approach

In the previous section, I argued that the capacity of the agent to whom an expectation is directed, the difficulty this agent would face in meeting the expectation, and the fairness of an expectation are not conditions on the reasonableness of an expectation. Therefore, the typical approaches to reasonable expectations in the literature do not offer an

⁸⁸ More specifically, I think unreasonable moral expectations fail to treat like cases alike, and are thereby unfair. So, if a moral expectation is unfair insofar as it fails to treat like cases alike, it is also unreasonable.

account of the conditions under which a moral expectation is reasonable. However, they do offer some guidance. As we saw, holding an agent to a demand, by threatening and imposing negative consequences aimed at preventing future infractions of it, is unjustified if pointless (i.e. if the agent lacks the capacity to ever bring it about, even with help from others, that she meet the expectation). Additionally, the reasonableness of an expectation must sometimes be sensitive to the difficulty an agent would face in meeting it, and, if moral expectations sometimes treat like cases unlike, they would be unfair and unreasonable due to this. An account of reasonable moral expectations should be able to accommodate and explain these claims.

In this section, I argue for an account of the conditions under which a moral expectation is reasonable or unreasonable. In particular, I argue that a moral expectation is reasonable if it meets one of the following individually sufficient conditions. It is reasonable if it requires an agent to (i) to refrain from acting in ways that treat other agents as lacking inherent or equal value, and (iii) it is not pointless, or if it requires an agent to (ii) to refrain from acting in ways that treat their interests as lacking or being of less relative importance compared to the interests of others and (iii) is not pointless. A moral expectation is unreasonable if it meets neither (i) nor (iii), or fails to meet (iii).

I develop this account by arguing that it is reasonable or unreasonable to expect things of others insofar as we stand in certain relationships with them. I examine two types of relationships in which we stand to others: “responsibility-based relationships,” and “participation-based relationships.” In responsibility-based relationships, we voluntarily assume responsibilities in virtue of entering the relationship. We thereby make it the case that it is reasonable for others to hold us to demands that we fulfill these responsibilities. In

participation-based relationships, by continuing to voluntarily participate in a relationship, we render it reasonable for others to hold us to demands which, if met, enable the relationship to continue to exist. I argue that the “moral relationship” is plausibly characterized as a participation-based relationship, and as a relationship of mutual regard. In virtue of our voluntary participation in this relationship, others can reasonably hold us to moral demands that we (i) or (ii) [that also meet (iii)], and cannot reasonably hold us to moral demands that require us to neither (i) nor (ii), or that do not meet (iii).

In section 4.5, I discuss possible objections to my account of the conditions under which a moral expectation is reasonable, and in particular, the objection that this account commits me to an overly demanding view of morality. In section 4.6 I discuss how the reasonableness of a moral expectation, on my view, is sometimes related to the culpability of an agent.

Reasonable Expectations and Relationships

The expectations we have of others are calibrated to the relationships in which we stand to them. For example, while I might expect my friend to console me after a difficult day, I certainly would not expect a stranger to do so. I expect my parents to make an effort to see me and my family, but do not expect distant cousins to do so. My children expect me to feed them dinner, but do not expect my neighbor to do so. Whether or not a particular expectation is reasonable also depends on the relationship that is its context. For example, it is reasonable of Sheila to expect her friend (qua friend) to offer her emotional support through a difficult illness. It is not reasonable of Sheila to expect her friend (all things equal) to donate his kidney to help her heal from the illness. Similarly, it is reasonable for Andy to

expect his spouse to perform his fair share of the housework. It is not reasonable for Andy to expect his spouse (all things equal) to throw over-the-top celebrations that far outreach their financial means, for his birthday each year.

We can begin to make sense of this phenomenon by considering *why* standing in a particular relationship with others makes it reasonable for them to expect certain things of us. In the most straight-forward cases, this is because in entering some relationships, we voluntarily assume a set of responsibilities. In doing so, we make it reasonable for others to hold us to demands that we fulfill these responsibilities. I'll call these responsibility-based relationships. Within the context of responsibility-based relationships, an expectation is reasonable if in holding it, we hold the agent in question to the demand that she fulfill the responsibilities she voluntarily assumed. An expectation is unreasonable if in holding it, we hold the agent to demands she fulfill responsibilities has not voluntarily assumed.

One example of a responsibility-based relationship is the one an instructor stands in with her students. When an instructor agrees to teach a class, she voluntarily assumes the responsibilities associated with doing so. She thereby makes it reasonable for others to expect that she fulfill these responsibilities. Imagine, e.g., that the instructor regularly fails to show up to teach (without a good explanation). Her students could reasonably demand that she show up. It would be appropriate, e.g., for them to rebuke her for her absences with the aim of preventing her future absences, "Get your act together! We are here, where are you? You can't keep skipping class." If she were to question the reasonableness of them holding her to this demand, it would be appropriate for them to respond, "But *you* chose to teach the class. You took on this responsibility!" It would be unreasonable of her students to expect her to fulfill responsibilities she did not voluntarily assume in virtue of agreeing to

teach the course. E.g., it would be unreasonable for one of her students to expect her (qua instructor) to listen to the details of his most recent break-up.

In other cases, our voluntary actions within the contexts of our relationships make it reasonable for others to expect certain things of us. In particular, in some cases our continued voluntary participation in the relationship makes it reasonable for others parties to expect us to meet what Shoemaker (2011) calls “relationship-defining demands,” or demands which, when met, “make the very existence of the relationship itself possible” (623). I’ll call these participation-based relationships. Within the context of participation-based relationships, an expectation is reasonable insofar as in holding it, we hold participants to the relationship-defining demands of relationship in question. An expectation is unreasonable if in holding it, we hold participants to non-relationship defining demands. Consider, for example, the relationship of friendship. Our friends can expect things from us that others cannot. This is plausibly explained by the fact that in voluntarily continuing to participate in the friendship, we make it reasonable for our friends to expect us to meet the relationship-defining demands of friendship.⁸⁹ Friendship is a relationship marked by

⁸⁹ We voluntarily continue to participate in a friendship in virtue of voluntarily engaging in the characteristic activities of friendship (e.g. seeking out the company of our friend, confiding in our friend, seeking advice and comfort from our friend, and so on).

In addition to being influenced by Shoemaker (2011), my thoughts here are heavily influenced by Jeske’s (1998) account of how personal relationships give rise to special obligations. Her view is that although most personal relationships are not cases in which we voluntarily assume responsibilities, these relationships do come about through our voluntary choices (e.g. the intimacy of friendship develops slowly, through a number of voluntary choices). Once a personal relationship is formed the parties now have a shared project—maintaining the relationship. This shared project is a source of obligations. She writes:

So after a friendship is established, each friend has a project that essentially involves the other and demands her continued participation. Whereas other projects sometimes create moral permissions, but not moral requirements, for agents to pursue the projects at the cost of the general interest, friendship creates obligations, types of moral requirements, to continue to care and to sustain intimacy, because of

intimacy and a resulting special, non-romantically and non-familial based regard.⁹⁰ The corresponding relationship-defining demands of friendship are that (a) we not act in ways that are inimical to this intimacy and regard (e.g. betraying confidences, reacting with indifference to exciting news), and (b) that we sometimes act in ways that support its continued existence (e.g. sharing personal details about our lives, emotions, and thoughts, , spending time together, and so on). In virtue of their continued voluntary participation in our friendship, we can reasonably hold our friends to demands that they (a) and/or (b). This is *all* that we can reasonably expect of our friends, qua friends. If an expectation falls outside the scope of (a) and (b), it will be unreasonable when addressed to friends, qua friends. For example, it would be unreasonable to expect that friends donate kidneys when we need them, loan us large sums of money, watch our children for free on numerous occasions, or quit their jobs to travel the globe with us. These demands fall outside the scope of (a) and (b); they demand that the agent do things that go beyond refraining from acting in ways that are inimical to the bonds of friendship, or in ways necessary for maintaining these bonds.

We stand in many different types of relationships with other human beings, some of which cannot be properly characterized as responsibility-based relationships or participation-based relationships. It is highly plausible that in some cases, the fact that other agents can reasonably expect things of us in virtue of our standing in a relationship with

its essentially shared character... while the voluntary nature of the friendship renders the demands placed on the parties in conformity with the requirement of voluntarism [the view that we only have special obligations or choices as the result of our voluntary choices], it is the shared character of the project of an intimate relationship that creates the demands to continue to care and to sustain the project (Jeske, 1998, p. 540).

⁹⁰ Nothing in particular hangs on this particular of this analysis of friendship being correct. What is important is that one accepts that participation-based relationships exist, and that they follow the above scheme.

them, will have nothing to do with our voluntary actions or assumptions of responsibility.⁹¹ However, we need not explore this possibility here. I think it is highly plausible that we can understand the “moral relationship”—the relationship in which we stand to other agents qua moral agents—as a practice-based relationship, and that we can glean a partial account of the conditions under which moral expectations are reasonable or unreasonable, by exploring its status as such.

We voluntarily participate in the moral relationship by voluntarily engaging in the practices of morality (i.e. blaming others, praising others, demanding fair or just treatment for ourselves and others, and so on). The moral relationship is plausibly construed as one of mutual regard; it is one in which all parties are recognized to be of inherent and equal value, and their interests to be of equal relative importance compared to the interests of other agents. The corresponding relationship-defining demand of the moral relationship is that agents maintain this regard by not acting in ways inimical to it—i.e. by refraining from actions that treat other agents as lacking inherent or equal value, or that treat their interests as lacking or being of less relative importance compared to the interests of others.⁹² In virtue

⁹¹ The relationship in which we stand to fellow citizens or to our government might be an example of such relationship.

⁹² One might wonder if we should not also include in the relationship-defining demand of the moral relationship that agents not hold beliefs or attitudes that are inconsistent with the equal standing of other moral agents. Isn't refraining from holding such beliefs and attitudes a way of respecting agents' equal moral standing? While the latter may be true, I do not think we can demand that others hold certain beliefs or attitudes, but only that they act so as to bring it about that they do so. Believing is not something we can do on demand; it makes no sense to demand that we do so (or to hold others to demands that they do so). We *can* act to bring it about that we believe something (e.g. by examining evidence, listening to other views, reflecting, and so on), and the relationship-defining demand of the moral relationship will sometimes require us to act so as to bring it about that we hold certain beliefs. Holding an attitude (e.g. caring, valuing), I think, often involves having a number of different beliefs and dispositions. For this reason, holding an attitude is also not something we can do on demand. Although, again, we can bring it about that we hold an attitude, and

of their continued participation in the moral relationship, we can reasonably expect other moral agents to meet this demand. This is all we can reasonably expect of other agents in virtue of their voluntary participation in it.

A natural question to ask here is whether or not the set of reasonable moral expectations is larger than the set of those made reasonable in virtue of an agent's continued voluntary participation in the moral relationship. I think it is not. I will not offer a full argument for this claim here, but I will explain my main motivation for believing this, which is two-fold. (1) I think it is highly plausible that other types of expectations (e.g. expectations we have of friends, landlords, employers, coworkers, and so on) are rendered reasonable or unreasonable almost exclusively in reference to the relationships that are their contexts.⁹³ It would be odd if moral expectations were notable exceptions to this rule. (2) To claim that it is reasonable to expect something of P is to claim that we're entitled to hold P to some demand, or if you prefer, that it is appropriate for us to do so. The question, then, is what else, besides P's continued voluntary participation in the moral relationship, could be the source of this entitlement or appropriateness. Though there are many other potential sources of this entitlement, I do not think any of them are plausible sources of it. Initial candidates include our inherent value (or the inherent value of others), the admirable content of the demand to which an agent is held, or the fact that rational agents would consent to being held to the demand in question. But, our entitlement (or lack thereof) to hold an agent to a demand often has nothing to do with the value of any of the agents who it concerns, or the admirable content of the demand it involves. For example, it is highly plausible that it is

sometimes the relationship-defining demand of the moral relationship will require that we do so.

⁹³ As noted previously, I think that a condition on any type of expectation being reasonable is that it is not pointless.

unreasonable to expect Billie Mac (from Chapter 3) to care about the human being who fell to his death and splattered on the ground below him. However, this does not plausibly have anything to do with the value of the agent issuing the demand, Billie Mac, or the deceased. Nor does it have anything to do with the content of the demand that he care (which is quite admirable). And, I at least hope that our entitlement to hold an agent to a moral demand does not depend on whether or not rational agents would consent to being held to it. It is sometimes very unclear whether or not this is the case—e.g. it is unclear whether or not rational agents would consent to being held to the demand that they do something they ought to do, but which is irrational from their own perspectives.

As I said, I recognize that this is not a sufficient argument for the claim that the set of reasonable moral expectations is comprised only of those rendered reasonable (almost exclusively) in virtue of an agent's voluntary participation in the moral relationship. Hopefully what I have said motivates this claim. If I am right about this, then we can shape an account of reasonable moral expectations based on the description of the moral relationship offered above.⁹⁴ The above description suggests two individually sufficient conditions for the reasonableness of a moral expectation. It suggests that a moral expectation will be reasonable if it requires an agent (i) to refrain from acting in ways that treat other agents as lacking inherent or equal value, or (ii) to refrain from acting in ways that treat their

⁹⁴ One might worry that if I am right, and we can reasonably expect things of moral agents only in virtue of their voluntary participation of the moral relationship, then we cannot reasonably hold agents to moral demands who are not voluntary participants of it. This may seem very unintuitive. But, we voluntarily participate in the moral relationship in virtue of engaging in the practices of morality—e.g. blaming, praising, demanding fair treatment, expressing our reactive attitudes. It will be *very* rare that a fully functioning adult human being does not voluntarily participate in any of these activities, and if he does not, it does not seem to me to be so odd to think that we cannot reasonably hold him to moral demands.

interests as lacking or being of less relative importance compared to the interests of others.⁹⁵

We must add to this account the condition on the reasonableness of any expectation, uncovered in the previous section: an expectation is reasonable only if (iii) it is not pointless. Combining (iii) with (i) and (ii) above, we arrive at the following account of the conditions under which a moral expectation is reasonable or unreasonable. A moral expectation is reasonable if it meets (i) and (iii) or if it meets (ii) and (iii). A moral expectation is unreasonable if it meets neither (i) nor (ii), or if it fails to meet (iii).

Although I do not think it is the case, it is possible that our voluntary participation in the moral relationship commits us to more than (i) and (ii), in which case it may be a source of further sufficient conditions for the reasonableness of a moral expectation that I have not considered here.⁹⁶ It is also possible that there are necessary conditions for the

⁹⁵ We can treat an agent as lacking inherent or equal value, or her interests as lacking or being of less relative importance in both active and passive ways. We do so actively if we treat her as if she has less dignity or worth than other moral agents or as being less deserving of certain treatment, or by actively assigning her interests less relative weight in our deliberations about what to do. We do so passively if we fail to allow an agent's inherent equal value, or the equal relative importance of her interests to act as appropriate limits on our actions. For example, Mr. S (discussed in Chapter 3) actively treats women as lacking equal dignity and worth as men; he treats them as financial assets and sources of free labor, rather than dignified, autonomous beings. A company would passively treat women as lacking equal value compared to men if they allow their managers to consistently offer much higher salaries to men than women, without justification. They would fail to allow the equal value of women to limit their actions. An agent would actively treat the interests others have in not being involved in a car accident as if they are of less relative importance than his interests in not angering his boss by being late, if he actively considers these interest while deliberating about whether or not to drive recklessly fast to work, and assigns the interests others have in not being involved in an accident less weight than his own interest in not being late to work. He would passively do so if he does not include their interests in his deliberations and drives recklessly fast; he would fail to allow their interests to limit his actions. On my view, an agent could be reasonably expected to refrain from all of these actions.

⁹⁶ And thereby further additions to the sufficient condition for the unreasonableness of a moral expectation.

reasonableness of expectations in general [like (iii)] that I have not considered. However, as I'll argue in section 4.6, even in its potentially incomplete state, my account can still be helpfully applied to the moral ignorance debate.

There are many concerns one might have with this account, which I will address shortly, in section 4.5. However, I do want to address one concern here. One might worry that this account suggests that many expectations are unreasonable which are not plausibly construed as such. E.g. one might worry that it suggests that expectations that an agent remember to clock out at the end of each shift, or that an agent remain seated while others toast him, are unreasonable insofar as they meet neither (i) nor (ii). But, these expectations are not obviously unreasonable. Fortunately, my account does not suggest that they *are*. My account applies only to *moral* expectations—i.e. to expectations addressed to moral agents qua moral agents, which are taken to be supported by distinctly moral reasons. The above expectations are not moral expectations; in order to determine if they are reasonable, we would need to consider them within the relationships that are their contexts.

Moreover, it is highly intuitive that most *moral* expectations that meet neither (i) nor (ii), are unreasonable. Moral expectations typically either require us to act in ways that treat others and their interests as being of inherent and equal value and importance, or in ways that treat them as lacking or being of unequal inherent value and importance.⁹⁷ If an

⁹⁷ For example, expectations that concern lying, stealing, unjustifiably harming others, and acting so as to avoid holding certain objectionable moral attitudes and beliefs are plausibly construed as doing so. I am aware that the very least, there will be cases that do not neatly fit this description. For example, imagine that while Peggy does not realize this, whenever she passes a black man on the street, she clutches her purse and scurries quickly past him. Greg, a black man who knows Peggy, holds her to the demand that she stop doing this. I think this is plausibly construed as a moral expectation. But, it does not obviously require Peggy to act in ways that treat others and their interests as being of inherent and equal value and importance, or in ways that treat them as being of unequal inherent value

expectation requires the former, it will be reasonable insofar as it meets (i) or (ii) and also (iii). If it requires the latter, it will be unreasonable because it will meet neither (i) nor (ii).⁹⁸

However, I take it that this is quite intuitive. Another way of putting this is that in most cases, moral expectations that meet neither (i) nor (ii) will require the opposite of either (i) or (ii).⁹⁹ And, expectations that require the opposite of (i) or (ii) are intuitively unreasonable moral expectations. To see what I have in mind here, consider the following example.

Felicity confides in Georgina that her husband is regularly physically abusive towards her and their children. Felicity expects Georgina to respect her privacy and to tell no one. It is plausible that this is a moral expectation, and it is one that meets neither (i) nor (ii). Instead, as I claim will most often be the case, it requires the opposite of (i) or (ii). It requires the opposite of (ii); it requires Georgina act in a way that treats the interests of some (i.e.

Felicity's interest in privacy) as if they are of greater relative importance than the interests of others (i.e. the interests her children have in being protected). This expectation is unreasonable on my view, and I take it that this is intuitively plausible. Something similar

and importance. However, I think Greg could be understood as a holding Peggy to the demand that she not have this fear; i.e. that she act so as to bring it about that she no longer has this fear (now that she is aware of it). She might do this, e.g., by educating herself and being cognizant of this fear, and being vigilant against it. Understood in these terms, I think Greg's expectation is one that requires Peggy (ii) to refrain from acting in a way that treat the interests of others as if they are of less relative importance compared to the interests of others. It is an expectation that she not act with complacency with respect to her fear, and thereby that she not treat the considerable interests others have in not being met with unjustified fear and suspicion as if they are of less relative importance than the (minimal) interests she has in not confronting and remedying this fear.

⁹⁸ If an expectation requires the opposite of (i) or (ii), it cannot also meet the other conditions [e.g. it cannot be that an expectation requires \sim (i), but meets (ii)].

⁹⁹ I leave open that there may be exceptions to this rule, though I do not think there will be many. If there are exceptions, I will be committed to saying that these expectations are unreasonable even if they are not intuitively so.

will be true with respect to most moral expectations that are unreasonable insofar as they meet neither (i) nor (ii).

As noted at the beginning of this section, a view of reasonable moral expectations should also be able to accommodate and explain the other lessons learned from my consideration of typical philosophical approaches to reasonable moral expectations: if an expectation fails to treat like cases alike, it is unfair and thereby unreasonable, and the reasonableness of a moral expectation is at least sometimes sensitive to the difficulty an agent would face in meeting the expectation. My account can accommodate and explain both lessons.

My account offers some insight into the way in which a moral expectation might fail to treat like case alike. If I am right, and the majority of the moral expectations that are unreasonable on my view, insofar as they meet neither (i) nor (ii), will also require the opposite of (i) or (ii), then many unreasonable moral expectations will also fail to treat like cases alike. An expectation that requires an agent to act in ways that treat some agents as if they lack inherent or equal value, or in ways that treat their interests as if they lack or are of unequal moral importance, will fail to treat like cases alike. Such expectations will not only be unreasonable because they extend beyond what can be reasonably expected of an agent, given her voluntary participation in the moral relationship [i.e. because they meet neither (i) nor (ii)], but also because they are unfair. However, in determining whether or not a moral expectation is unreasonable, we can focus exclusively on whether or not an expectation meets neither (i) nor (ii). It happens that most moral expectations that fall out of out of the scope of the relationship-defining demands of the moral relationship are also typically unfair

because they fail to treat like cases alike. But it is the latter, and not the unfairness of the expectation, that serves as the primary explanation of its unreasonableness.

My account can also explain why the reasonableness of a moral expectation is sometimes sensitive to the difficulty an agent faces in meeting it. The difficulty an agent faces in meeting an expectation sometimes impacts the reasonableness of the expectation, insofar as it sometimes makes it the case that the expectation will meet neither (i) nor (ii), and that it instead requires the opposite of (ii). In particular, it sometimes makes it the case that an expectation requires an agent to treat the interests of others as if they are of greater relative importance than her own. E.g. if my husband expects me to spend months traveling the Amazon rain forest in search of a rare plant that will eradicate the minor seasonal allergies he experiences, he holds me to a demand that meets neither (i) nor (ii), and which instead requires me to treat his interests as if they are of greater relative importance than my own. His expectation is thereby unreasonable. We can see this by considering how inappropriate it would be for him to reprove me if I refuse, “Don’t you care about my watery, itchy eyes? Your other obligations can wait.”

In other cases, the difficulty an agent faces in meeting a moral expectation does not affect its reasonableness in this way. Recall, for example, Samantha, who suffers from kleptomania. I argued that the shopkeeper’s expectation that she not steal from his store is not rendered unreasonable by the fact that it would be very difficult for her to meet this expectation. The shopkeeper has significant interests in Samantha not stealing from his store. E.g. he has interests in not losing inventory, in having his property rights respected, in not devoting time and energy to preventing or punishing her theft, and so on. These interests are plausibly of greater relative importance than the interests Samantha has *in* stealing—e.g.

it relieves the great tension she feels prior to stealing, and seeking treatment or avoiding situations that trigger her urge to steal might be costly and burdensome. Therefore, although it is very difficult for Samantha to meet the expectation that she not steal, the expectation is reasonable on my view; it requires Samantha to (ii) refrain from acting in ways that treat agents' interests as lacking or being of less relative importance compared to the interests of others.

4.5 Defending my account against objections

There are two main worries one might have about the account I have offered of the conditions under which a moral expectation is reasonable or unreasonable. (1) One might worry that it is not anymore helpful than an intuitive approach to what can be reasonably expected of a moral agent. And, (2) one might worry that my account of reasonable moral expectations commits me to an overly demanding, impartial view of morality. In response to (1), I argue that my account is no more difficult to apply than an intuitive approach to reasonable moral expectations. In response to (2), I argue that my account does not commit me to an overly demanding, impartial view of morality, and that even if it does involve some revision to ordinary moral views, these revisions are not as drastic as they may at first seem.

(1)

One might worry that my account of reasonable moral expectations is not any more helpful than an intuitive approach to reasonable moral expectations. In the last two chapters, I argued that, in part, we need a general account of reasonable moral expectations, due to the fact that intuitions are often divided regarding the claim that an expectation is unreasonable. E.g. some find the claim that agents cannot be reasonably expected to do something they see

no reason to do, to be highly plausible, while others do not. My hope has been that a general account of the conditions under which a moral expectation is reasonable or unreasonable would settle debates like this. However, one might think that my account is just as difficult to apply. In particular, one might think that it is difficult to determine whether or not an expectation meets the second disjunct of (ii), which holds that an expectation is reasonable if it requires an agent to refrain from acting in ways that treat the interest of some agents as if they are of unequal relative importance compared to the interests of others. Determining this requires one to determine the general importance of the interests at stake, and whether or not one agent's interests are, in fact, of greater relative importance than another agent's interests. Our judgment about the latter may be no clearer than our intuitions about particular cases.¹⁰⁰

In fact, I think our judgments about the relative importance of interests are clearer. We are highly practiced at weighing interests. In ordering our own lives, we regularly weigh our own interests against each other, as well as against the interests of our loved ones, acquaintances, and strangers. While it may be difficult to determine the relative importance of interests in some cases, this is not typically a difficult task. Moreover, our judgments regarding the relative importance of interests are not clouded in the ways that our intuitions regarding the reasonableness of moral expectations are clouded. The latter are often muddled by our intuitions concerning whether or not an agent is to blame for the act in question, or whether or not he had a moral obligation to perform it. The former are not clouded by these concerns.

¹⁰⁰ I thank Matt King for raising this worry at the 2015 meeting of the Alabama Philosophical Society.

(2)

Perhaps of greater concern is the worry that according to my view, and in particular, according to the second disjunct of (ii), we can be reasonably expected to refrain from acting in ways that treat the interests of some as if they are of unequal relative importance compared to the interests of others. One might worry that this entails that it is reasonable to expect us to be impartial towards the interests of others, and even towards our own interests. But, this seems to entail a view of morality that is too demanding. However, I do not think my account entails this view.

Let's first consider the worry that the second disjunct of (ii) entails that we can be reasonably expected to refrain from treating the interests of others partially. I think this worry is misguided, and arises out of the mistaken thought that in treating one agent's interests partially *for a good reason*, we thereby treat her interests as if they are of greater relative importance than the interests of others, or in the mistaken belief that reasons for acting partially must be grounded in the supposed greater importance of the interests of the individual to whom we're partial. However, neither is true. In virtue of treating their interests partially, we treat the interests of some as being of greater relative importance than the interests of others, only if we do so without good reason, or we intentionally harm the interests of others so as to benefit the interests of those to whom we are partial. Consider the following example. I worry about and tend to my daughter's nutrition and emotional wellbeing. I do not (generally) worry about and tend to the nutrition and emotional wellbeing of other children. I have good reasons for this. My daughter can reasonably expect me to do this, given our relationship of mother to young child. Doing so is an integral expression of my love for her. Doing so promotes my own interests; I am emotionally

rewarded by her doing well (because I love her and am emotionally invested in her wellbeing), and I have an interest in this monumental project of parenthood going well. In light of these reasons, I think it is implausible that in treating her interests partially, I thereby treat them as if they are of greater relative importance than the interests of other children. I *would* plausibly do so if I intentionally harmed other children in order to promote the interests of my daughter (e.g. stole their food to give to my daughter), or if I promoted her interests but not theirs without legitimate reasons for doing so. Moreover, notice that my reasons for treating my daughter partially are not grounded in the supposed greater importance of her interests. Of course it is more important *to me* that she do well; I am personally invested in this. But, this does not entail that I hold that her interests in doing well are objectively more important than the interests other children have in doing well. If treating an agent's interests partially, with good reason (and without intentionally harming others in order to do so), does not involve treating their interests as if they are of greater relative importance compared to the interests of others, then the second disjunct of (ii) does not entail that we can be reasonably expected to refrain from doing this.

This same line of reasoning can be used to address the worry that the second disjunct of (ii) requires that we not act in ways that treat our own interests partially. All the second disjunct of (ii) requires is that we not treat the interests of some agents partially *without good reason*, and that we not intentionally harm the interests of some to promote the interests of others to whom we are partial. This does not entail that we can be reasonably expected to act in ways that treat our own interests impartially. We arguably often have good reasons for treating our interests as such; at the very least, doing so is often integral to the success of our projects, and something in which we are deeply emotionally invested.

However, one might worry that we often preference our interests and the interests of others without good reason, and do not typically think we thereby act in a way that others could reasonably demand we not. E.g. I buy my children pretty shoes and dresses that they clearly do not need, even though I know that many children in the Lake Chad region of Nigeria are on the brink of starvation, and my money could instead be used to help save them. Many of us engage in similarly frivolously purchases, with the same knowledge that our money could be put to (*much*) better use. In doing so, aren't we treating the interests of some partially, without good (or at least sufficiently good) reason? Doesn't this suggest that we often, many times a day even, act in ways that, on my account (by the second disjunct of (ii)), other agents could reasonably expect we do not? Isn't this at odds with ordinary moral thinking?

Fully addressing this worry is difficult in that it requires an account what constitutes a good reason for being partial. I have offered some examples of potential good reasons for being partial—e.g. that others could reasonably expect us to be partial to particular interest of theirs in virtue of another relationship in which we stand with them, that being partial is an integral expression of our love, that being partial is integral to the success of projects that are important to us. There may be other good reasons for being partial as well. But, even if my account suggests that we can be reasonably expected not to, e.g., make utterly frivolous purchases when we could instead devote that money to save another agent's life, I do not think this involves as radical a revision to our ordinary moral thinking as it may at first seem. Our resistance to the claim that we could be reasonably expected to refrain from such actions is rooted, I think, in the mistaken assumption that the mere fact that we can be

reasonably expected to X entails that we have an obligation to X, or that we can be blamed for failing to X. But neither is true.

In the following section I will argue that in many cases, blaming an agent for X is justified only if the agent could have been reasonably expected to avoid X at the time she performed it. However, that an agent could have been reasonably expected to X is not a sufficient condition for the justification of blame. To claim that it is reasonable to expect an agent to X means that it is appropriate for other agents to hold him to the demand that he X, by guarding against his infractions of X through threatening or imposing negative consequences aimed at doing so. Nothing about this claim entails that the agent is blameworthy for failing to do what he could have been reasonably expected to do. In order for an agent to be culpable for failing to X other things must also be true. E.g. it is plausible that it must also be true that in failing to X the agent in question also wronged another person, failed to fulfill his moral obligations, or that his actions manifested a character vice or objectionable moral attitude. None of these conditions are necessarily met in virtue of an agent failing to do what he could be reasonably expected to do.¹⁰¹ For example, I think it is highly plausible that a starving child could reasonably hold me to the demand that I not buy a \$4 coffee drink and instead offer him life-saving aid. It would be appropriate, e.g., for him to reprove me for buying the coffee, in the hopes of preventing my future infractions of his demand. “You’re choosing a latte over the life of a child? Please donate to Oxfam instead. Surely my life is more important than a coffee!” However, it does not follow from this that I am blameworthy for instead buying the latte if doing so does not manifest a character vice (e.g. greed, callousness or selfishness), does not violate any moral obligations I have to help

¹⁰¹ They often *are* met, insofar as agents can reasonably expect us to fulfill our moral obligations and to refrain from harming them.

the child, does not otherwise wrong the child, and does not manifest an objectionable moral attitude (e.g. a lack of respect for the value of his life).

It is also not the case that the fact that we could be reasonably expected to X entails that we have an obligation to X. Again, to claim that we could be reasonably expected to X, is to claim that it would be appropriate for other agents to hold us to the demand that we X, by imposing and threatening negative consequences aimed at preventing our future infractions of this demand. Whatever one's view is about the origins of our moral obligations, it is hopefully clear that we do not have them simply insofar as other moral agents can reasonably expect certain things of us in virtue of the moral relationship in which we stand to each other.

Once we realize that the fact that an agent can reasonably expect us to X does not entail that we are obligated to X, or that we are culpable for failing to X, the idea that we can be reasonably expected to not treat our own interests or the interests of other partially, without good reasons, is more palatable. Most of us do think, e.g., that even though we do not violate our obligations or do something blameworthy when we spend money on utterly frivolous items with the knowledge that the money could be used for much better causes, there is something morally "off" about doing so. My account can explain this "offness." In doing so, even though we don't violate our obligations or act in ways that are blameworthy, we act in ways that other moral agents could reasonably expect we do not.

4.6 Reasonable Moral Expectations and Culpability

Now that I have an account of the conditions under which a moral expectation is reasonable or unreasonable, I can explain how the reasonableness of a moral expectation is

sometimes related to whether or not we are justified in blaming an agent for failing to meet it. As I argued in section 4.4, most moral expectations are unreasonable insofar as they meet neither (i) nor (ii). Such expectations most often also require the opposite of (i) or (ii), and thereby fail to treat like cases alike. When an expectation is unreasonable in this sense—i.e. when it meets neither (i) nor (ii) and treats like cases unlike—we would be unjustified in blaming an agent for failing to meet it. To put this in slightly different terms, if it would have been unreasonable, in this sense, to expect an agent to avoid X at the time she performed X, then we would be unjustified in blaming her for X. This is because when we blame an agent in virtue of her having failed to meet some demand, we at least implicitly endorse this demand. E.g. when I blame Fred for lying to me, I implicitly endorse the demand that he not lie. But if a demand fails to treat like cases alike, in endorsing it—in blaming the agent for failing to meet it—we fail to treat like cases alike. If I am right that most moral expectations are unreasonable insofar as they meet neither (i) nor (ii), and I am right that in virtue of this, most expectations fail to treat like cases alike, then in most cases, we would be unjustified in blaming an agent for failing to X if she could not have been reasonably expected to avoid X, because in doing so, we would fail to treat like cases alike.

Although this explanation of the relationship between the reasonableness of an expectation and the justification of blame is a bit clunky, I think it is fairly intuitive. We often think that we are unjustified in blaming an agent for that which she could not have been reasonably expected to avoid, insofar as in doing so, we either endorse an unfair moral demand, or treat the agent herself unfairly. On my view, in blaming an agent for failing to meet a demand that would have required her to treat *others* or their interests as if they lack or are of unequal value or relative importance, we endorse an unfair moral demand. In

blaming an agent for failing to meet a demand that would have required her to treat *herself* as if she lacks or is of unequal value, or her interests as if they lack or are of less relative importance, we treat *her* unfairly. In support of this, consider the following examples. Recall Felicity and Georgina. Felicity expects Georgina to refrain from reporting to others that Felicity and her children are being physically abused by Felicity's husband. It is unreasonable to expect Georgina to refrain from reporting this; this expectation neither requires Georgina (i) to refrain from acting in ways that treat other agents as lacking inherent or equal value, nor (ii) to refrain from acting in ways that treat their interests as lacking or being of less relative importance compared to the interests of others. Moreover, this expectation requires Georgina to do the opposite of the second disjunct of (ii). In doing so, it fails to treat like cases alike. If Georgina were to report the abuse (to the proper authorities), we would not be justified in blaming her for doing so. In blaming her, we would implicitly endorse the demand that she not do so; we would endorse an unfair demand. This seems to me to be in line with our intuitions. If someone blamed Georgina for reporting the abuse, we would like speak up by pointing out that she was right to have done so, that she had to protect the children. In pointing this out, I think we would be pointing out that the demand that Georgina not report the abuse misrepresents the balance of interests at stake; it misrepresents the interests Felicity has in not having the abuse reported as being of greater importance than the children's interests in being protected. It thereby treats the children unfairly, and is not a demand that we should implicitly endorse, via blaming Georgina for failing to meet it.

Alternatively, imagine that Sally has promised Fred that she will be at his party tonight. However, while getting dressed for his party, Sally trips, falls, and breaks her ankle.

She is in intense pain, and needs medical care. Her only option is to go to the emergency room, where she'll have to wait many hours, and therefore will miss Fred's party. Sally regrets breaking her promise, but goes to the E.R. We could not have reasonably expected Sally to keep her promise to Fred. This expectation would meet neither (i) nor (ii), and would require the opposite of the second disjunct of (ii); it would require Sally to treat Fred's interests in her attending the party/keeping her promise, as if they are of greater relative importance than her own interests in being healthy, and free from pain. In blaming Sally for failing to keep her promise to Fred, we endorse the expectation that she do so. However, in endorsing this unreasonable moral expectation, *we* treat Fred's interests as if they are of more weight than Sally's. We thereby treat Sally unfairly. Again, I think this is in line with our intuitions. Were someone to blame Sally for failing to keep her promise, we would likely speak up by pointing out that it is not fair to expect her to keep a promise if nothing very important depends on her doing so, and doing so would require her to endure intense physical pain and possibly exacerbate her injury.

4.7 Reasonable Moral Expectations, Culpability, and the Moral Ignorance Debate

I am finally in a position to return to the discussion of the moral ignorance debate. In this section, I apply my partial account of reasonable moral expectations first to the skeptical argument, and then to the culpability views. In this section, I apply my view of reasonable moral expectations first to the skeptical argument, and then to the culpability views. I argue that it suggests we ought to reject the skeptical argument, and accept only revised versions of culpability views.

Reasonable Moral Expectations and the Skeptical Argument

The success of the skeptical argument depends on the defensibility of its premise that agents are only culpable for their moral ignorance if it is the result of a knowing mismanagement of their moral beliefs. Proponents of the skeptical argument depend on this premise in order to reach the conclusion that agents are almost never culpable for their moral ignorance insofar as they very rarely knowingly mismanage their moral beliefs. In order to defend this premise, they must defend the further claim that agents can only be reasonably expected to avoid their moral ignorance if it is the result of a knowing mismanagement of their moral beliefs. However, using my account of reasonable moral expectations, I argue that this further claim is mistaken. The skeptical argument is therefore indefensible.

In Chapter 2, I discussed two objections to the claim that agents are only culpable for their moral ignorance, if it is the result of a knowing mismanagement of their moral beliefs. One is offered by Smith (2015), and the other by FitzPatrick (2008). Smith (2015) argues that we need not trace culpability for objectionable moral attitudes (like moral ignorance) to culpability for some past actions. As I explained in Chapter 2, her view is that,

We trace only in order to “find a judgment-sensitive locus of responsibility that links the agent to the thing for which he is being morally criticized” (Smith, 2015, pg. 124). E.g. in the case of a drunk driver, we hold the agent to be culpable because her action of getting drunk in the first place (without a plan to get home safely) reflects an objectionable attitude. Smith (2015) argues that this justification for tracing does not apply to attitudes. Attitudes, “*by their very nature*, are judgment-sensitive states, so there is no need to try to ‘trace back’ to some *other* judgment-sensitive locus of responsibility” (124) (Zinn, 2016, p. 34).

Because moral ignorance constitutes an objectionable moral attitude, on Smith’s view, we need not trace culpability for it to some past action; we can be directly culpable for it.

In chapter 2, I argued that proponents of the skeptical argument might respond to Smith by offering a different account of tracing, according to which we trace culpability for moral ignorance to culpability for some past action in order to determine whether or not an agent could have been reasonably expected to avoid her moral ignorance. However, in order for this response to act as a defense of their claim that agents are only culpable for their moral ignorance if it is the result of a knowing mismanagement of their moral beliefs, they must also argue that agents can only be reasonably expected to avoid their moral ignorance if the latter is true.

FitzPatrick (2008) argues that the claim that agents are only culpable for their moral ignorance if it is the result of a knowing mismanagement of their moral beliefs is too strong. Instead, he suggests, moral ignorance is culpable if (and only if) the agent could have been reasonably expected to avoid it. He argues that there is no reason to endorse the stronger claim, and that his alternate view explains our intuitive reactions to particular cases.

The success of this objection also depends on the truth of the claim that agents can only be reasonably expected to avoid their moral ignorance, if it is the result of a knowing mismanagement of their moral beliefs. If the former is true, then FitzPatrick's (2008) claim that agents are culpable for their moral ignorance if they can be reasonably expected to avoid it, does not challenge the claim that agents are culpable for their moral ignorance only if it is the result of a knowing mismanagement of their moral beliefs.

Rosen (2004) and Levy (2009) argue that agents cannot be reasonably expected to avoid their moral ignorance, unless it is the result of a knowing mismanagement of their

moral beliefs.¹⁰² They hold that as long as an agent has not knowingly mismanaged her moral beliefs she does not do anything she believes to be morally wrong when managing her moral beliefs, nor does she fail to do anything she believes she ought to do with respect to them. Instead, she manages her moral beliefs in exactly the way she believes she ought. In order to avoid developing her moral ignorance, she would either have to act against her own best judgment, or she would have to do something she sees no reason to do. It is unreasonable to expect an agent to act in either manner, and therefore it is unreasonable to expect an agent to avoid her moral ignorance, if she developed it unwittingly.

An expectation that an agent take steps to remedy his moral ignorance, or to avoid it in the first place, is plausibly construed as a moral expectation. E.g. the expectation that an agent stop and consider counter-evidence, not make a hasty-generalization, or refrain from being overconfident *in order to bring it about that she has true moral beliefs or remedies false ones*, is plausibly construed as an expectation issued to a moral agent qua moral agent, and one which is taken to be supported by distinctly moral reasons. In order to determine if Rosen and Levy are right that it is always unreasonable to expect agents to act in ways through which they could remedy or avoid their moral ignorance, but which appear irrational from their perspectives, we need to determine if it is always the case that such expectations meet neither (i) nor (ii), or that they never meet (iii). It is not plausible that such expectations never meet (iii). Agents who are morally ignorant are in the grip of a false moral view, and may see no reason to act so as to remedy this fact, but it is not the case that they lack the capacity (at any relevant time) to ever bring it about (even with the help of others) that they hold true moral beliefs (or at least lack false ones). It is also not plausible

¹⁰² See Rosen (2004), pgs. 602-603. Levy (2009) argues for these claims at length, in response to FitzPatrick's (2008) argument.

that it is always the case that such expectations meet neither (i) nor (ii). In fact, I think in many cases, they will meet the second disjunct of (ii); they will require an agent to refrain from acting in ways that treat his interests as if they are of greater relative importance than the interests of others.

To see this, consider the following example. Bill owns and runs a plant where employees are regularly exposed to chemicals known to cause serious forms of cancer. He does only what the law requires him to do in order to protect his employees from exposure to these chemicals, even though he has good reason to believe these measures are inadequate and he could easily expand them without much expense. He sees no problem with his actions. On his view, as a business owner, his only responsibility is to turn as high a profit as possible, and to obey the law. It is the government's responsibility to protect workers. Bill is morally ignorant, but let's stipulate that this is not because he has knowingly mismanaged his moral beliefs. He is aware that others disagree with his views, but thinks that they are too soft, and simply cannot understand what it is like to own a successful business. He was raised with these values by his parents, who owned his company before him, and they have been reinforced by other business owners he knows. It has never occurred to him to doubt their judgment, especially since no one else he knows and respects, questions it.

According to Levy's and Rosen's view, it would have been unreasonable to expect Bill to avoid his moral ignorance. In order to do so, he would have had to act in a manner that was irrational from his own perspective, or for no reason at all. For example, he may have avoided his moral ignorance if he had seriously questioned the positions of his father or his own positions once they were developed. However he saw no reason to do so. Perhaps this, in turn, is due to overconfidence on his part, and could have been avoided if he had

doubted his own ability to evaluate moral claims. However, he did not know that he was overconfident. In order to correct this fault, Bill would have had to question his own moral reasoning abilities, despite judging that there is no reason to do so.

Is it true that we could not have reasonably expected Bill to question the views of his parents, or his own positions once they were developed (or even his own abilities), given the fact that he saw no reason to do so? I think it is not. In failing to question these views, especially in light of the fact that he was aware that others disagreed with them, Bill is guilty of moral complacency. He adopted the dominant moral view of his community because he saw good reason to do so, and no good reason not to do so. He then settled into this view, and refrained from subjecting it to serious scrutiny. Instead, he used it to dismiss the views of its opponents as overly soft or sentimental. I think we could have reasonably expected him not to do this. We could have reasonably expected him to question his adopted moral beliefs and to scrutinize them in response to disagreement, *even though he saw no reason to do so*. That is, we could have reasonably expected him to have done so as standard practice, not because he saw any particular reason to do so, but because, in order to guard against the possibility of holding false moral beliefs (and being morally complacent), we have to reflect on our own moral beliefs from time to time, and scrutinize them in response to conflicting views that are not based on obviously terrible premises.¹⁰³ We typically can reasonably

¹⁰³ By “obviously terrible” I have in mind views that largely depart from our shared sense of morality—e.g. views that it is permissible to round up and murder large groups of individuals in order to “purify” a particular race. Of course, we could offer a more refined view of when agents act complacently, but I think it is at least plausible that they do so when they fail to scrutinize their own moral views in response to conflicting views held by segments of their own communities, or when they unreflectively adopt the dominant moral view of their community. I take it that complacency does not involve a knowing negligence or wrongdoing. So, in expecting agents to avoid complacency, we expect them to act in ways that seem irrational from their own perspectives.

expect this of any moral agent. We certainly have interests in being complacent. It is easier than regularly scrutinizing our moral beliefs in response to disagreement. Thinking about morality takes time and effort. There is comfort in complacency. There may be an uncomfortable cognitive dissonance in doing something we believe we either have no reason to do, or have good reason not to do (e.g. like seriously considering a view we believe to be false). But, I think it is plausible that in many cases these interests are not of greater relative importance than the interests others have in us refraining from moral complacency, and thereby avoiding or remedying our moral ignorance.¹⁰⁴ The expectation that others not be complacent with respect to their moral beliefs is therefore often reasonable; it often meets the second disjunct of (ii) by requiring that agents refrain from acting in a way that treats the interests of others as if they are of less relative importance than their own.

As further support for this argument, notice that we regularly expect that agents avoid complacency in other arenas as well. Imagine, for example, that Eleanor has been an airplane mechanic for 25 years. She has performed hundreds of thousands of pre-departure safety checks on aircrafts. She has never made a mistake, and by this point in her career, she is so familiar with airplanes that she is 100% confident in her ability to spot a safety issue by quickly glancing a plane over. She sees no reason to painstakingly check each wire connection, or to look at each and every bolt on the wing of the plane. However, I think very few would agree that it would thereby be unreasonable to expect her to do these things (as

¹⁰⁴ I grant that some agents may have significant interests that are served by being morally complacent. E.g. agents who would face serious alienation, hostility or aggression if they were to reject the dominant moral view of their community plausibly have serious interests in being morally complacent. Whether or not it is reasonable to expect such agents to refrain from being complacent will depend on what interests others have in them avoiding or remedying their false moral beliefs.

part of a pre-departure safety check).¹⁰⁵ Even if she does not believe it possible, there is a chance she could miss something this time. Others have serious interests in her not missing anything, and no significant interest of Eleanor's is served by not checking the plane carefully. We could reasonably expect her to do it; we could reasonably expect her not to be complacent with respect to our safety.¹⁰⁶ Similarly, no matter how confident we are that our moral beliefs are accurate, there is a chance we could be wrong (any cursory exploration of human history indicates that we often are). We often have no serious interests in being morally complacent. Others often have serious interests in us holding true moral beliefs. We can therefore be reasonably expected to avoid moral complacency; we can be reasonably

¹⁰⁵ One might argue that this isn't a moral expectation, but rather a role-expectation. However, I think it is plausible that we could expect Eleanor to perform the safety check carefully either based on her role, or simply as a moral agent. Moreover, we could swap this example with one of an agent who is 100% confident in his ability to drive safely while texting. I unfortunately know someone who suffers from this delusion. He is aware that research from reputable sources indicates that individuals who text while driving are much more likely to cause an accident than those who do not. He believes he is an exception to this rule (and has some evidence to support this). He therefore sees no reason to refrain from texting while driving, and has many reasons to text while driving. But, it would be implausible to suggest that we could not reasonably expect him to refrain from doing so because this is irrational from his perspective.

¹⁰⁶ One might be tempted to think that we could reasonably expect Eleanor to carefully check the plane only if she were aware, on some level, that it was possible that she might make a mistake. Only if she recognized this could it be said that she should have known better, and that she should have carefully checked the plane. I take the reasoning here to be that if she is aware that she could miss something, then it is not in fact internally irrational for her to carefully check the plane; at some level, even *she* knows she ought to do so. I see no reason to grant that we could reasonably expect Eleanor to carefully check the plane only if she is aware, on some level, that it was possible that she might make a mistake in glancing it over. But, notice that even if we do grant this, and then apply this lesson to the moral ignorance cases, we would still have an argument against Levy and Rosen. We would not be able to argue that they are wrong in holding that agents cannot be reasonably expected to do what would be irrational from their own perspective, but we could argue that it is extremely rare for it to be the case that it is internally irrational for agents to avoid moral complacency insofar as it is very rare (does it ever happen?) that agents are not aware that it is *possible* that their moral beliefs are mistaken.

expected to reflect upon and question our moral beliefs, and to scrutinize them in response to disagreement, even if we see no reason to do so.

One might worry that in claiming that we can be reasonably expected to do this, I am claiming that we can all be reasonably expected to be moral philosophers. This would be too demanding. But, I am claiming no such thing. An agent does not have to be a moral philosopher in order to seriously question and scrutinize his moral beliefs. In fact, I think most people do a fairly good job of this. E.g. they might not see any reason to believe that there is something wrong about buying antibacterial hand soap, but they are willing to listen to, and really consider why someone else does (and perhaps even change their mind). It may have never occurred to them that it could be immoral to buy wine that is grown in California, given the State's water crisis, and indeed they might find the idea preposterous, but they don't dismiss the claim outright. They listen, and think about it, and reconsider their own views. We don't need to be moral philosophers to do these things; we just have to be reflective and open-minded.¹⁰⁷

I also do not mean to suggest that all moral ignorance is due to moral complacency, or that by reflecting on our moral views, and scrutinizing them in response to criticism (even if we see no reason to do so), we will always arrive at true moral beliefs (or remedy false ones). I only mean to argue against Levy and Rosen's view that we can never reasonably expect agents to take actions through which they could remedy or avoid their moral ignorance, if such actions would be irrational from their own perspectives.

¹⁰⁷ One might worry more about the claim that agents need to question their core moral beliefs. I don't think we must necessarily do this directly; by questioning the implications of our core moral beliefs, we question these core beliefs themselves.

If I am right, then proponents of the skeptical argument cannot successfully defend the claim that agents are culpable for their moral ignorance only if it is the result of a knowing mismanagement of their moral beliefs against Smith's and FitzPatrick's objections. It is implausible that we must trace culpability for moral ignorance to a past knowing mismanagement of her moral beliefs, in order to determine whether or not the agent could have been reasonably expected to avoid it. This doesn't establish that Smith (2015) is right, and that we need not trace at all in order to determine if an agent is culpable for her moral ignorance.¹⁰⁸ But, so far there is no reason to think that we must trace culpability for it to a past, knowing mismanagement of one's moral beliefs. And, without the claim that agents can only be reasonably expected to avoid their moral ignorance if it is the result of a knowing mismanagement of their moral beliefs, they cannot defend their view against FitzPatrick's suggestion that the weaker requirement that agents are culpable for their moral ignorance if they could have been reasonably expected to avoid it, is preferable.

Given the fact that proponents of the skeptical argument cannot defend the premise that agents are only culpable for their moral ignorance if it is the result of a knowing mismanagement of their moral beliefs, and the fact that this premise leads to the highly implausible conclusion that agents are very rarely culpable for their moral ignorance, we ought to reject it. We ought to reject the skeptical argument.

One might think that proponents of the skeptical argument could defend their premise on other grounds. One might think that because it is true that agents are only

¹⁰⁸I discuss Smith's non-volitional view in the following section. While nothing I have to say in the dissertation confirms the non-volitionalists' claim that agents are directly culpable for their moral ignorance, nothing disconfirms it either. For my part, I find it highly plausible.

culpable for their *factual* ignorance, if it is a result of a knowing mismanagement of their factual beliefs, this is true with respect to moral ignorance as well. Alternatively one might think that if moral ignorance cannot be traced to a knowing mismanagement of one's moral beliefs, we will face an infinite regress, and be unable to determine if the agent is culpable for her moral ignorance. Neither claim is plausible.

It is not plausible that we must trace culpability for factual ignorance to a *knowing* mismanagement of one's moral beliefs. Imagine, for example, that Greg does not know that the substance he is adding to the soup is rat poison, and not salt. He stores both the rat poison and the salt in clear plastic containers near the stove. In doing so, he acts negligently, and is thereby culpable for the fact that he does not know he is adding rat poison to the soup. It need not also be the case that he was aware that he ought not to store the poison next to the salt, in an indistinguishable container. It is plausibly enough that he should have known this, i.e. that he could have been reasonably expected to know this.

It is also not plausible that we face an infinite regress if an agents' culpability for their moral ignorance is not traced to a knowing mismanagement of their moral beliefs. Proponents of the skeptical argument argue for this claim in the following manner. Culpability for moral ignorance must be traced to culpability for some past act or omission. Moral ignorance is either the result of a knowing or unknowing mismanagement of one's moral beliefs. Because agents are only culpable for their acts/omissions performed from moral ignorance if they are also culpable for their ignorance, agents are only culpable for an unknowing mismanagement of their moral beliefs if they are also culpable for *this* bit of ignorance. So, culpability for moral ignorance can only be traced to culpability for some past, unknowing mismanagement of one's moral beliefs, if one is culpable for the ignorance

from which she mismanaged her beliefs. Otherwise, the mismanagement would not be culpable. In order to be culpable for the ignorance from which she mismanaged her moral beliefs, it must be traced to a past, culpable act or omission, i.e. a knowing mismanagement of one's moral beliefs. If it is instead the result of an unwitting mismanagement, in order to determine if this mismanagement is culpable, we must determine if the agent is culpable for this further bit of ignorance. This pattern continues unless the agent's original moral ignorance "bottoms out" in a knowing mismanagement of her moral beliefs.

This argument rests on two main claims. (1) Agents are only indirectly culpable for their moral ignorance; they are only culpable for it by way of being culpable for something else. (2) Agents are only culpable for actions performed from moral ignorance if they are also culpable for their ignorance. Arguing against (1) would require us to take a position concerning whether or not agents are directly or indirectly culpable for their moral beliefs, which requires us to take a position concerning responsibility for beliefs in general. This would take us far afield from the discussion of culpability for moral ignorance. Fortunately, we need not tackle this project. If (2) is implausible, this would show that the above argument is mistaken.

Many simply assume that (2) is true.¹⁰⁹ However, I am not convinced this is right. Zimmerman (1997, 2008) and Rosen (2004) offer a defense of (2). But, neither defense is

¹⁰⁹ The only exception of which I am aware is Harman (2011, 2014). Harman's (2011) denies "that that the blameworthiness for the wrong action *is derivative* from the blameworthiness for the belief," but grants that "the action is blameworthy *only if* the belief is blameworthy as well" (459). Actions performed from culpable moral ignorance are culpable for the same reason that moral ignorance is culpable: they constitute a failure to fulfill our obligation to care sufficiently about something of moral significance. Non-culpable moral ignorance (i.e. that rooted in non-culpable factual ignorance) does not reflect this, and neither do actions performed from non-culpable factual ignorance. Therefore, "While both the beliefs and the actions are blameworthy, the actions are not blameworthy

successful, and no other defense of it is clearly available. Zimmerman (1997, 2008 Ch. 4) holds that agents are only culpable for their actions performed from moral ignorance if they are also culpable for their ignorance because, lack of ignorance is a “root requirement” for responsibility (424, 177). He supports the latter claim by pointing to cases of factual ignorance, in which we judge that agents are only culpable for their action performed from ignorance if they are also culpable for their ignorance. Rosen (2004) defends (2) by arguing that it would be unreasonable to expect an agent to act in a way that she (blamelessly) believes to be acceptable, and therefore that agents are not culpable for their actions performed from moral ignorance unless they are also culpable for their ignorance (302-303).

While it is true that we hold that agents are not culpable for their actions performed from factual ignorance, unless they are also culpable for their ignorance, this does not ground the claim that lack of (culpable) ignorance is a root requirement for responsibility. There are many possible reasons we hold the former to be true. For example, we may hold this to be true insofar as, if an agent is not culpable for her factual ignorance, her action performed from factual ignorance (all things equal) does not reflect an objectionable moral attitude. E.g. If Frank non-culpably believes you are not allergic to anchovies, his act of serving them to you in homemade salad dressing does not (all things equal) reflect an objectionable moral attitude. If this is the explanation of the fact that agents are only culpable for actions performed from moral ignorance, if they are culpable for their moral ignorance, then not only does it not support the claim that lack of culpable ignorance is a root requirement for culpability (lack of an objectionable moral attitude would be), it also does not extend to moral ignorance.

because the beliefs are blameworthy. Rather, the actions and the beliefs are blameworthy for similar reasons” (Harman, 2011, pp. 459).

Rosen's (2004) defense of the claim that agents are only culpable for their actions performed from moral ignorance if they are also culpable for their moral ignorance is also unsuccessful. We *can* reasonably expect agents to act in ways that they blamelessly, falsely believe to be permissible. Even if we stipulate that an agent is non-culpable for his moral ignorance, it does not follow that we cannot reasonably expect him to avoid acting on it. This would, of course, require him to act in a manner that appears irrational from his own perspective, but as we saw above, this is not necessarily unreasonable to expect. Imagine, for example, that Stan is non-culpable for his racist beliefs, insofar as we could not reasonably expect him to avoid them. Even if this is true, one could reasonably expect him not to make comments he recognizes to be racist to persons of color, even if he non-culpably believes this to be permissible. In expecting him to do so, we would be requiring him to refrain from acting in a way that treats his interests in making the comments (whatever these may be) as if they are of greater relative importance than the interests persons of color have in not being exposed to racist comments.

It's not clear how else we could defend the claim that agents are only culpable for their actions performed from moral ignorance if they are also culpable for their moral ignorance.¹¹⁰ Without this claim, it is not true that *we must either trace culpability for

¹¹⁰ Note that there are also positive reasons to reject this claim. Without it, we can better explain our intuitive reactions to some cases. In some cases, we tend to think that an agent is not culpable for his moral ignorance, insofar as he could not be reasonably expected to avoid it, but is culpable for acting from this ignorance. E.g. we might think that, given the horribly abusive condition of an agent's upbringing, he could not be reasonably expected to avoid forming the false moral belief that human life is fairly worthless, and is therefore not culpable for doing so. However, this does not suggest that he could not be reasonably expected to avoid acting from this ignorance, by murdering another individual, and thereby be non-culpable for doing so. On my view, there will very rarely (never?) be an instance in which expecting an agent to refrain from murdering another agent would require him to treat their interests as if they are of greater weight than his own. It will therefore very rarely or

moral ignorance to culpability for some past, knowing mismanagement of one's moral beliefs, or face an infinite regress. Without *, there is no reason to think that an agent cannot be culpable for mismanaging his moral beliefs, even if he did so unknowingly. If this is right, then even if we grant that agents are only culpable for their moral ignorance if it is the result of a past culpable act or omission, we need not also grant that agents are only culpable for their moral ignorance if it is the result of past, knowing mismanagement of their moral beliefs. Again, I conclude that we ought to reject the latter, and thereby the skeptical argument.

Reasonable Moral Expectations and Culpability Views

My account of reasonable moral expectations impacts culpability views less forcefully. However, it does suggest that we should accept them only if they are significantly revised. Proponents of culpability views argue that agents are always culpable for their moral ignorance insofar as it always manifests an insufficient care for things of moral significance, or objectionable moral attitudes. In Chapter 3, I argued that it is not plausible that moral ignorance *always* manifests insufficient moral care or objectionable moral attitudes. Moreover, I argued there is good reason to doubt the weaker claim that agents are always culpable for their moral ignorance *insofar* as it manifests insufficient moral care or objectionable moral attitudes. I considered two objections to this weaker claim: (1) some agents are not in an epistemic position to form true moral beliefs, and therefore are not culpable for their resulting moral ignorance, and (2) in some cases, agents cannot be reasonably expected to care sufficiently, or to avoiding holding objectionable

never be unreasonable to expect him to refrain from murdering an agent. This is true regardless of what he (blamelessly) believes about the permissibility of doing so.

moral attitudes, and thereby the beliefs that manifest them. I argued that a full evaluation of (1) and (2) requires an account of the conditions under which an agent can be reasonably expected to X (and would therefore have to wait until this chapter). I now have such an account, and so can fully evaluate (1) and (2). I argue that (1) is sometimes plausible, and that it also sometimes plausible that agents cannot be reasonably expected to form true moral beliefs, given the cultural circumstances in which they find themselves. I argue that (2) is correct.

In light of these objections, culpability views ought to be modified to hold that agents are always culpable for their moral ignorance insofar as it manifests insufficient moral care or objectionable moral attitudes, *and they could have been reasonably expected to avoid their moral ignorance, and to avoid the insufficient care or objectionable moral attitudes it manifests*. I argue proponents of culpability views ought to accept this modification of their view as a welcome amendment; it allows them to offer a more plausible view of culpability for moral ignorance, without sacrificing the main aims and characteristics of their views.

(1)

The objection that some agents are not in the position to form true epistemic beliefs is offered by Wieland (2015). I described this objection in Chapter 3:

Wieland (2015) argues that some agents are not in a position to hold epistemically justified true moral beliefs insofar as agents do not always have access to evidence that would enable them to form justified true moral beliefs. He argues that there are three potential sources of evidence supporting true moral beliefs: testimony, mere reflection, and direct experience (Wieland, 2015, p. 9). It is not always the case that agents have access to evidence via these sources. Testimony does not always

provide one with evidence that tells against the morality of a practice. E.g it may not do so in cases in which an immoral practice is widely accepted (Wieland, 2015, p. 10). Some agents simply lack the capacity to conclude, through mere reflection, that a practice is wrong (10). Wieland (2015) holds, e.g. that not everyone has the capacity to see “by mere reflection that the suffering of slaves outweighs the benefits of slavery for a given society (such as Athens)” (11). This is because not everyone “is in a position to make such rather philosophical inferences” (Wieland, 2015, p. 11). Finally, not everyone will have access to direct experience that provides evidence for true moral beliefs (e.g. not all Athenians will have direct contact with slaves) (Zinn, 2016, p. 68).

I argued in Chapter 3 that while some agents may lack access to testimony or direct experience that would provide them with evidence for true moral beliefs, it is implausible that fully-abled, normally functioning adult agents lack access to evidence for true moral beliefs via reflection. However, I granted the possibility that some agents, through no fault of their own, may not be very good at accessing evidence for true moral beliefs through mere reflection, and therefore if they also lack access to evidence for them via testimony and direct experience, it seems like they could not be reasonably expected to form true moral beliefs. I offered the following fictionalized example of a potential such agent, and the following analysis.¹¹¹

Mr. S is from a rural pocket of Malawi. He falsely believes that women and girls are of less moral value than men and boys. His belief is similar to the rather widely held belief that non-human animals have less moral value in virtue of being members of a different species; Mr. S believes that women and girls have less moral value in virtue of being a different sex. This is the prevailing belief in Mr. S’s community, where women and girls are widely accepted to be similar in value to livestock. Some men in Mr. S’s community treat their wives and daughters with great love and respect, but Mr. S and other men in his community do not take this to be evidence that they are equals.

¹¹¹ As noted in Chapter 3, this example is based on an actual agent, Mr. Simbeye, discussed by LaFrainere (2005).

Instead they understand these men to have a particularly strong emotional attachment to their wives and daughters, similar to the attachment that some individuals have to pets. Mr. S isn't very empathetic or imaginative, and therefore is not very good at arriving at moral truths via mere reflection.

One might argue that Mr. S could not be reasonably expected to avoid holding his false moral belief that women and girls are of less moral value than men and boys. Testimony confirms this belief. Direct experience arguably does not provide evidence for or against it. Just as there is not much evidence via direct experience that counts for or against a view that moral value depends on one's species membership, there is not much that counts for or against a view on which moral value depends on one's sex. If Mr. S is poor at arriving at moral truths via mere reflection, due to his diminished capacities, one might argue that Mr. S cannot be reasonably expected to avoid his false moral belief about the moral value of women and girls. (Zinn, 2016, p. 70).

We are now in a position to determine if this is right—if Mr. S and agents relevantly like him could not be reasonably expected to form true moral beliefs in virtue of lacking evidence for them via testimony or direct experience, and insofar as they have diminished capacities, through no fault of their own, to access evidence for them through mere reflection. While I do not think this is true with respect to Mr. S in particular, it is plausible that it is sometimes the case that we cannot reasonably expect agents in this type of epistemic situation to form true moral beliefs. Agents with diminished capacities for moral reflection will have to go to greater lengths in order to bring it about that they have true moral beliefs. Whether or not we can reasonably expect them to do this depends on the interests other agents have in them forming true moral beliefs.

Consider Mr. S in particular. It would have been more difficult for Mr. S to form true moral beliefs than it would have been for someone with greater capacities for reflection. E.g. let's say that when Mr. S reflected on whether or not women and girls have equal moral value, he tried to consider the matter from their perspective, but all he came up with was that

it would be bad to be a woman or girl and to be treated as if one were livestock. If he were a woman or girl, he would not want to be treated this way. But, livestock don't particularly like being made to perform labor, or having their movement otherwise constrained, and many do not thereby conclude that they are of equal moral value as human beings. Mr. S reasoned that from the mere fact that he would not like to be treated as women and girls are treated in his community, it does not follow from that women and girls have equal moral value. For Mr. S to have seen that they do, he may have needed to "supplement" his reflection. E.g. he may have needed to ask women and girls what it is like to be them, and how they experience the treatment to which they are subjected. He may have needed to talk to them about whether or not *they* think they have equal moral value. Alternatively, he may have needed to seek out other men in his community who believe women to have equal moral value and ask them their reasons for believing this.

I think that in this particular case, given the very significant interests women and girls in Mr. S's community have in him forming true moral beliefs about their value, Mr. S. could have been reasonably expected to perform the above tasks in order to remedy his false moral beliefs. I think that in this particular case, given the very significant interests women and girls in Mr. S's community have in him forming true moral beliefs about their value, Mr. S. could have been reasonably expected to perform the above tasks in order to remedy his false moral beliefs. One might think that in reflecting upon the issue at all, Mr. S did all he could have been reasonably expected to do in order to bring it about that he had true moral beliefs. Moreover, he saw no reason to do more than this, so we could not have reasonably expected him to do more. While efforts like Mr. S's may be all that we can reasonably expect of moral agents in ordinary cases, I think it is highly plausible that we can

reasonably expect agents to go to extraordinary lengths to ensure that moral beliefs are accurate when these moral beliefs seriously affect the lives of many other individuals. We can reasonably expect them to subject their moral beliefs to a higher level of scrutiny, even if they see no reason to do so. In many cases, even if the agents in question have interests in not doing so, these interests will not plausibly be of greater relative importance than the interests others have in them arriving at true moral beliefs. In expecting them to refrain from exposing their moral beliefs to only an ordinary level of scrutiny, we expect them to refrain from acting in a way that treats the interests of others as if they are of less relative importance than their own. Returning to Mr. S, in expecting him to have refrained from settling for ordinary efforts at arriving at true moral beliefs, we would have expected him to refrain from acting in a way that treats the interests of women and girls in his community as if they are of less relative importance than his own interests in not submitting his moral beliefs to a high level of scrutiny. This would have been a reasonable moral expectation. The moral gravity of his situation called for more serious efforts on his part to arrive at true moral beliefs.

But, it is not plausible that it is *always* true that agents can be reasonably expected to form true moral beliefs, despite lacking evidence for such beliefs via testimony and direct experience, and having diminished capacity for reflection. It may not be true when nothing very significant depends on an agent having true moral beliefs, and the agent would have to go great lengths to develop them. In expecting an agent to do this, we would not plausibly be holding her to a demand that requires (i) or (ii).

Moreover, it is plausible that sometimes, given their cultural situation, it may be unreasonable to expect agents to form true moral beliefs or to remedy their false moral

beliefs, insofar as doing so requires them to treat the interests of others as if they are of greater relative importance than their own. I.e. in expecting this, we would hold them to a demand that meets neither (i) nor (ii), and instead requires the opposite of the second disjunct of (ii). For example, it might be unreasonable for this reason to expect an agent who is raised in an isolated, misogynistic culture, to avoid having false moral beliefs concerning the place and value of women for reasons unrelated to the efforts he would have to exert in order to bring this about. Imagine, instead, that this agent faces a number of different challenges in overcoming the moral ignorance perpetuated by his culture. Breaking with the misogynistic values of his culture may come at great costs. E.g. he might be rejected by his family and/or religious community, may face persecution and bodily harm, and may be barred from certain educational and employment opportunities. In such cases, where much is at stake should the agent form true moral beliefs that are at odds with the dominant views of his culture, it is plausible that demanding that he nevertheless do so may, *in some cases*, require that he treat the interests of others as if they are of greater moral importance than his own. Whether or not this is true must be determined on a case-by-case basis, and involves an analysis of what is at stake for other agents, should he not avoid or remedy his false moral beliefs, in comparison to what is at stake for him if he does.

(2)

It is also unreasonable, in some cases to expect agents to avoid the insufficient care or objectionable moral attitudes manifested by their moral ignorance. When this is true, it is not plausible that they are culpable for their insufficient moral care or objectionable moral attitudes, or for their moral ignorance, insofar as it manifests them. In Chapter 3, in support

of the claim that it is sometimes unreasonable to expect agents to avoid caring insufficiently about things of moral importance, or to avoid holding objectionable moral attitudes, I pointed out that psychologists have found that in response to traumatic events, some agents engage in “emotional numbing” (Zinn, 2016, p. 72-73). They suppress the normal emotions they might feel about actions or events in self-defense; allowing themselves to feel these feelings would be too painful. The practice of “emotional numbing” can become habitual, and as a result, can lead to a “withdrawal of empathy and the cultivation of callousness toward others” (Kerig et. al, 2012, p. 277). It can lead to what can be properly characterized as insufficient moral care or objectionable moral attitudes. I considered the following specific examples of this (Zinn, 2016, p.72-73):

...the following self-report from a veteran, Billie Mac, cited in Gilliland and James (2013).

Looking back on it, I cannot believe how callous I have become. SOP [standard operating procedure] was “It don’t mean nothin’, screw it, drive on.” This would be right after a B-40 round had blown your buddy’s brains all over you. You had to put it behind you to survive. A guy fell off the construction site I was working on last fall and splattered himself all over the pavement. I sat on a steel beam about 30 feet above the guy and just kept eating my lunch. No big deal (160)!

Consider also the following report of a young woman, charged with assault, who has a history of trauma.

Sadhanna is a 22-year-old woman mandated to outpatient mental health and substance abuse treatment as the alternative to incarceration. She was arrested and charged with assault after arguing and fighting with another woman on the street. At intake, Sadhanna reported a 7-year history of alcohol abuse and one depressive episode at age 18. She was surprised that she got into a fight but admitted that she was drinking at the time of the incident. She also reported severe physical abuse at the hands of her mother’s boyfriend between ages 4 and 15. Of particular note to the intake worker was Sadhanna’s matter-of-fact way of presenting the abuse history. During the interview, she clearly

indicated that she did not want to attend group therapy and hear other people talk about their feelings, saying, “I learned long ago not to wear emotions on my sleeve.”

Sadhanna reported dropping out of 10th grade, saying she never liked school. She didn’t expect much from life. In Sadhanna’s first weeks in treatment, she reported feeling disconnected from other group members and questioned the purpose of the group. When asked about her own history, she denied that she had any difficulties and did not understand why she was mandated to treatment. She further denied having feelings about her abuse and did not believe that it affected her life now. Group members often commented that she did not show much empathy and maintained a flat affect, even when group discussions were emotionally charged (Substance Abuse and Mental Health Services Administration, 2014, p. 62).

I argued in Chapter 3 that it is *plausible* that Billie Mac and Sadhanna could not be have been reasonably expected to avoid their insufficient moral care or objectionable moral attitudes insofar as they are the products of emotional numbing, and they could not have been reasonably expected to avoid engaging in this practice. I am now in a better position to support these claims. In many cases, expecting agents like Billie Mac and Sadhanna to have refrained from engaging in emotional numbing, and to thereby avoid their insufficient moral care or objectionable moral attitudes, would not have plausibly required them to (i) or (ii). In fact, I think it would have required them to do the opposite of the second disjunct of (ii); it would have required them to treat the interests of others as if they are of greater weight than their own. If this is right, then they are not plausibly culpable for engaging in emotional numbing, their resulting insufficient moral care or objectionable moral attitudes, or their moral ignorance insofar as it manifests them.¹¹²

¹¹² It is possible that agents like Billie Mac and Sadhanna could have been reasonably expected to refrain from emotional numbing on non-moral grounds. E.g. perhaps their

Consider Billie Mac in particular. In holding Billie Mac to the demand that he care sufficiently about the loss of a human life (or that he not lack regard for it), we would expect to have not engaged in the practice that gave rise to this insufficient care or objectionable moral attitude—i.e. we expect him to have not engaged in emotional numbing. I do not think we could have reasonably expected him to have done so. Others do have significant interests in Billie Mac caring sufficiently about the loss of a human life. E.g. they have interests in him caring enough about it that he would prevent it if he could, or that he react appropriately to others who are distressed by the loss of a human life.¹¹³ However, it is not clear that these interests are of greater relative importance than the interests that Billie Mac had in engaging in emotional numbing, which resulted in his failure to care sufficiently about the loss of human life. He suggests that doing so was (is?) necessary for his emotional survival, and I see no reason to doubt this. The interests Billie Mac had in engaging in emotional numbing are plausibly of greater relative importance than the interests others have in him caring care sufficiently about the loss of human life.¹¹⁴ It would have been unreasonable of us to expect him to refrain from engaging in the practice of moral numbing insofar as it would have required the opposite of the second disjunct of (ii); it would have required him to treat the interests of others as if they are of greater relative importance than his own. He is therefore

spouses or children could have reasonably expected them to do so. If this is right, they may be culpable as a spouse or parent, while not also being morally culpable.

¹¹³ The description of Billie Mac offered above gives no indication that his lack of care about the loss of a human life would lead to more result in him actively harming another, so we need not include preventing such actions in the interests others have in him caring sufficiently about the loss of human life.

¹¹⁴ Again, this analysis depends on the particular features of this case, which do not indicate that Billie Mac was inclined to harm others in virtue of his lack of sufficient moral care.

not plausibly culpable for not doing so, or for the insufficient care that resulted from this practice.

I think many cases like Billie Mac's and Sadhana's will follow this pattern. Unless the interests other agents have in them not caring insufficiently or not holding objectionable moral attitudes are very compelling, it will be unreasonable to expect agents who engage in emotional numbing in response to traumatic events to refrain from doing so, and to thereby avoid or remedy their relevant insufficient care or objectionable moral attitudes. If this is right, then in some cases, we cannot justifiably blame agents for caring insufficiently about things of moral importance or for holding objectionable moral attitudes, and therefore cannot blame them or their moral ignorance in virtue of the fact that it manifests these things.

Application to culpability views

If what I have argued is right, it has significant implications for culpability views. It would be unreasonable to expect some agents to avoid their moral ignorance because, given their epistemic or cultural circumstances, doing so would not require them to (i) or (ii), but instead would require the opposite of the second disjunct of (ii). It would require them to treat the interests of others as if they are greater relative importance than their own. It would be unreasonable to expect other agents to care sufficiently about specific things of moral significance, or to not hold specific objectionable moral attitudes insofar as this would have required them to have not engaged in emotional numbing, which in turn would also have required the opposite of the second disjunct of (ii); it would have required them to treat the interests of others as if they are greater relative importance than their own. In these cases,

agents cannot be justifiably blamed for their moral ignorance. In blaming them for their moral ignorance, we would endorse a demand that they treat the interests of others as if they of greater relative importance than their own. We would thereby treat their interests as such—we would treat them unfairly.

Application to Harman's View

These observations suggest that Harman is mistaken in holding that agents are always culpable for their moral ignorance, as long as it manifests a failure to care sufficiently about things of moral significance. Her view ought to be amended (at least) to hold that agents are culpable for their moral ignorance insofar as it manifests insufficient care *and it is the case that they could have been reasonably expected to avoid their moral ignorance and the insufficient care it manifests*. But I do not think this amendment threatens the core elements or aims of Harman's culpability view. Harman develops her view primarily in response to the skeptical argument; she develops it in response to the claim that in order to be culpable, moral ignorance must be traced to a past, knowing mismanagement of one's moral beliefs. She rejects this claim, arguing that agents can be directly culpable for their moral ignorance insofar as it manifests a failure to care sufficiently about things of moral significance. For everything I have said, it may be true that agents have an obligation to care sufficiently about things of moral significance, and to hold moral beliefs that manifest this care. What I have added is that they are only culpable for failing to do so if they could have been reasonably expected to care sufficiently about specific things of moral significance, and to avoid moral ignorance that manifests insufficient care. In order to accommodate this modification, Harman need only give up the claim that agents are *always* culpable for their

moral ignorance insofar as it manifests a failure to care sufficiently about things of moral significance. She can do this without giving up the claim that agents are directly, and most often culpable for their moral ignorance.

Applications to non-volitional culpability views

In light of these objections, proponents of non-volitional culpability views must also revise their views that agents are always culpable for their moral ignorance insofar as it manifests objectionable moral attitudes. They ought to instead hold that agents are always culpable for their moral ignorance insofar as it manifests objectionable moral attitudes, *and they could have been reasonably expected to avoid their moral ignorance and the attitudes it manifests*. This will require proponents of non-culpability views to modify the non-volitional views of blame that motivate their views. But, again, I do not think this requires proponents of culpability views to give up the main elements or aims of their views.

As noted previously, non-volitional culpability views are motivated by non-volitional views of blame. Proponents of non-volitional views of blame take issue with their counter-parts, volitional views of blame, which hold that agents are only culpable for X if is the product of their voluntary choices. As I explained in Chapter 3,

according to some non-volitional views, blame consists of a judgment that an agent has or has displayed an objectionable moral attitude, and as is the case with any judgment, it is justified as long as this judgment is accurate, and we are justified in reaching it, given our evidence. E.g. when we blame someone for driving recklessly, we judge him to have displayed a lack of respect for the lives and safety of others. This judgment is justified if he in fact displayed a lack of respect, and we have sufficient evidence for judging him to have done so. According to other non-volitional views of blame, when we blame someone for X, we implicitly demand that she justify X; in doing so we demand that she justify the moral attitudes reflected by X, and if she cannot, we blame her for

them. For example, when I blame Sam for cheating on an exam in my class, I demand that Sam justify this action. To do so, he must defend his evaluation of the reasons for cheating. If he can do so (e.g. perhaps he had good reason to believe his life depended on passing this exam, and rightfully judged his life to be more valuable than being honest in this particular scenario), then I ought to rescind my blame. Alternatively, if he cannot do so (e.g. because he knows he behaved badly, or because he falsely believes that his grades are more valuable than being honest), my blame is justified (Zinn, 2016, p.75).

I neither defend, nor object to the conditions under which proponents of non-volitional culpability views hold agents are culpable for their moral ignorance. However, in light of the objections discussed in this section, I think they ought to amend these conditions. They ought to hold the following revised views. (A) An agent is culpable for X if X is, or manifests an objectionable moral attitude, our judgment that X is or has done so is justified, *and the agent could have been reasonably expect to avoid X or the objectionable moral attitude it manifests*. (B) An agent is culpable for X if X is, or manifests an attitude for which the agent lacks an adequate defense, *and she could have been reasonably expect to avoid X and/or the objectionable moral attitude it manifests*. If proponents of culpability views refuse these modifications, they will, on my view, be committed to unjustifiably blaming agents insofar as they will hold that agents can be culpable for that which they could not be reasonably expected to avoid.¹¹⁵

But, as was the case with Harman's view, I do not think these modifications alter the spirit or aims of non-volitional culpability views. In particular, they do not commit them to

¹¹⁵ In Chapter 3, I discussed Hieronymi's (2004) argument that the fact that it would be unreasonable to expect an agent to X does not make it unfair to blame her for X. Hieronymi (2004) argues that this is because blame is only unfair if the judgment involved in blame is false. I hope it is clear how, on my view, blame can be unfair in another sense. It can be unfair if it fails to treat like cases alike by blaming agents for failing to meet unreasonable expectations.

volitional views of blame. This would be true only if, in granting that agents are culpable for X only if it was not unreasonable to expect them to avoid X, they also grant that an agent is culpable for X only if it is the result of her voluntary choices. However, there is no reason that they must do so. There is no reason, on my view, to think that whether or not an expectation is unreasonable is necessarily related to an agent's voluntary choices. Moreover, accepting the modifications I have suggested will allow proponents of non-volitional culpability views (and of non-volitional views of blame more generally) to avoid the less savory implications of their views, e.g. that agents are culpable for actions that manifest objectionable moral attitudes even if such actions are the result of severe brain injuries or implantation.¹¹⁶

4.8 A Moderate and Intuitive Approach to Culpability for Moral Ignorance

I have argued for an account of reasonable moral expectations, and their relationship to culpability, which I hope will be both practically and theoretically useful. Many of our practical judgments and their consequences rely on the further judgment that an agent could have been reasonably expected to avoid X. My account is hopefully of use in clarifying and sharpening these judgments, and also usefully applied to other debates in the literature on moral responsibility.

I have also provided what amounts to a moderate and intuitive approach to agents' culpability for their moral ignorance. We have no reason to accept the skeptical argument's conclusion that they are almost never culpable for it, or the culpability views' claim that

¹¹⁶ Levy (2005) and Miller (2014) argue that Angela Smith is committed to this view, despite her (weak) efforts to avoid it. Scanlon (1998) openly embraces this view, arguing that even if a shift in attitudes was due to a brain injury, if the shift is permanent, the agent will be culpable for these attitudes (279).

they are almost always culpable for it. Instead, because agents sometimes cannot be reasonably expected to avoid their moral ignorance, they sometimes are not culpable for it. This must be determined on a case-by-case basis, and it may turn out to be quite rare. It depends both on how costly it would be for the agent to avoid moral ignorance, and the likely affects that the agents' ignorance may have on other agents. While we can reasonably expect an agent to assume high costs in order to avoid moral ignorance that is likely to lead to severe consequences for other agents (e.g. a belief that human life is worthless), we cannot expect agents to assume high costs in order to avoid moral ignorance that will likely only have minor effects on other agents (e.g. a belief that people cannot be trusted with their jewelry). This tracks our intuitions concerning the matter. We tend to think that agents most often are culpable for their moral ignorance, although there are exceptions, particularly when it is the result of unfortunate formative or cultural circumstances (in which case, the agent must endure high costs to avoid moral ignorance).

References

- Adams, Robert. (1985). Involuntary Sins. *Philosophical Review* 94: 3-31. doi: 10.2307/2184713
- Beah, Ishmael. (2014, Jan. 7). The Making, and Unmaking, of a Child Soldier. *The New York Times Magazine*. Retrieved from: http://www.nytimes.com/2007/01/14/magazine/14soldier.t.html?_r=0
- Benson, Paul. (2001). Culture and Responsibility: A Reply to Moody-Adams. *Journal of Social Philosophy* 32:4: 610-620. doi: 10.1111/0047-2786.0011
- Björnsson, Gunnar. Forthcoming. Explaining (away) the epistemic condition on moral responsibility. *Responsibility—The Epistemic Condition*. Ed. by Robichaud & Wieland. Oxford University Press.
- Burns, Jeffrey M. et. al. (2003). Right Orbitofrontal Tumor With Pedophilia Symptom and Constructional Apraxia Sign. *Arch Neurol*. 2003;60(3):437-440. doi:10.1001/archneur.60.3.437
- Buruma, Ian. (2006, September 18). War and Remembrance: Shrouded by the Günter Grass controversy is an extraordinary new memoir. *The New Yorker*. Retrieved from: <http://www.newyorker.com/magazine/2006/09/18/war-and-remembrance>
- Calhoun, Cheshire. (1989). Responsibility and Reproach. *Ethics* 99:2: 389-406. doi: 10.1086/293071
- Ciurria, Michele. (2014). Moral Responsibility and Mental Health: Applying the Standard of the Reasonable Person. *Philosophy, Psychiatry, and Psychology* 21 (1):1-12. doi: 10.1353/ppp.2014.0007
- FitzPatrick, William. (2008). Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge. *Ethics* 118: 589-613. doi: 10.1086/589532
- Fricker, Miranda. (2014). What is the Point of Blame? A Paradigm Based Explanation. *Nous* 50:1:165–183. doi: 10.1111/nous.12067
- Gallup Editors. (2014, Dec.). Gallup Review: Black and White Differences in Views on Race. Retrieved from: <http://www.gallup.com/poll/180107/gallup-review-black-white-differences-views-race.aspx>.
- Gilliand, Burl and Richard James. (2016). *Crisis Intervention Strategies*. Cengage.
- Gist, Karin (writer), & Chandra Wilson (director). (2016, April 7). I Wear the Face [*Grey's Anatomy*]. Betsy Beers (producer). Los Angeles, CA: ABC studios.

Grass, Günter. (2007). *Peeling the Onion*. Translated by Michael Henry Heim. Harcourt, Inc.

Grant, Jon E. & Suck Won Kim. (2002). Clinical Characteristics and Associated Psychopathology of 22 Patients With Kleptomania. *Comprehensive Psychiatry*. 43:5: 378-384. doi: 10.1053/comp.2002.34628

Guerrero, Alexander. (2007). Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution. *Philosophical Studies* 136 (1):59-97. doi: 10.1007/s11098-007-9143-7

Harman, Elizabeth. (2011). Can Moral Ignorance Exculpate? *Ratio* 24 (4):443-468. doi: 10.1111/j.1467-9329.2011.00511.x

Harman, Elizabeth. (2014). Ethics is Hard: What Follows? Unpublished Manuscript.

Hieronymi, Pamela. (2004). The Force and Fairness of Blame. *Philosophical Perspectives* 18: 115-148. doi: 10.1111/j.1520-8583.2004.00023.x

Hieronymi, Pamela. (2008). Responsibility for Believing. *Synthese* 161 (3):357-373. doi: 10.1007/s11229-006-9089-x

Jeske, Diane. (1998). Friends, Family, and Special Obligations. *Canadian Journal of Philosophy* 28(4):527 – 555. doi: 10.1080/00455091.1998.10715984

Kerig, Patricia et. al. (2012). Nothing really matters”: Emotional numbing as a link between trauma exposure and callousness in delinquent youth. *Journal of Traumatic Stress* 25:3: 272–279. doi: 10.1002/jts.21700

King, Matt. (2009). The Problem With Negligence. *Social Theory and Practice* 35 (4):577-595. doi: 10.5840/soctheorpract200935433

Koutsikou, S., Crook, J. J., Earl, E. V., Leith, J. L., Watson, T. C., Lumb, B. M., & Apps, R. (2014). Neural substrates underlying fear-evoked freezing: the periaqueductal grey–cerebellar link. *The Journal of Physiology*, 592(10), 2197–2213. doi: 10.1113/jphysiol.2013.268714

Kukla, Rebecca and Mark Lance. (2009). *‘Yo!’ and ‘Lo!’: The Pragmatic Topography of the Space of Reasons*. Harvard University Press.

LaFraniere, Sharon. (27 Nov., 2005). Forced to Marry Before Puberty, African Girls Pay Lasting Price. *The New York Times*. Retrieved from: <http://www.nytimes.com/2005/11/27/world/africa/forced-to-marry-before-puberty-african-girls-pay-lasting-price.html>

Levy, Neil. (2005). The Good, the Bad and the Blameworthy. *Journal of Ethics and Social Philosophy* 1(2): 2-16.

- Levy, Neil. (2009). Culpable Ignorance and Moral Responsibility: A Reply to FitzPatrick. *Ethics* 119 (4):729-741. doi: 10.1086/605018
- Ludden, Jennifer. (8, February, 2011). Ask For a Raise? Most Women Hesitate. *National Public Radio*. Retrieved from <http://www.mprnews.org/story/npr/133599768>
- Lynch, Kevin. (2016). Willful Ignorance and Self-deception. *Philosophical Studies* 173(2):505–523. doi: 10.1007/s11098-015-0504-3
- Macmillian, Malcom. (2002). Restoring Phineas Gage: a 150th Retrospective. *Journal of the History of the Neuroscience* 9:1: 46-66. doi: 10.1076/0964-704X(200004)9:1;1-2;FT046
- MacNamara, Coleen. (2011). Holding Others Responsible. *Philosophical Studies* 152:81–102. doi: 10.1007/s11098-009-9464-9
- Mason, Elinor. (2015). Moral Ignorance and Blameworthiness. *Philosophical Studies* 172 (11):3037-3057. doi: 10.1007/s11098-015-0456-7
- Miller, Daniel. (2014). Answerability, Blameworthiness, and History. *Philosophia* 152:81–102. doi: 10.1007/s11406-013-9510-x
- Moody-Adams, Michele. (1993). On the Old Saw that Character is Destiny. Pp.111-131 in *Character and Morality*, ed. by Owen Flanagan and Amelie Rorty. MIT Press.
- Moody-Adams, Michele. (1994). Culture, Responsibility, and Affected Ignorance. *Ethics* 104:2: 291-309. doi: 10.1086/293601
- Montmarquet, James. (1999). Zimmerman on Moral Ignorance. *Ethics* Vol. 109:4: 842-845. doi: 10.1086/233949
- Rosen, Gideon. (2003). Culpability and Ignorance. *Proceedings of the Aristotelian Society, New Series* 103:61-84. doi: 10.1111/j.0066-7372.2003.00064.x.
- Rosen, Gideon. (2004). Skepticism about Moral Responsibility. *Philosophical Perspectives* 18 (1):295–313. doi: 10.1111/j.1520-8583.2004.00030.x
- Rosen, Gideon. (2008). Kleinbart the Oblivious and Other Tales of Ignorance and Responsibility. *Journal of Philosophy* 105(10): 591-610. doi: jphil20081051023
- Santelices, Maria Veronica and Mark Wilson. (2010). Unfair treatment? the case of Freedle, the SAT, and the standardization approach to differential item functioning. *Harvard Educational Review* 80(1):106-142. doi: 10.17763/haer.80.1.j94675w001329270

- Scanlon, T.M. (1998). *What We Owe Each Other*. Boston: Belknap Press of Harvard University.
- Sher, George. (2001). Blame for Traits. *Nous* 35 (11):146-161. doi: 10.1111/0029-4624.00291
- Shoemaker, David. (2003). Caring, Identification, and Agency. *Ethics* 114(1): p. 88- 118. doi: 10.1086/376718
- Shoemaker, David. (2011). Attributability, Answerability, and Accountability: Toward a Wider Theory of Moral Responsibility. *Ethics* 121:3: 602-632. doi: 10.1086/659003
- Sykes, Gresham M. and David Matza. (1957). Techniques of Neutralization: A Theory of Delinquency. *American Sociological Review* 22 (6): 664-670.
- Smith, Angela. (2005). Responsibility for attitudes: activity and passivity in mental life. *Ethics* 115:2: 236-271. doi: 10.1086/426957
- Smith, Angela. (2007). On Being and Holding Responsible. *Journal of Ethics* 11 (4):465 - 484. doi: 10.1007/s10892-005-7989-5
- Smith, Angela. (2008). Control, Responsibility, and Moral Assessment. *Philosophical Studies* 138: 367-392. doi: 10.1007/s11098-006-9048-x
- Smith, Angela. (2012). Attributability, Answerability, and Accountability: In Defense of a Unified Account. *Ethics* 122(3): 575-589. doi: 10.1086/664752
- Smith, Angela. (2015). Attitudes, Tracing, and Control. *Journal of Applied Philosophy* 32 (2):115-132. doi: 10.1111/japp.12107
- Substance Abuse and Mental Health Services Administration. (2014). Trauma Informed Care in Behavioral Health Services. *Treatment Improvement Protocol (TIP) Series 57*. HHS Publication No. (SMA)13-4801. Retrieved from: <http://store.samhsa.gov/shin/content/SMA14-4816/SMA14-4816.pdf>
- Tognazzini, Neil. (2015). The Strains of Involvement. In Randolph Clarke, Michael McKenna & Angela M. Smith (eds.), *The Nature of Moral Responsibility* (19-44). Oxford: Oxford University Press.
- Willem von Hippel & Robert Trivers. (2011). The evolution and psychology of self-deception. *Behavioral and Brain Sciences* 34: 1-56. doi: 10.1017/S0140525X10001354
- Wallace, R. Jay. (1996). *Responsibility and the Moral Sentiments*. Boston: Harvard University Press.
- Wallace, R. J. (2008). Emotions, expectations and responsibility. In M. McKenna & P.

Russell (Eds.), *Free will and Reactive Attitudes* (157–185). Surrey, England: Ashgate.

Watson, Gary. (1996). Two Faces of Responsibility. *Philosophical Topics* 24:2: 227-248. doi: 10.5840/philtopics199624222

Wieland, Jan. (2015). What's Special About Moral Ignorance? *Ratio* 29 (1): 1-14. doi: 10.1111/rati.12127

Vargas, Manuel. (2005). The Trouble With Tracing. *Midwest Studies in Philosophy* 29 (1):269–291. doi: 10.1111/j.1475-4975.2005.00117.x

Zimmerman, Michael. (1997). Moral Responsibility and Ignorance. *Ethics* 107 (3):410- 426. doi: 10.1086/233742

Zimmerman, Michael. (2002). Controlling Ignorance: A Bitter Truth. *Journal of Social Philosophy* 33 (3):483–490. doi: 10.1111/0047-2786.00154

Zimmerman, Michael. (2008). *Living With Uncertainty*. Oxford: Oxford University Press.

Zinn, Carmen. (2016). Reasonable Moral Expectations and the Moral Ignorance Debate (Unpublished Doctoral Dissertation). Santa Barbara, California: University of California, Santa Barbara.