

UNIVERSITY OF CALIFORNIA

Santa Barbara

Responsibility Beyond Belief: The Epistemic Condition on Moral Responsibility

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

by

Christopher Michael Cloos

Committee in charge:

Professor Aaron Zimmerman, Chair

Professor Matthew Hanser

Professor Pamela Hieronymi

Professor John Greco

September 2018

The dissertation of Christopher Michael Cloos is approved.

Matthew Hanser

Pamela Hieronymi

John Greco

Aaron Zimmerman, Committee Chair

September 2018

Responsibility Beyond Belief: The Epistemic Condition on Moral Responsibility

Copyright © 2018

by

Christopher Michael Cloos

ACKNOWLEDGEMENTS

First and foremost, I wish to express my deep gratitude and thanks to my wonderful wife Kristie. Without her steadfast love and support this project and pursuing my PhD would have not been possible. This dissertation is dedicated to her.

Next I thank the professors at UCSB that have shaped my development as a philosopher and thinker. I thank C Anthony Anderson for helping shape my mind logically and analytically. I thank the late Anthony Brueckner for teaching me the importance of intellectual humility and his willingness to patiently explore with me literature in epistemology. I thank Aaron Zimmerman for modeling intellectual perseverance and courage. Aaron's ability to think autonomously in a search for deeper understanding has left a mark on my intellectual life. I also thank him for serving as my advisor. I thank Matthew Hanser for teaching me how to explore ethical issues with depth and rigor. Matthew's penetrating insights into complex moral issues has shaped my thinking and approach to ethics. I thank Kevin Falvey for enabling me to see the importance of action theory and philosophy of mind as it relates to the topics of my interest. I thank Tom Holden for modeling excellence in teaching and in modern philosophy, especially in giving me a deeper appreciation of the importance of Hume.

I thank professors outside UCSB for their guidance philosophically. Thank you to Pamela Hieronymi for serving on my committee and providing excellent comments that provided many helpful revisions. Additionally, thank you to John Greco for serving on my committee and providing excellent guidance and comments. Also, thank you to John Martin Fischer for feedback on earlier iterations of my

dissertation. Though John could not finish serving on my committee, his feedback on related papers and this project proved invaluable. I also thank Christopher Buford for working with me as a co-author and explorer of the literature in epistemology.

I wish to thank my fellow colleagues and graduate students that have shaped my thinking and teaching while at UCSB. Colleagues that come to mind are Rick Stoody, Keith Hess, Alex Bundy, Nat Tabris, Tim Butzer, Arnel Blake Batoon, Jonathan Caravello, Juliana Lima, Zachary Rentz, and David Mokriski.

Lastly, I wish to thank family and friends who have contributed to my academic journey in many ways. I thank my loving and supportive mother, Diana Cloos. Her shepherding allowed me to pursue higher education. She has encouraged me to approach life with curiosity and creativity—skills needed for adept philosophy. I thank my uncle Samuel Kampschmidt. He is a father figure that has always come along side me intellectually and pointed out possible avenues of discovery. He is a philosopher in his own right. I thank another father figure in my life—David Yee. David encouraged me to pursue my academic dreams. He was always available with wise council along the way. I thank my good friend Philip Myers for being there for me as a willing confidant and interlocutor. Philip encouraged me to think independently in the pursuit of my own ideas. I thank my sister Katie for being my biggest fan and a reminder of what achieving a PhD means for our family. Her love and support has meant the world to me. I thank my grandmothers Laura Wilborn and Ellie Kampschmidt. Their love and support enabled me to pursue my dreams. Lastly, I thank my late father Mickie Cloos. By example, he taught me how to teach skills to people, which served me in the classroom. He will always be my “coach.”

VITA OF CHRISTOPHER MICHAEL CLOOS

September 2018

EDUCATION

Bachelor of Arts in Philosophy, Cal Poly, San Luis Obispo, June 1999

Master of Arts in Philosophy, San Jose State University, June 2009

Doctor of Philosophy in Philosophy, UC Santa Barbara, September 2018

PROFESSIONAL EMPLOYMENT

2011-16: Teaching Assistant, Department of Philosophy, UC Santa Barbara

2016-17: Teaching Associate, Department of Philosophy, UC Santa Barbara

PUBLICATIONS

“Responsibilist Evidentialism,” *Philosophical Studies* (2015), 2999-3016

“A Dilemma for the Knowledge Despite Falsehood Strategy,” (with Christopher Buford), *Episteme* (2017), 1-17

AWARDS

Paul Wienpahl Award for Excellence in Teaching, UC Santa Barbara, Department of Philosophy

FIELDS OF STUDY

Major Fields: Epistemology and Ethics

ABSTRACT

Responsibility Beyond Belief: The Epistemic Condition on Moral Responsibility

by

Christopher Michael Cloos

Under what conditions is a person morally responsible for something they have done? Two conditions commonly endorsed as requirements for moral responsibility are freedom and knowledge. The person must have acted freely, and they must have known what they were doing. Corresponding to the two main requirements are excusing conditions. They absolve agents of moral responsibility for what they have done. A common excuse condition on the freedom side is external causal forces. If a strong gust of wind blows my car door into someone thereby injuring them, then I am not morally responsible for injuring them. I did not have the right sort of control over the car door hitting them to be responsible for injuring them. A common excuse condition on the knowledge side is ignorance. If I did not know that I was spooning cyanide instead of sugar into someone's coffee, then I am not morally responsible for poisoning them. But, there's a catch. External forces and ignorance fail as excuses when I am culpable for the forces having their influence or for putting myself in a position to be ignorant. In this dissertation, I focus on whether ignorance is an excuse. This move is strategic. Under the influence of the perennial free will debate, most of

the focus in the literature on moral responsibility has been on the freedom or control condition. This is unfortunate. It has left the knowledge or epistemic condition massively underdeveloped by comparison. My dissertation aims to contribute to rectifying this imbalance in the literature.

In the introductory chapter, I clarify the concept of moral responsibility at issue. It involves holding people morally responsible for bad behavior because they deserve blame for what they have done. Deserving blame is a matter of it being appropriate to target the person with certain reactions. These reactions are attitudes, such as being indignant or resentful that the person performed the bad action. On this model, moral responsibility is a decidedly social concept and is reflected in our everyday practices of judging people blameworthy for bad behavior. I round out this chapter looking at the nature of ignorance and the skeptical challenge that is the motivator for searching for a better characterization of the epistemic condition.

The first chapter presents the skeptical challenge. It claims that we are never warranted in judging with sufficient confidence that someone is morally responsible. This is troubling because our ordinary practices assume that we can form warranted or justified judgments of responsibility. If we cannot judge with reasonable confidence that a person is blameworthy for bad behavior in any specific case, this undermines a foundational way that we conduct ourselves morally. We regularly and naturally respond to bad behavior with reactive attitudes. This dissertation fleshes out a reasonable and permissive epistemic condition on moral responsibility to argue that we are warranted in our ascriptions of responsibility in all sorts of cases, even when a person is ignorant that bad behavior is in fact morally bad.

Chapter two gets into the nitty gritty of the literature on the epistemic condition. In this chapter I defend a position called “externalism” from an objection. I argue that culpable ignorance (i.e. ignorance that does not excuse) originates not just in an act whereby a person knowingly did not do something they thought they should do, such as knowing they should investigate something but failing to do so. The origin of culpability is also found in acts of vice, even if a person does not think of their vices as vices. And, acts of vice afford warranted ascriptions of responsibility.

In the third chapter, I embrace a view of moral responsibility that supports the thesis that moral ignorance never exculpates. I consider a defense of this thesis by Elizabeth Harman. However, Harman’s view is subject to a counterexample. I accept this counterexample and propose my own modified position. This position supports the thesis that moral ignorance never exculpates. Thus, this chapter is a further widening of the conditions under which the epistemic condition is satisfied, and this affords another avenue for justified ascriptions of responsibility.

The final chapter considers a counterexample to the view of moral responsibility developed in chapter three. Given that the concept of moral responsibility at issue is constituted by our ordinary practices of judging people blameworthy for their bad behavior, I look at an empirical study indicating how people ascribe responsibility in relation to the supposed counterexample. This neutralizes the counterexample, and it leads me to develop my own principle capturing the epistemic condition on moral responsibility. Then, I show the theoretical utility of my epistemic condition on moral responsibility by showing how it affords a novel approach to the problem of moral luck.

TABLE OF CONTENTS

INTRODUCTION	1
0.1 An Imbalance and a Skeptical Challenge	1
0.2 The Concept of Moral Responsibility	4
0.3 Ignorance as an Excuse Condition	9
0.4 The Nature of Ignorance	13
CHAPTER 1: A SKEPTICAL CHALLENGE	18
1.0 Introduction	18
1.1 Factual and Moral Ignorance	19
1.2 Duties of Inquiry	23
1.3 Conditions on Blameworthiness	27
1.4 Internalism and the Regress of Blameworthiness.....	31
1.5 A Skeptical Challenge	36
1.6 Two Responses to the Skeptical Challenge.....	41
1.7 Conclusion	45
CHAPTER 2: THE EXTERNALIST RESPONSE TO THE SKEPTICAL CHALLENGE	46
2.0 Introduction.....	46
2.1 Externalism’s Response to the Skeptical Challenge	46
2.2 Levy’s Objection to Externalism.....	61

2.3 Robichaud’s Two-pronged Reply to Levy’s Objection	68
2.4 Levy’s Rejoinder to Robichaud’s First Prong.....	79
2.5 Levy’s Rejoinder to Robichaud’s Second Prong.....	95
2.6 Weakening the Internalist Commitments.....	102
2.7 Externalism’s Akrasia Condition	110
2.8 Conclusion	119
CHAPTER 3: MORAL IGNORANCE IS NEVER EXCULPATORY.....	120
3.0 Introduction.....	120
3.1 Obligations and Blameworthiness	124
3.2 A Challenge to Harman’s Defense of No Excuse.....	135
3.3 New Attributionism	137
3.4 New Attributionism Applied	146
3.5 An Objection to New Attributionism.....	151
3.6 When Knowledge Grounds Blameworthiness.....	155
3.7 Interlude on the Nature of Moral Responsibility	163
3.8 An Epistemic Frankfurt-style Case	167
3.9 Conclusion	182
CHAPTER 4: EVIL DICTATORS, EPISTEMIC DIFFICULTY, AND MORAL LUCK	184
4.0 The Skeptical Challenge Redux.....	184

4.1 Real Self Views	186
4.2 Wolf's JoJo Case	187
4.3 Testing the JoJo Intuition	191
4.4 A Scalar Deep Self View	195
4.5 The Difficulty with Difficulty	199
4.6 Justified Responsibility Ascriptions.....	204
4.7 Summary of Key Results	210
4.8 The Problem of Moral Luck.....	213
4.9 Denying the Epistemic Argument	222
4.10 Accepting Constitutive Moral Luck	226
4.11 Chapter and Dissertation Conclusion	230
REFERENCES	232

INTRODUCTION

0.1 An Imbalance and a Skeptical Challenge

There is an imbalance in the literature on moral responsibility. Moral responsibility is comprised of at least two necessary conditions: a control condition and a knowledge condition.¹ Yet, in the literature on moral responsibility, there is a sprawling literature on the control (or freedom) condition and a comparatively small literature on the knowledge (or epistemic) condition. Lack of robust development of the epistemic condition is problematic because, as Jan Willem Wieland explains,

there is no complete account of moral responsibility absent some epistemic condition. Even if you are more concerned about other issues concerning responsibility, you will not be able to give a final verdict about an agent's responsibility without checking the epistemic condition. After all, as most will agree, the epistemic condition might render agents fully blameless all by itself, no matter what other conditions are met. (Wieland 2017a: 5)

The epistemic condition may absolve an agent of responsibility for wrongdoing.² If an agent is not aware that she is doing something morally wrong at the time of her action or omission, and her ignorance of this wrongfulness is blameless, then she is not morally responsible for her action or omission even if she satisfies other conditions

¹ The provenance of these conditions is found in the work of Aristotle. See *Nicomachean Ethics* iii.1-5. For excellent discussion of Aristotle's view of the epistemic condition see Sauv  Meyer (2011: Appendix I) and Eche ique (2012: Ch. 6).

² An example of this is from Aristotle's reading of Oedipus. When Oedipus sleeps with his mother and kills his father, he acts with control and so acts voluntarily, but given that Oedipus does not realize he is doing such things, he performs those actions involuntarily.

for moral responsibility such as possessing the right sort of control.³ This dissertation is part of a growing trend that recognizes the importance of developing the epistemic condition as part of offering a complete theory of moral responsibility.⁴

Though the literature on the epistemic condition is comparatively small, there is a growing debate over the epistemic condition.⁵ This debate centers on a skeptical challenge to moral responsibility—in any case of wrongdoing done in ignorance we are not warranted in judging that the agent is morally responsible for the action or omission.⁶ This is not a global skeptical challenge to the possibility of moral responsibility.⁷ Rather, according to this challenge, even if moral responsibility is possible such that people can satisfy metaphysical and epistemological conditions for being morally responsible, ascriptions of moral responsibility in particular cases cannot be made with sufficient confidence to warrant such ascriptions. We are never justified in judging that people are morally responsible. Roughly, this is because we

³ This claim is endorsed by all parties in the literature. It is the general statement that blameless ignorance excuses. However, it can also serve as a criterion of theory choice. The best theory will provide the best explanation of why it is the case that blameless ignorance excuses. In this dissertation, I do not explicitly offer an account of why blameless ignorance excuses. Although, the position I endorse at the end of this project could be formulated as such an explanation. Blameless ignorance excuses because the subject does not perform the unwitting wrongful act with a poor quality of will and judgment. For more on the claim that blameless ignorance excuses and how this might serve as a criterion of theory choice, see Wieland (2017a: Sect. 6). I thank John Greco for requesting clarification on the status of this claim.

⁴ Growth in the literature on the epistemic condition is evidenced by the forthcoming edited volume on the subject. See Robichaud and Wieland (2017).

⁵ I suspect this growth is due to the shift toward considering the nature and norms that govern the application of the reactive attitudes, including focus on the nature of blame. For more on this trend see Coates and Tognazzini (2013). Focus is broadening away from fixation on the free will debate.

⁶ Warrant is the property that turns a true belief into knowledge. The skeptical claim is that we cannot know that people are morally responsible for specific acts. Actually, the literature is imprecise on this point. The skeptical argument claims that we cannot ever know that people are morally responsible. However, we might possess less than knowledge-level warrant for truly believing that people are morally responsible. If this is true, we can confidently assert that people are morally responsible for their actions. It would take a further claim to argue that possessing knowledge is a requirement for making legitimate assertions. The knowledge norm of assertion is something people argue for. However, this debate would lead us too far afield. I simply flag this imprecision in the literature and note it as a possible further avenue to explore in addressing the structure of the skeptical argument.

⁷ For a global challenge to the possibility of moral responsibility see Galen Strawson (1994).

must locate the origin of responsibility somewhere in the causal chain that led to the wrongful act performed in ignorance. Establishing such an origin of responsibility requires identifying an act of knowing wrongdoing (i.e. an akratic act).⁸ But, for various reasons, we cannot with sufficient confidence identify such acts of knowing wrongdoing, thus we cannot form justified judgments that people are morally responsible for wrongful acts done out of ignorance that they are doing something morally wrong. Why is such a skeptical challenge troubling?

The skeptical challenge to moral responsibility is troubling because it delegitimizes a key foundation on which society regulates itself morally. Were ascriptions of responsibility unwarranted in *all* specific cases, this would undermine the legitimacy of people using ascriptions of praise and blame within the moral community to instruct members of the community away from bad behaviors and toward good behaviors. A foundation for legitimate applications of praise and blame within the moral community is desert. The person must deserve praise or blame before a person in the moral community with proper standing can target that person with a warranted application of the reactive attitudes.⁹ For instance, if I am not justified in ascribing blameworthiness to my child for unwittingly hitting his sister, then targeting my child with blame is misplaced. I aim my reactive attitude at a target I am not

⁸ I recognize that knowing wrongdoing is not equivalent to an act of akrasia. Knowing wrongdoing can occur without knowing that it is wrong all-things-considered, which is what akrasia requires. I follow the literature in using “knowing wrongdoing” loosely and interchangeably with “akrasia.” Thanks to Pamela Hieronymi for suggesting clarification of this point.

⁹ This is the traditional understanding of desert and praise and blame. Desert of the reactive attitudes is prior to legitimate application of such attitudes. However, someone like R. Jay Wallace (1996) reverses this order. Holding responsible via application of the reactive attitudes is prior to (or synonymous with) being deserving of such attitudes. Such a view may succumb to the skeptical challenge in a different way. I will not explore this here. Instead, I simply note that the skeptical challenge to justified ascriptions of desert concerning responsibility focuses on the traditional picture.

justified in believing deserving of that reaction. Blaming my child for hitting his sister is a way of instructing him that such behavior is morally problematic. If all applications of praise and blame within the moral community are mistargeted, they lose their legitimacy and the moral community is impoverished of the legitimate use of a key tool for moral instruction and education.

0.2 The Concept of Moral Responsibility

One way of characterizing the debate concerning the nature of moral responsibility is that it involves theorists suggesting that their characterization of moral responsibility is the genuine conception of moral responsibility because it best captures the basic notion of moral responsibility at issue. What is the basic notion?

It is helpful to hone in on the basic concept of moral responsibility by looking at related but non-equivalent forms of responsibility. One form of responsibility is causal responsibility. Moral responsibility is not equivalent to causal responsibility. A person is causally responsible for something when they are part of the causal sequence that produced that thing. Causal responsibility does not just apply to persons. A tornado that decimates a town is causally responsible for decimating that town. But, we would not say that the tornado is morally responsible for decimating the town. Legal responsibility is another type of responsibility that can come apart from moral responsibility. For instance, sexual infidelity during a committed, monogamous relationship may make the adulterer morally blameworthy, but such an act may not

make the adulterer subject to criminal liability or legal responsibility in that sense.¹⁰ So, moral responsibility is not equivalent to legal responsibility.

Now that we have roughly identified the basic notion of moral responsibility by seeing what it is not, it is helpful to mention two general approaches to characterizing what the basic notion of moral responsibility is. At a general level, there are forward-looking and backward-looking approaches to moral responsibility. On forward-looking approaches, people are morally responsible and rightly held morally responsible because of the knock-on effects of doing so. People are held morally responsible for their actions because doing so tends to promote future good effects (e.g., moral development, reconciliation, and rehabilitation). The downside with this forward-looking approach to capturing the basic notion of moral responsibility is that it fails to do justice to our robust social practices of holding people morally responsible. People are not judged morally responsible and held morally responsible just because doing so tends to promote future good effects. Rather people are judged morally responsible and held morally responsible for what they have done. They are subject to negative social reactions, such as the reactive attitudes of resentment and indignation, given that what they have done is deserving of such reactions. Such social practices are not legitimized merely by what future effects they may or may not secure. Rather, they are a fitting response to what has already occurred. Given the importance of desert (i.e., what a person deserves) in holding people blameworthy for their bad actions, backward-looking accounts of moral responsibility are often called desert-

¹⁰ This depends on the jurisdiction. Thanks to Aaron Zimmerman for reminding me that in many Muslim countries adultery is a criminal offence. Most Western states, though, show that it is possible for moral responsibility for adultery to come apart from legal (criminal) responsibility.

entailing accounts. Our judgments that people are blameworthy for their morally wrong actions entails that such people are deserving of blame. The literature on the epistemic condition on moral responsibility is an effort to capture the cognitive aspect of this backward-looking, desert-entailing characterization of the basic concept of moral responsibility. So, for the purposes of this project, I set aside the forward-looking approach.¹¹ Instead, the relevant conception of responsibility discussed in the literature is desert-entailing, and Derk Pereboom provides an apt description of this view of moral responsibility:

For an agent to be morally responsible for an action in the...desert entailing sense is for it to belong to her in such a way that she would deserve to be the recipient of an expression of moral resentment or indignation if she understood that it was morally wrong, and she would deserve to be the recipient of an expression of gratitude or praise if she understood that it was morally exemplary. The desert at issue here is basic in the sense that the agent, to be morally responsible, would deserve to be the recipient of the expression of such an attitude just because she has performed the action, given sensitivity to its moral status; not, for example, by virtue of consequentialist or contractualist considerations. Moral responsibility in this sense is presupposed by our

¹¹ Though, if you are interested in the forward-looking approach, see Derk Pereboom (2014: 134-5). Pereboom embraces a forward-looking conception of the basic notion of moral responsibility because of problems he raises for the backward-looking conception. Against Pereboom's forward-looking approach, Bruce Waller (2017: 30) argues that, "Pereboom's forward-looking moral responsibility is a gracious concession to the proponents of moral responsibility, but the result is a Tychonic system of moral responsibility that preserves elements of the just deserts moral responsibility system—a system that is not worth preserving in whole or in part." Waller's reference to a "Tychonic system" is a reference to the attempt by astronomer Tycho Brahe to still keep the earth at the center of the solar system yet make some concessions in the direction of the Copernican system. Waller's analogy is that just as the Tychonic view of the cosmos clung to a failed system (i.e., the Ptolemaic system, which embraced geocentrism), the forward-looking view trying to account for just deserts is likewise problematic.

attitudes of resentment, indignation, gratitude, and moral praise, since having such an attitude essentially involves the supposition that the agent in question basically deserves to be the recipient of its expression. In P. F. Strawson's account, moral responsibility is essentially tied to these reactive attitudes, and hence the basic-desert entailing sense is plausibly the variety that he brings to the fore. (Pereboom 2013: 611)

The relevant conception of moral responsibility at issue in the debate over the nature of moral responsibility and the epistemic condition is desert-entailing, non-consequentialist, and it involves a broadly Strawsonian conception of responsibility.¹² This conception of responsibility is inspired by Peter F. Strawson (1962) in that it emphasizes the importance of the reactive attitudes. It emphasizes the importance of social practices in making legitimate attributions of responsibility. However, this view is not something Strawson himself endorsed, as he did not think fittingness for application of the reactive attitudes necessary for blameworthiness. Strawson thought just desert of application of the reactive attitudes was, "internal to the structure...The existence of the general framework of attitudes itself is something we are given with the fact of human society...it neither calls for, nor permits, an external 'rational' justification."¹³ This contrasts with the contemporary conception of moral responsibility inspired by Strawson according to which it is possible to justify the

¹² Here I follow Fischer and Ravizza (1998: 5-8) in embracing a Strawsonian view of responsibility.

¹³ This quote is from the reprinted essay in P. F. Strawson (1993: 64), and it continues, as Strawson memorably quipped, "there still seems to remain a gap between [the framework of the reactive attitudes] applicability in particular cases and its supposed moral consequences. Sometimes [the compatibilist or incompatibilist] plugs this gap with an intuition of fittingness—a pitiful intellectual trinket for a philosopher to wear as a charm against the recognition of his own humanity." Thanks to Zimmerman for recalling this passage from Strawson.

reactive attitudes from a standpoint external to them. In this project, I embrace this contemporary Strawsonian conception of moral responsibility.

There are two main variants of the desert-based notion of moral responsibility at issue—accountability and attributability.¹⁴ Later I discuss these notions in more detail. Let me mention them here. Thinking of moral responsibility in terms of accountability involves focusing on whether it is fair, at least in principle, to hold a person accountable for their attitudes and actions.¹⁵ On most accountability views it is fair to hold a person accountable given certain historical conditions are satisfied. A popular historical requirement is that it must have been reasonable to expect the agent to have corrected the ignorance out of which she committed a wrongful act. By contrast, responsibility as attributability does not have a historical requirement. Even if an agent could not reasonably have been expected to rectify her ignorance this does not necessarily excuse her from moral responsibility. If the action aligns with what the person values and thinks good and true, then that action is attributable to the agent. She is morally responsible for it. It reflects her commitments at the time of action. What she could or could not have done in the past is not relevant on this conception of responsibility.

¹⁴ Attributability views are variously characterized. They are sometimes called “self-disclosure” views, “real self” views, or “quality of will” views. The basic idea is that moral blameworthiness rightly targets agents that bear a certain relation to their bad behavior. This relation is that the behavior reveals something of the agent’s mental life and what they value. Attributability theorists include Harry G. Frankfurt (1988), Robert Merrihew Adams (1985), and Nomy Arpaly (2003), to mention a few. Accountability views think attributability is not enough for responsibility. It must be fair to require the agent to give an account for her bad behavior. It must be fair to require the agent to align with the moral demands of her moral community. However, it can be unfair to consider an agent accountable for her bad behavior when her formative environment and conditioning made the behavior unavoidable or inevitable to some degree. Accountability theorists include Paul Benson (1987), R. Jay Wallace (1996), Fischer and Ravizza (1998), and Susan Wolf (1990), to mention a few.

¹⁵ Regarding the fairness requirement as involving the “adequate opportunity to avoid” something, both Watson (1996) and Scanlon (1998) endorse this idea, though differing slightly in details.

Historically, the literature on the epistemic condition on moral responsibility has focused on responsibility as accountability. However, this is changing. There is a shift toward attributability. This dissertation follows this pattern. I start out talking about the accountability notion and propose additional conditions for satisfaction of the epistemic condition that align with this notion of responsibility. But I argue that shifting to responsibility as attributability is fruitful. It affords better handling of skeptical threats to the effect the moral responsibility judgments lack truth-value. Attributability also reflects ordinary folk judgments of responsibility.

0.3 Ignorance as an Excuse Condition

Aristotle is commonly regarded as the first philosopher to give a theory of moral responsibility. In giving his account of virtue Aristotle discusses the preconditions of virtue. A precondition for virtuous action is voluntariness. If an agent involuntarily performed an action, then the agent is pardoned, excused, or pitied for performing the action. In such a case, the agent does not deserve praise or blame. Only actions that are voluntary are subject to the reactive attitudes. For Aristotle, actions that arise due to force or ignorance are involuntary actions. If I am in a parking lot and a strong gust of wind blows my car door out of my hand and into the door of another car, thereby damaging the other car, then a force external to myself caused the damage. I will owe compensation to the other person to cover the repairs and inconvenience, but I am not morally responsible for damaging the other person's car. Wholly external forces makes the action involuntary—if we're dealing here with an action at all—and it excuses me from moral responsibility.

Ignorance can also make action involuntary. Aristotle thinks lacking knowledge of what one is doing when one is doing it or lacking knowledge of the outcome for which one is doing it can make an action involuntary. An example of the former is when, “someone says that [the secret] slipped out while he was speaking” (*NE*, 1111a9).¹⁶ The person was unaware that he was revealing a secret with the speech content he mistakenly believed to be permissible to convey to others. An example of the latter is when we, “by giving someone a drink to save his life we might kill him” (*NE*, 1111a14). The person was unaware that giving the person a drink, which he mistakenly believed would save the person’s life, would result in the person’s death. When someone acts in ignorance of what he is doing (under some morally relevant description or conceptualization of his action) he is neither blameworthy nor praiseworthy for his action (under that description).¹⁷

In the contemporary literature, a distinction is drawn between performing an action that is morally wrong and being *blameworthy* for performing an action that is morally wrong. As Holly H. Smith (1983) explains, many moral philosophers draw a distinction between:

The objective *wrongness* of an action and the agent’s *blameworthiness* for performing it. No matter how terrible the act may be from an objective point of view, the agent is not blameworthy if he had an excuse for what he did. Thus a

¹⁶ I abbreviate *Nicomachean Ethics* ‘NE’. Quotes are from the translation by Terence Irwin (Aristotle 1985). The literature on moral responsibility largely assumes Aristotle’s use of “voluntariness” is translatable to the modern notion of “moral responsibility.”

¹⁷ It is debatable whether Aristotle thought moral ignorance exculpates. There are places in *NE* where he suggests all moral ignorance is culpable and places where he suggests moral ignorance can exculpate or render an action involuntary under some morally relevant description. Also, regarding factual ignorance, ignorance that excuses is ignorance of particulars under some relevant description or conceptualization of the action. For more on Aristotle and moral responsibility see Sauvé Meyer (2011) and Echeñique (2012).

doctor who treated a premature infant's respiratory distress in 1954 by exposing her to unnecessarily high concentrations of oxygen, and so caused severe eye damage, did something terrible. But if the doctor did not realize that oxygen enrichment would have this effect, he is not to blame for the baby's blindness. Ignorance of the nature of one's act is the pre-eminent example of an excuse that forestalls blame. (543)

Even if an act is morally wrong, an agent can escape being blameworthy for performing the action given that the agent was ignorant that he was doing something wrong. The doctor Smith mentioned did not recognize his act of treating a patient through oxygen enrichment under the description of exposing the patient to risk of severe harm. The doctor was ignorant of the circumstantial fact that exposure to elevated levels of oxygen can result in badly damaged eyes. However, such ignorance tout court may not excuse the doctor of blameworthiness for harming his patient. Ignorance of the circumstantial fact can be blameless or culpable. For instance, if the state of medicine in 1954 did not realize the link between oxygen enrichment and eye damage, then the doctor was blamelessly ignorant. However, if the medical literature in 1954 recognized this link, and the doctor should have kept abreast of such medical facts, then the doctor is culpably ignorant. Culpable ignorance may not constitute an excuse condition for wrongful action (i.e., action that is objectively wrong due to causing unnecessary harm to a person's welfare, unjustly violating a person's rights, and so on).¹⁸

¹⁸ H. Smith (1983: 548) notes that Aristotle echoes this sentiment concerning citizens and enforcers of law when he says, "ignorance is itself no protection against punishment if a person is thought to be responsible for his ignorance."

Aristotle pioneered an account of culpable ignorance. One's ignorance is culpable when one put oneself in a compromised epistemic position. Javier Echeñique (2012) provides a good summary of when a person is culpable for putting themselves in or being in a compromised epistemic position such that the person is culpable for the ignorance that results:

First, Aristotle makes it clear that what the agent is unaware of (both dispositionally and operationally) must be a 'necessary' piece of information...failing to know things one could reasonably be expected to know or anticipate...'reasonably' being defined from the perspective of the virtuous spectator....Another set of considerations that Aristotle uses to define 'culpably' resolves around the question of whether it was 'easy'...or 'not difficult'...to acquire or exercise a given piece of knowledge. Even if one is expected to advert to a certain piece of information, or to have acquired it in the first place, it may have been unreasonably difficult for one to do so, or even impossible....This is not all, for Aristotle also seems to assume that besides this 'formal' definition of culpability....there are certain causal factors corresponding to faults in the agent that account for his unfavorable epistemic position: these are the vices of impure agency. The culpable agent will be someone who, at the very least, will be blamed as 'careless', 'negligent', etc. because his culpable getting into the unfavorable epistemic position expresses a faulty disposition. (161)

Though Aristotle develops an account of culpability for being in a state of ignorance it is less clear that Aristotle embraces a tracing account of responsibility for acts done out of such ignorance. That is, it is an open question whether Aristotle thinks an agent

is doubly blameworthy—blameworthy for the act done from ignorance due to blameworthiness for getting into a state of ignorance. Some theorists argue that Aristotle holds that all acts done from factual error are involuntary. On this account, the agent is culpable for the ignorance but not culpable for the act done as a result of such ignorance. Others hold that, “in those cases where the error is not excusable, the act done as a result is not excusable either” (Echeñique 2012: 162). I will not attempt to resolve this dispute in Aristotle exegesis here. Instead, I note that Aristotle recognized that not all ignorance is blameless, and he proposed conditions for when ignorance is culpable.

To wrap up this section, the contemporary literature on the epistemic condition on moral responsibility investigates what conditions are necessary and sufficient for wrongful actions done out of ignorance to excuse an agent from being blameworthy for performing such actions. Moral responsibility, roughly put, involves being a legitimate target or candidate for moral appraisal on account of one’s actions. When one is not a candidate for the reactive attitude of blame for what one has done, as one has an excuse that prevents such a moral assessment from being legitimate, then one is not morally responsible for what one has done.

0.4 The Nature of Ignorance

The debate over the epistemic condition centers on cases involving ignorance. But, what is ignorance? And, what are the ways of being ignorant of a proposition?

There is some debate concerning the nature of ignorance. Some theorists embrace the common view that you are ignorant of a proposition when you fail to

know that proposition.¹⁹ This makes knowledge and ignorance complements. Knowledge and ignorance are, “mutually exclusive and exhaustive: for person *S* and fact *p*, either *S* knows that *p* or is ignorant that *p*” (Le Morvan 2011: 32). Someone is ignorant that *p* just in case *S* doesn’t know that *p*. However, other theorists have argued that the “ignorance = lack of knowledge” view is wrong. As Rik Peels (2014) points out there are several ways a person *S* might lack knowledge of a proposition *p*:

- (i) *p* is false;
- (ii) *S* disbelieves the true proposition *p*, that is, believes that $\sim p$ while *p* is true;
- (iii) *S* suspends belief on the true proposition *p*, that is, *S* has considered the true proposition *p*, but neither believes nor disbelieves that *p*;
- (iv) *p* is true and *S* neither believes that *p*, nor disbelieves that *p*, nor suspends belief on *p*;
- (v) *S* believes the true proposition *p*, but *S*’s belief that *p* lacks warrant, where warrant is that which turns true belief into knowledge. (484)

Are epistemic situations (i)-(v) all cases of ignorance? Arguably not. One can only be ignorant of true propositions. Take situation (i). Situation (i) does not indicate whether *S* has any attitude toward the false proposition. If I believe the falsehood that Kim Jong-un is dictator of South Korea, then I am not ignorant that Kim Jong-un is dictator of South Korea, as that proposition is false. However, (i) indicates that *p* is false. Let’s assume it is true that *p* is false. Then, I can be ignorant of this truth by

¹⁹ Theorists who implicitly or explicitly embrace ignorance as lack of knowledge include Driver (1989), Haack (2001), Zimmerman (1988), Houlgate (1968), Unger (1975), Fields (1994), and Anscombe and Morgenbesser (1963).

believing that p , suspending judgment on p , or having never considered whether p (i.e., neither believing that p , nor disbelieving that p , nor suspending judgement that p). When this occurs, I will be ignorant that p is false because I will lack true belief that p is false. Situations (ii)-(iv) involve ignorance. In such cases, there is a true proposition that one disbelieves, suspends belief concerning, or has never considered (i.e. neither believes, disbelieves, nor suspends belief on). However, epistemic situation (v) is debatable. Let me briefly say more about this.

Whether knowledge is the compliment of ignorance depends on whether situation (v) involves ignorance. If situation (v) involves ignorance, then all ways of lacking knowledge result in ignorance. If not, then at least one way of lacking knowledge does not result in ignorance. Some philosophers argue that ignorance is not lack of knowledge; rather, ignorance is lack of true belief.²⁰ Yet, as previously noted, many philosophers equate ignorance with a lack of knowledge. I am not defending a conception of the nature of ignorance in this project, so I will be ecumenical regarding the nature of ignorance and allow that (v) is a possible source of ignorance. This results in four ways that a person can be ignorant of a proposition:

Disbelieving Ignorance. S is disbelievingly ignorant that p iff (i) it is true that p , and (ii) S disbelieves that p .

Suspending Ignorance. S is suspendingly ignorant that p iff (i) it is true that p , and (ii) S suspends belief on p .

²⁰ Theorists who think of ignorance as the absence of true belief include Peels (2010), Peels (2011), Goldman (1986), Guerrero (2007), Rivera-López (2006), and Van Woudenberg (2009).

Deep Ignorance. *S* is deeply ignorant that *p* iff (i) it is true that *p*, and (ii) *S* neither believes that *p*, nor disbelieves that *p*, nor suspends belief on *p*.

Warrantless Ignorance. *S* is warrantlessly ignorant that *p* iff (i) it is true that *p*, (ii) *S* believes that *p*, and (iii) *S* does *not* know that *p*. (Peels 2014: 485)

It is worth keeping in mind the different ways one can be ignorant of a proposition. In most cases in the literature on the epistemic condition the protagonist in the cases has either *Disbelieving Ignorance* or *Deep Ignorance* of a true proposition. In many cases, the true proposition is that they are performing a morally wrong action.²¹ In the case of a wrongful omission the true proposition is that they are wrongfully omitting to do something, which is something they ought to have done. They believe that they are not performing a morally wrong action, or they believe that it is not wrongful to omit the performance of some act. Alternatively, deep ignorance occurs when the subject has no attitude toward the true proposition. The ignorance is deep because the subject has not considered the matter. This frequently occurs when moral ignorance is widespread and ingrained in cultural practices. The question of the moral status of one's wrongful actions and omissions in such a scenario may simply not arise. If the subject is blamelessly ignorant of the true proposition identifying an act's moral valence, and is either deeply ignorant or disbelievingly ignorant, then it is possible that such ignorance excuses her of the responsibility she would otherwise have for the action.²² Contrarily, some people argue that deep ignorance, which stems from moral ignorance

²¹ Ignorance of this proposition is moral ignorance. There is also factual ignorance, which is ignorance of the facts that make the action wrong. In chapter 3, I argue that moral ignorance is no excuse.

²² Peels (2014) argues for these claims. However, he argues that warrantless ignorance, even when the ignorance is blameless, never constitutes an excuse for moral responsibility.

due to widespread wrongful practices, never constitutes an excuse for moral responsibility.

CHAPTER 1: A SKEPTICAL CHALLENGE

1.0 Introduction

When a mother gives her infant a fatal dose of medicine it is natural to react by judging that the mother is blameworthy for the death of her child. Yet, imagine that the mother believed she was giving her infant medicine to treat an infection, not prescription narcotics that could poison the infant. Suppose, that is, that the mother acted from ignorance.²³ Now we can ask, was the mother culpable for her ignorance?²⁴ Imagine that right before administering the dose to her baby someone switched the liquid antibiotics with a visibly indistinguishable liquid form of painkillers. Given such a scenario, the mother's ignorance excuses her from moral responsibility for the death of her baby. Alternatively, the mother might be culpably ignorant. Imagine that the mother kept the bottle of prescription painkillers right next to the bottle of antibiotics and failed to check the label before administering her infant the dose. Among other things, the mother is blameworthy for failing to inform herself prior to administering the dose.²⁵ Adopting a term from Holly Smith's (1983) hallmark paper in the literature on culpable ignorance, a *benighting act* occurs when someone positively impairs or

²³ It is an interesting question what the difference is between acting *from* ignorance as opposed to merely acting *in* ignorance. Aristotle uses an example involving a drunk person to illustrate the difference (*NE*, 1110b25-30). Though a (very) drunk person may hurt another person in ignorance that they are doing so, it is not the ignorance that causes the drunk person to hurt the other person. It is the drunkenness. For Aristotle such an act is not rendered involuntary by the ignorance. The drunkenness caused the harmful action, not the ignorance. The contemporary literature on culpable ignorance does not discuss this distinction much. I take it this is because the scenarios in question involve ignorance as the (at least partial) cause of the wrongful act or omission. In this project I will not draw a hard distinction between when an act is done in ignorance as opposed to from ignorance. As a result, I will interchangeably refer to acts done in ignorance and from ignorance as acts that are a result of ignorance.

²⁴ I adopt the convention in the literature in using 'culpable' and 'blameworthy' interchangeably. I also use 'responsibility' and 'blameworthiness' interchangeably because we are focused on wrongful acts and the consequences of such acts. For more on these conventions, see Rosen (2004: 296).

²⁵ The mother would act in *moral* ignorance if she knew that she was giving her infant painkillers but she was unaware that it was wrong to risk grave harm to her daughter by giving her painkillers.

fails to improve their epistemic position.²⁶ This frequently occurs by failing to satisfy a duty of due diligence to gather readily available information relevant to one's action. The mother's benighting act of omitting to read the medicine bottle's label caused her to perform what Smith calls an *unwitting wrongful act*²⁷—unknowingly doing what it was wrong to do under the relevant description (i.e., administering prescription narcotics).

This chapter involves articulating the primary components of the skeptical challenge to moral responsibility from ignorance. I will discuss the distinction between circumstantial and moral ignorance, duties of inquiry, conditions on blameworthiness, Internalism and the regress of blameworthiness, the skeptical argument, and two responses to the skeptical challenge.

1.1 Factual and Moral Ignorance

The skeptical challenge from ignorance comes from cases where people do something wrong in ignorance. Such ignorance comes in two varieties: factual (aka circumstantial) ignorance and moral (aka normative) ignorance. You act in factual

²⁶ A benighting act is an impairment of one's epistemic position that can occur due to an act or an omission. An omission that impairs one's epistemic position is technically a benighting act. Forgetting to carefully look at the research is an omission to look at the research. But, in that it positively impairs or fails to advance one's epistemic state, it is also a benighting act. Thanks to Hieronymi for requesting this clarification.

²⁷ The literature assumes wrongness is objective. This is because the goal in the literature is to explain how even under the assumption that acts are objectively wrong agents may or may not be responsible for such acts depending on the status of the ignorance from which the act was done. Wieland (2017a: 6) explains subjective wrong versus objective wrong by way of an example, "[I]f a doctor was ignorant about the hazards of a certain drug, it would not be subjectively wrong for her to prescribe it. Indeed, if she believed that it was the best treatment option available, the subjectivist about wrongness might hold that she is even obligated to prescribe it, despite the unknown hazards. According to objective accounts of wrongness, in contrast, the wrongness of [the act] is not threatened by ignorance. On this view, the doctor's ignorance would not suffice to neutralize the wrongness of prescribing the hazardous drug. We will assume, with most participants in this debate, that some kind of objective account is true. Objective accounts allow for the cleaner separation of issues of wrongness and blameworthiness."

ignorance when you did not know that you were doing something that was wrong to do under the relevant description of that act. You act in moral ignorance when you knew what you were doing (under the relevant description), but you just did not realize it was wrong. A successful response to the skeptical challenge will address both types of ignorance. To illustrate the difference between types of ignorance consider the following variations on slaveholder cases.

Slaveholder 1: Subject S keeps slaves. She knows she could let them go, yet she does not do this because she hardly pays any attention to them, and is ignorant that they suffer.

Slaveholder 2: S keeps slaves. This time, she knows they suffer. Still she does not let them go because she is ignorant that she could run her business without them.

Slaveholder 3: S keeps slaves. This time, she knows they suffer, and that she could run her business without them. Still, she does not do this because she is ignorant that she is doing anything wrong. For all she knows, slavery is a given fact of nature, and if she was unlucky enough, she could have been a slave herself.²⁸

²⁸ These cases are from Wieland (2017a: 2). I thank Hanser for pointing out that the cases are under-described in important ways. The cases assume that factual ignorance about the suffering of slaves is leading to moral ignorance of the wrongfulness of the slaveholder's acts. We need to stipulate that she does not know that in cases 1-3 unwitting wrongful acts are committed. This is because ignorance of the slaves suffering may not entail ignorance of the wrongness of slavery. Slavery might be wrong because it is a violation of autonomy even if the slaves do not suffer at all. Also, it is unclear why the slaveholder in case 2 does not know that she can run her business without keeping slaves. She may not know that she ought to release the slaves so, given an ought-implies-can principle, she does not know that she can release the slaves. Or, there might be a psychological reason at play. Fear of poverty might prevent her from thinking too hard about things. She might know she ought to release the slaves but fail to draw the inference that she can release the slaves given her psychology.

All three slaveholders commit an unwitting wrongful act—the enslavement of persons—but they do so ignorant of different things. Slaveholder 1 and 2 exhibit circumstantial (or factual) ignorance. Slaveholder 1 is ignorant of the fact that her act of keeping slaves has the consequence of causing them suffering. Slaveholder 2 is ignorant of the fact that she does not need to keep slaves to run her business.²⁹ She does not know that her keeping of slaves is wrong because she does not know that she could have acted otherwise in running her business. Slaveholder 3 exhibits moral ignorance. She is ignorant of the moral wrongness of keeping slaves. Some cases involve a combination of moral and factual ignorance.

Let's focus on factual ignorance. There are several ways one might perform a benighting act that hinders one's epistemic position and causes an unwitting wrongful act. As a consequence of such a benighting act one's ignorance is culpable. H. Smith (1983) illustrates how culpable ignorance prevents ignorance from being an excuse forestalling blameworthiness by way of three cases involving factual ignorance. In these cases, the agents perform wrongful acts while falsely believing they are morally permissible or morally required acts.

Deficient Investigation: The doctor who failed to read his medical journal is an example of this kind of case: he ought to have read it, and if he had, he would

²⁹ It is worth also mentioning that this case put forward by Wieland (2017a) is under-described. Is Slaveholder 2 ignorant of the fact that she cannot run her business effectively, at current scale, or at all without the use of slaves? What sort of business is she running such that it could not exist without the use of slaves, at least in her mind? Wieland intended these to be toy examples illustrating a simple point, but there are various ways the cases are under-described. If anything hangs on the way they are under-described I will note that as we proceed. Otherwise, I will assume they illustrate different ways of being ignorant and performing wrongful acts from that ignorance.

have discovered the use of high oxygen concentrations to be unnecessarily harmful to the infant.

Preventing Subsequent Discovery: A person is slightly near-sighted, but not legally required to wear glasses while driving. Late for work one foggy morning, and unable to find her glasses quickly, she leaves home without them. Subsequently she swerves to avoid hitting a dog on her left, and seriously injures a child walking in the street on her right. Had she worn her glasses, she would have seen the child in time not to swerve.

Deficient Inference: On Monday a real estate agent tells her husband she will need the family car on Wednesday in order to show a client some property. Their conversation recedes from the husband's consciousness, and on Wednesday he does not ask himself whether his wife might need the car. Had he asked himself, he would have remembered her request. Not remembering, he takes the car rather than the bus to work, and the agent is forced to cancel her appointment. (H. Smith 1983: 544-5)

As discussed in the introduction to this dissertation, a person is ignorant when they lack a true belief or knowledge. The cases above illustrate different ways agents can poorly gather and handle information such that, at the time of action, they are ignorant of the fact that they are doing something wrong. For instance, in the Deficient Inference case, the husband neglected to infer that he should not take the car to work because his wife needs the car on that day. His forgetting the conversation with his wife where she told him she will need the car on Wednesday is the benighting act. It resulted in his failing to improve his cognitive position and infer from his background

beliefs that he ought not take the car to work—that it is wrong of him to take the car to work. Additionally, the husband’s cognitive slip is not excusable on other grounds, such as it being the first time he forgot to query himself whether his wife needs the car, where this is the natural result of his aging. Instead, his wife would justly judge that her husband should have thought about whether she may need the car and remembered the conversation about her needing it. If the husband is culpable for the benighting act, then, according to H. Smith (1983: 547), the husband is culpable for the consequences of the act done in ignorance—the unwitting act. But this raises the question: why are benighting acts wrong? For instance, why was it wrong that the doctor failed to read the relevant medical journal, the driver did not wear her glasses, or the husband neglected to remember that he ought not take the car?

1.2 Duties of Inquiry

Benighting acts are wrong because they violate duties of inquiry.³⁰ Smith’s cases involved deficient investigation, deficient inference, and preventing subsequent discovery. The agent’s performed acts or omitted to perform acts that resulted in impairment of their ability to attain true beliefs concerning the wrongness of their acts. They impaired their epistemic position (i.e., their position with respect to the attainment of truth and knowledge). We can group these different ways of impairing one’s epistemic position under the violation of duties of inquiry. But, why focus on violation of duties of inquiry? Why not be more direct and claim that the agents

³⁰ This claim is made by Wieland (2017a) in his introduction to the state of the art on the epistemic condition. I follow Wieland on this point and give some reasons for thinking it is the way to go. Should a better approach to capturing this emerge, I remain open-minded in that regard. I do not intend much to hang on endorsing duties of inquiry as the best way to explain why benighting acts are wrong. Thanks to John Greco for requesting clarification of the status of this claim.

violated duties of belief? The agents ought to have believed that they were doing something wrong; they are guilty of failing to have beliefs that they ought to have.

A common reason for focusing on duties of inquiry as opposed to duties of belief in assigning culpability for unwitting wrongful acts is that belief-formation is not under direct voluntary control. If belief-formation were under direct voluntary control, an agent could decide to believe something and, on that basis alone, come to believe that thing. Most epistemologists do not think we can form beliefs at will in such a manner.³¹ Thinking that it is most reasonable to hold people accountable for things they can readily control prompts the idea that it is not most reasonable to hold people accountable for having or lacking particular beliefs.³² Instead, it is reasonable to hold people accountable for doing or failing to do things that can indirectly impact the beliefs they come to have. The notion of inquiry involves mental or physical acts people have direct voluntary control over that can indirectly influence the beliefs agents come to have. Such acts include reflection, reasoning, introspection, and evidence gathering. Belief-forming agents have duties of inquiry. They are reasonably held responsible for gathering or failing to gather evidence, reflecting or failing to reflect, and so on. What is the nature of such duties of inquiry?

Inspired by a suggestion made by Holly H. M. Smith (2014) that duties of inquiry are objective and derivative, Wieland (2017a: 9) observes:

³¹ For defense of belief-formation being voluntary see Weatherson (2008) and Steup (2008). For an argument against doxastic voluntarism see McHugh (2011), and for an account of why we are unable to believe at will see Hieronymi (2006) and Hieronymi (2009).

³² This assumption will not necessarily be endorsed by theorists that focus on moral responsibility in terms of attributability or answerability. Some theorists in this camp think the involuntariness of belief-formation is an important aspect of what makes us responsible for the things we believe. It is because our beliefs cannot be changed at will that they are rightly regarded as reflecting our commitments or the things we value. For more on this perspective, see Hieronymi (2008).

Pharmaceutical company directors, doctors, engineers, and parents have the duty not to expose certain classes of people to avoidable risk of harm, and the fact that performing certain inquiries would allow them to satisfy this duty is what grounds a duty to carry out these inquiries. In other words, it would be wrong for these agents to expose others to an avoidable risk of harm and to fail to perform prior actions that would diminish this risk of exposure. In short: *the benighting [act] is wrong because the unwitting [act] is wrong*. The wrongness of a doctor's failure to read her journals is derived from the wrongness of prescribing harmful drugs. Cases of moral ignorance are subject to a similar treatment. It is wrong for the slaveholder to fail to question her practices because it is wrong to keep slaves.

An agent performing a benighting act violates a duty of inquiry. Failing to improve one's cognitive position might involve failure to gather relevant evidence, failure to scrutinize one's beliefs or evidence through reflection, introspection, or reasoning. The wrongness of such a benighting act stems from the wrongness of the unwitting act.³³ Why is this the case? The duty of inquiry stems from the fact that satisfying that duty will facilitate satisfying other duties. Satisfying the doctor's duty to read the relevant research will facilitate the doctor satisfying the duty to help patients and not expose them to avoidable risk of harm. We can make this more precise:

Derivativeness of Duties of Inquiry: S has a duty to inquire D1 because she has some other duty D2 that is such that conforming to D1 will enable her to see

³³ Recall that a goal in this literature is to explain how objective wrongness can come apart from blameworthiness. Given that we are considering wrongful actions in the objective sense, the benighting act is objectively wrong due to inheriting its wrongness from the wrongness of the unwitting act, which is by stipulation objectively wrong.

that she has D2 (or will enable her to see *how* she can conform to D2). (Wieland 2017a: 8)

We can relate this back to the relationship between benighting acts and unwitting wrongful acts. Benighting acts are wrong because unwitting wrongful acts are wrong and performing a benighting act, such as failing to satisfy one's duty to inquire, puts one in a position such that one cannot see that the unwitting act is wrong. What type of duties are these duties?

The duties or obligations at issue are epistemic obligations that receive their authority from their connection to moral duties. They are epistemic obligations that are part of a standard of due care or, as Gideon Rosen (2004) classifies them, they are part of taking "precautions against negligent harm." For Rosen the duties are instantiated with reference to the person of reasonable prudence:

As you move through the world you are required to take certain steps to inform yourself about matters that might bear on the permissibility of your conduct. You are obliged to keep your eyes on the road while driving, to seek advice before launching a war and to think seriously about the advice you're given; to see to it that dangerous substances are clearly labeled, and so on. These obligations are your procedural epistemic obligations. Again, they are impossible to codify. But again, the person of ordinary prudence provides a serviceable heuristic. In any given case we can ask whether the agent's ignorance derives from a failure to do what any reasonably prudent person in his circumstances would have done in order to see to it that he was adequately informed...these procedural obligations are always 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 direction, etc. The procedural obligation is not itself an obligation to know or believe this or that. It is an obligation to take steps to ensure that when the time comes to act, one will know what one ought to know. (Rosen 2004: 301)

Duties of inquiry (or procedural epistemic obligations) are obligations that serve moral ends and put one in a position to act in an informed manner at the time of action. Violating these epistemic obligations, which derive their wrongness from their connection to moral obligations, results in one being culpably ignorant. This analysis raises an important question: What necessary and sufficient conditions hold concerning blameworthiness for benighting acts, which involve violations of duties of inquiry, and blameworthiness for subsequent unwitting wrongful acts?

1.3 Conditions on Blameworthiness

Orthodoxy on culpable ignorance and moral responsibility endorses two general conditions concerning blameworthiness for unwitting wrongful acts.³⁴ The first condition indicates that being blameworthy for ignorance that the unwitting act is wrong is necessary for being blameworthy for the unwitting act. The second condition indicates that being blameworthy for the benighting act that resulted in the unwitting act is necessary and sufficient for being blameworthy for ignorance that the unwitting act is wrong. More explicitly, the conditions are:

³⁴ Proponents of such an orthodox stance include Michael Zimmerman (1997), Ginet (2000), Rosen (2004), FitzPatrick (2008), and Levy (2009).

(I) S is blameworthy for the unwitting wrongful act only if S is blameworthy for her ignorance that the unwitting wrongful act is wrong; and

(II) S is blameworthy for her ignorance that the unwitting wrongful act is wrong iff S is blameworthy for a benighting act (at least one such act) that led to the unwitting wrongful act.³⁵

Applying (I) and (II) to the case of the mother who gave her infant a pill unwitting of the wrongness of doing so, as unbeknownst to her the pill contained a fatal dose of a narcotic painkiller, we get: the mother is blameworthy for giving her infant the pill only if the mother is blameworthy for her ignorance that giving her infant the pill is wrong³⁶; and, the mother is blameworthy for her ignorance that giving her infant the pill is wrong just in case she is blameworthy for at least one benighting act (e.g., failing to read the label prior to administering the medicine, or keeping the infant's medicine too close to her pain medicine) that resulted in giving her infant the pill.

There are further orthodox commitments in the literature. It is important to specify these orthodox commitments. When people argue against orthodoxy it helps to bear in mind what claims they are rejecting. These commitments are related to conditions (I) and (II), but they are not equivalent to them. These commitments focus on what it takes to establish derivative or indirect blameworthiness, which is when

³⁵ Wieland (2017a: 9) proposes these principles. He uses abbreviation 'A1' for the benighting act, and abbreviation 'A2' for the unwitting act. I opted to not use the abbreviations to avoid confusion.

³⁶ Recall the mother is ignorant of the circumstantial fact that she is not giving her infant medicine but is instead administering a high dose of a narcotic.

blameworthiness for the unwitting act derives from blameworthiness for the benighting act. Here are the conditions.³⁷

Transfer Condition: If an agent is blameworthy for a benighting act (or the ignorance that results from such an act), and the agent has no further excuses for the unwitting wrongful act, then the agent is blameworthy for the unwitting wrongful act (or a consequence of the act).

Necessary Condition: An agent is blameworthy for an unwitting wrongful act only if an agent is blameworthy for a benighting act.

Explanatory Condition: An agent is blameworthy for an unwitting wrongful act (partly) because an agent is blameworthy for a benighting act.

Parity Condition: The Transfer Condition, Necessary Condition, and Explanatory Condition apply to both factual ignorance and moral ignorance.

Responsibility as Accountability Condition: The conditions above apply to responsibility as accountability.³⁸

According to the Transfer Condition, blameworthiness for a benighting act (or the resulting ignorance) can transfer to blameworthiness for an unwitting wrongful act (or a consequence of that act). This connects to tracing conditions on moral responsibility. If the Transfer Condition is true, blameworthiness for an unwitting wrongful act traces back to blameworthiness for a benighting act. For this transfer of

³⁷ These conditions come from Wieland (2017a). Wieland is the first to systematize them. I have adapted the claims by placing labels on them and putting them in descriptive language instead of the abbreviated symbols used by Wieland.

³⁸ Responsibility as accountability is a species of a desert-entailing conception of responsibility. Roughly, a person deserves blame for performing an action given that it is appropriate to hold the person accountable for their action by way of the reactive attitudes. Thanks to Greco for requesting clarification of this point.

blameworthiness to occur the Necessary Condition and Explanatory Condition must hold. That is, blameworthiness for a benighting act is both necessary for and explanatory of blameworthiness for an unwitting wrongful action.³⁹ The Transfer Condition does not automatically fall out of conditions (I) and (II).⁴⁰

By contrast, the Necessary Condition falls directly out of conditions (I) and (II).⁴¹ The Necessary Condition implies how ignorance can serve as an excuse condition for responsibility for a wrongful action. If the agent is not blameworthy for a benighting act, then the agent is not blameworthy for the resulting unwitting wrongful act. The ignorance is blameless and thereby qualifies as an excuse condition.

The Explanatory Condition occurs because not all necessary conditions are explanatory conditions. Additionally, blameworthiness for a benighting act only partially explains blameworthiness for an unwitting wrongful act, as additional conditions may explain blameworthiness, such as a control condition or an autonomy/authenticity condition.

³⁹ For more on the Transfer Condition, see Robichaud and Wieland (2017: Ch. 16).

⁴⁰ Upon closer inspection, it is unclear that the Transfer Condition is derivable from (I) and (II) without an assumption that reverses the necessity found in principle (I). That is, an agent being blameworthy for an unwitting act, such as prescribing the drugs, is necessary for that agent being blameworthy for the *ignorance* that the unwitting act is wrong. Such a principle seems necessary to derive the Transfer Condition from (I) and (II). Taking the right-to-left direction of the biconditional in (II) and this alternative necessity condition, plus performing a disjunctive syllogism, results in deriving the conditional found in the Transfer Condition. It is possible to endorse this alternative necessity principle or reject the Transfer Condition. This is to deny that blameworthiness for the unwitting act must trace to blameworthiness for a benighting act. I explore this approach later in the dissertation and allow original blameworthiness to terminate at blameworthiness for the unwitting act. It need not always inherit its blameworthiness from some other act that explains it or to which it can be traced. Wieland is the first to systematize these conditions, and it is unclear that many people have thought about the connections between (I) and (II) and deriving the additional conditions. It may turn out that all such conditions are not properly classified as orthodoxy. I thank Hanser and Hieronymi for urging me to consider the connections between (I) and (II) and the further conditions.

⁴¹ Using the left-to-right direction of the biconditional in (II) coupled with the conditional in (I) results in the conditional found in the Necessary Condition.

The Parity Condition indicates that the above conditions all apply to factual ignorance and moral ignorance. Wieland notes that, “One rationale for parity is the existence of duties to inquire into factual and moral issues. If factual ignorance traces to failures to inquire into factual matters, why not think that moral ignorance traces to failures to inquire into moral matters?” (Wieland 2017a: 11).

Lastly, the Responsibility as Accountability Condition indicates that all the conditions apply to the kind of responsibility referred to as accountability. On this understanding of responsibility, if an agent is blameworthy for a benighting act, then the agent could have reasonably been expected to have avoided performing the benighting act. Likewise, if an agent is blameworthy for an unwitting wrongful act, the agent could have reasonably been expected to have avoided performing the unwitting wrongful act. Recalling the slaveholder cases, “the slaveholder is blameworthy to the extent that she could reasonably have been expected to question her practices and, indeed, to stop them” (Wieland 2017a: 11). Responsibility as accountability tracks our reasonable expectations of the agent in light of the evidence the agent had (or could have easily had), the difficulty and complexity of inquiry and information gathering, and the pursuant difficulty of forming and revising beliefs.

1.4 Internalism and the Regress of Blameworthiness

Internalism is the view that blameworthiness for unwitting wrongful acts must trace back to, or find its origin in, witting wrongdoing (i.e., acts accompanied by the belief that one is doing wrong all-things-considered). The view is internalist because it requires the agent to mentally represent the wrongness of the act or omission for the

act or omission to generate culpable ignorance.⁴² This idea is captured by Michael Zimmerman's Origination Thesis:

Origination Thesis: Every chain of culpability is such that at its origin lies an item of behavior for which the agent is directly culpable and which the agent believed, at the time at which the behavior occurred, to be overall morally wrong. (Zimmerman 1997: 418; 2008: 176)⁴³

The Origination Thesis involves commitment to a tracing condition and what is called "clear-eyed akrasia." Regarding tracing, culpability for an act must always terminate in, or trace back to, a point in the chain of culpability where direct culpability is established. An agent is directly blameworthy for something when she is not blameworthy for that act or omission in virtue of being blameworthy for something else. On the other hand, derivative blameworthiness involves blameworthiness for one thing in virtue of being blameworthy for something else. So, the tracing condition endorsed by the Origination Thesis requires blameworthiness to always terminate in acts or omissions that one is directly, as opposed to derivatively, blameworthy for. The Origination Thesis is internalist in that it limits what qualifies as an act or omission for which an agent is directly blameworthy to acts involving clear-eyed akrasia. Such

⁴² This contrasts with the externalist who does not exclusively require the agent to mental represent—in the form of awareness or belief—the wrongness of the act or omission. The externalist allows that the agent may have no representation of the wrongness of the act or omission yet the act or omission may still be the terminus of culpability.

⁴³ For internalists like Zimmerman the Origination Thesis is an advance over the benighting requirement put forward by H. Smith (1983). Zimmerman points out that Smith, "fails to acknowledge that culpability requires, at bottom, a belief concerning wrongdoing" (1997: 417-8 n.12). Externalists deny this point. They think Smith's benighting requirement captures an important range of ways culpability can arise. Externalists grant that Zimmerman's doxastic condition (i.e., when one has a belief that one is doing wrong all-things-considered) is sufficient for culpability, but externalists hold the doxastic condition is not a necessary condition—as culpability can originate in other ways.

akrasia requires a belief that one is doing wrong at the time that one is doing it, and it requires having the facts of the situation in full view.⁴⁴ The akratic agent has such a belief that they are doing wrong, but they commit the act or omission anyway. Such an agent acts contrary to her all-things-considered judgment of what she morally ought to do. Thus, the Origination Thesis requires culpability to trace back to, or ground out in, an act of clear-eyed akrasia—an act the agent is directly culpable for.

We can apply the Origination Thesis to blameworthiness for benighting acts transferring to blameworthiness for unwitting wrongful acts. When looking at benighting acts such as those involving deficient investigation or deficient inference the agent must believe that it is all-things-considered wrong to omit to gather information or it is wrong to fail to make an inference supported by his or her background beliefs. In the absence of this belief in the all-things-considered wrongness of failing to discharge her epistemic duties, there would be no blameworthiness to transfer from the agent's benighting acts to his or her unwittingly wrongful acts.

Internalism coupled with orthodox commitments (I) and (II) generates a regress of blameworthiness. I will illustrate the regress using the case of the mother who killed her infant by unwittingly administering a high dose of a narcotic:

- (1) The mother is blameworthy for giving her infant the pill only if

⁴⁴ This includes the relevant moral and non-moral (descriptive) facts concerning what a person is doing. Typically, it requires awareness of such facts or beliefs concerning what one is doing.

(a) The mother believed that giving her infant the pill was wrong, or (b) The mother was then blameworthy for her ignorance that giving her infant the pill was wrong;

Condition (b) holds only if

(2) The mother is blameworthy for a past omission (e.g., failing to look at the label on the medicine bottle) that resulted in her ignorance of the wrongness of her act at the time she performed it;

Condition (2) holds only if

(c) The mother believed at the time of her omission that this omission was wrong (because a dereliction of her epistemic duties), or (d) The mother was then blameworthy for her ignorance of the wrongness of this omission;

Condition (d) holds only if...

The regress preys on the internalist commitment that blameworthiness must terminate in an akratic act for which the agent is directly culpable.⁴⁵ Though the regress is not vicious, as it can terminate in satisfaction of conditions like (a) or (c), it

⁴⁵ Internalism adds conditions (c) and (d) to the requirements for blameworthiness. However, condition (d) (i.e., the mother is blameworthy for her ignorance that the past omission is wrong) will only hold given that the mother believes that her ignorance that the past omission is wrong is wrong (what we might call "(f)", or the mother is blameworthy for her ignorance that her ignorance that the past omission is wrong is wrong. This chain of ever hard to track culpability keeps regressing because internalism requires writing wrongdoing to halt the regress. As a preview, Rosen's skeptical challenge arises because he embraces the internalist requirement to halt the regress and argues that it is very difficult from a first-personal or third-personal perspective to determine if, at the time of action, the person believed that what they were doing (or about to do) was wrong. He thinks it is rare that people's wrongful acts done in ignorance are accompanied by clear-eyed akrasia. Even if they are so accompanied, this is very difficult to determine given considerations of the opacity of mental states and so on. Thus, he doubts that we are ever in an epistemic position that affords warranted ascriptions of responsibility, which is the skeptical challenge to responsibility. Thanks to Pamela Hieronymi for urging me to clarify internalism and the regress.

can lead to a skeptical result, as I will show in the next section. So, halting the regress is important for avoiding skepticism about moral responsibility.

Recall that in introducing the regress I said that, “Internalism coupled with orthodox commitments (I) and (II)” generates the regress of blameworthiness. The questions naturally arise: Why is this a *regress* for Internalism in particular; and what bit is Internalism adding to generate the regress? Doesn’t the regress just fall out of orthodox commitments (I) and (II)?⁴⁶

Let me start by answering the question concerning what work Internalism is doing in generating the regress. In the regress, orthodox commitments (I) and (II) get us step (1), including (a) and (b), and step (2), minus (c) and (d). Commitment (II) does not specify what is required to make step (2) hold. Internalism adds (c) and (d) to indicate this, namely that the origin of direct blameworthiness for a benighting act is discovered when such an act occurs accompanied by a belief that one is doing wrong all-things-considered in committing the benighting act. This bit is added by the internalist embrace of the Origination Thesis—all chains of culpability must terminate in an act for which one is directly culpable, where direct culpability holds just in case, “the agent believed, at the time at which the behavior occurred, to be overall morally wrong” (Zimmerman 1997: 418; 2008: 176).

Now that it is clearer what Internalism adds to (I) and (II) to generate the regress, we are in a better position to indicate why this is a *regress* for Internalism in particular. Though the regress is not vicious, it is a regress considering the difficulty of locating—either from a first-personal or a third-personal perspective—whether a

⁴⁶ I thank Hieronymi for comments prompting these questions and the need for clarifications.

benighting act was an akratic act. The search for blameworthiness in the chain is apt to keep regressing given Internalism's narrow requirement on what counts as an act capable of being the origin of direct responsibility.⁴⁷

1.5 A Skeptical Challenge

In the prior section, we discovered that Internalism's commitments can lead to a regress of blameworthiness—a continuing search for an act of direct blameworthiness (i.e. clear-eyed akrasia), pushing the search further and further back in the chain of culpability. Those same internalist commitments that allow for such a regress can generate a skeptical challenge to moral responsibility.

Rosen (2004) appropriates Zimmerman's Origination Thesis to generate a skeptical challenge concerning judgments of moral responsibility in any particular case.⁴⁸ Recall, there is a distinction between original (or direct) responsibility and derivative (or indirect responsibility). According to the Origination Thesis, the origin of any sequence of culpability always terminates in an act of clear-eyed akrasia. Such akrasia occurs when one knows that the balance of one's reasons count against doing

⁴⁷ One might also wonder whether Internalism and its regress are equivalent to a tracing requirement, such that they should just be called a tracing requirement on blameworthiness? In short, I think Internalism and its regress are species of a tracing requirement, but I do not think they should simply be called a tracing requirement. Simply calling Internalism and its regress a tracing requirement places Internalism and the regress at the wrong level of generality, as the term "tracing requirement" is a general term that does not specify the content of when a previous act can secure blameworthiness for a subsequent act. Internalism, and its regress, specifies this content in terms of a propositional attitude the content of which is the wrongness of one's present act (i.e., that the act I am committing is morally wrong overall). Thanks to Hieronymi and Hanser for comments prompting this clarification about Internalism, the regress, and tracing requirements.

⁴⁸ It is worth noting that this is not a skeptical argument that moral responsibility is impossible. For such an impossibility argument see Galen G. Strawson (1994), and for an explanation why the impossibility argument fails see Randolph Clarke (2005). Rather, the claim is that even if moral responsibility is possible we never have a sufficient degree of confidence that warrants us in making judgments of moral responsibility in particular cases. This is also not skepticism concerning morality because the skeptical argument assumes common sense morality is roughly correct.

act A, and one judges that one ought not A, but one does A anyway. Rosen argues that identifying acts of clear-eyed akrasia in the sequence of culpability is extremely difficult. This is because the contents of one's mind or another person's mind is often opaque to introspection and reflection. To illustrate, given the opacity of the contents of someone's mental states, concerning someone who unwittingly broke a promise through lying it is unreasonable to judge that, "at the time of action, either he knew that he had decisive reason not to lie, or if he did not know this, that his ignorance was the upshot of some prior bad action done in full knowledge of every pertinent fact or norm" (Rosen 2004: 308). As Rosen explains:

Sometimes the relevant facts will be straightforwardly inaccessible. The surgeon neglects to check her patient's chart; this is an act done from ignorance. Is she culpable for the ignorance? That depends on whether it derives from some prior culpable act or omission. But how are you (a third party) supposed to approach that question? Perhaps you can identify the prior omission—perhaps you can show that given her track record on such matters, she should have asked a colleague to remind her to check the chart. Still, the question will be whether her failure to do so was itself a culpable failure; and that will depend on whether this episode was itself an akratic episode—that she failed to ask for the reminder even though she knew that she ought to ask for it—or traces back to an akratic episode. And it seems obvious that in any real case it will be impossible to resolve this question with any confidence.⁴⁹

⁴⁹ Rosen also points out that it will often be difficult to tell a genuine act of akrasia apart from a standard act of weakness of will:

It is possible to reinforce this line of thought with a more general ground for doubt. The agent is culpable for his bad action only if that bad action is, or derives from, an episode of genuine

Rosen's skeptical challenge to moral responsibility is somewhat qualified. It is skepticism about attributions of responsibility, not about the possibility of responsibility itself. It is a claim that we, as a third party, can never be sufficiently confident that akratic conditions are satisfied. The regress must bottom out in an act of clear-eyed akrasia for us to establish blameworthiness for the act in question. But given that we cannot get inside the head of the person at the time they made the decision and given that first person reports concerning whether at the time of action the person believed that what they were doing was wrong are unreliable, we can never possess justified judgments concerning whether a person was morally responsible for what they did (or failed to do). We can never justifiably halt the regress by forming a justified judgment that the chain has bottomed out in an act of clear-eyed akrasia.

William FitzPatrick (2008) precisifies Rosen's skeptical challenge. FitzPatrick more explicitly details the propositions involved in the skeptical challenge and how those propositions generate the skeptical conclusion.⁵⁰

(1) An agent S is responsible for a wrong act A only if either

akrasia. But genuine akrasia in this sense is extremely difficult to identify. The reason is that it is not readily distinguishable from an impostor: ordinary weakness of will. The akratic agent judges that A is the thing to do, and then does something else, retaining his original judgment undiminished. The ordinary moral weakling, by contrast, may initially judge that A is the thing to do, but when the time comes to act, loses confidence in this judgment and ultimately persuades himself (or finds himself persuaded) that the preferred alternative is at least as reasonable. Moreover, in between these two pure cases there is a continuum of cases; cases in which the agent suspends his original judgment without quite rejecting it; or cases in which it is simply indeterminate as the agent acts whether he in fact believes that all things considered he should do A. (Real action almost never involves an explicit practical judgment; and it is not hard to imagine that it may be indeterminate whether, in the absence of such a judgment, the agent is to be credited with an implicit belief that A is to be done.) (2004: 308-9)

This is another reason to maintain skepticism concerning judgments of moral responsibility.

⁵⁰ I present FitzPatrick's specification of Rosen's skeptical argument in slightly modified form. I do this to make more obvious the connection between the argument and the regress of blameworthiness facing Internalism. Though, as Greco rightly points out, the full logical structure of the argument has not been laid bare. Laying bare the full logic of the argument is a worthwhile future project.

(a) S is originally responsible for A, or (b) S is derivatively responsible for A by virtue of being originally responsible for something else that led to A.

(2) Thus, S is responsible for A only if either

(c) A itself is the origin of original responsibility, or (d) there exists an origin of original responsibility in A's causal history.

(3) Act A is an origin of original responsibility (such as case (2c) above) only if

(e) S knows the balance of reasons against doing A; ignorance of this—whether due to circumstantial ignorance or to normative ignorance—removes original responsibility for A.

(4) Thus, S is originally responsible for A only if his action is a case of clear-eyed akrasia (i.e., acting against his considered judgment about what there is most reason for him to do).

(5) If instead S is ignorant of the balance of reasons against doing A, and hence is not originally responsible for A, then S may still be derivatively responsible for A (case (2d) above) but only if S is culpable for the relevant circumstantial or normative ignorance by being originally responsible for whatever led to that state of ignorance.

(6) But, as in (3), S will be originally responsible for what led to his ignorance only if this amounted to a knowing failure to fulfill certain procedural epistemic duties—that is, knowing “negligence or recklessness in the management of his

opinion,” in this case related to securing knowledge of the balance of reasons against doing A.⁵¹

(7) So the only way for S to be responsible for what led to the ignorance that resulted in A would again be for S to have been involved in a form of clear-eyed akrasia in connection with the relevant epistemically debilitating behavior—that is, knowing that he had most reason to fulfill certain epistemic duties and yet failing to do so, knowingly being negligent in the management of his opinion.

(8) Thus, if S is responsible for A, then either A is itself a case of clear-eyed akrasia or it results from such akrasia associated with A’s causal antecedents.

(9) But it is not possible for us to know in any particular case whether such clear-eyed akrasia is really involved in the etiology of the action.

(10) Therefore, it is not possible to know in any particular case whether S is truly responsible for A, and we should thus suspend judgment about it.

The skeptical argument concludes that suspension of judgment is warranted concerning responsibility for acts. This includes responsibility for unwitting wrongful acts, but it is not limited to them. Yet, given the scope of the argument, in any specific case where we are trying to judge that an agent is morally responsible for an unwitting wrongful act we cannot know that the agent is responsible for that act.⁵²

⁵¹ Otherwise, as FitzPatrick indicates, “if [S] has been duly thoughtful and reflective all along, and his ignorance is merely a result of poor available information, bad upbringing, or being in the grip of a false normative view despite his best efforts, then the ignorance leading [S] to do A isn’t his fault: he blamelessly believes what he believes” (FitzPatrick 2008: 592).

⁵² There are various ways of indicating why we lack such knowledge. One way of explaining this, which aligns with Rosen’s comments, is that we cannot have sufficient confidence in our judgments of responsibility, where such confidence (all else being equal) licenses knowledge or puts us in a position

1.6 Two Responses to the Skeptical Challenge

In this section I present two ways of responding to the skeptical challenge. These responses are united in holding that blameworthiness for unwitting wrongful actions need not always bottom out in clear-eyed akrasia. The responses diverge concerning their underlying commitments, but they ultimately respond to the skeptical challenge in a similar way from a structural perspective (i.e., they both deny the same step in the argument). Yet, there are differences between the views when it comes to how they handle cases involving moral ignorance. Given this, I think it best to start with externalism and end up embracing attributionism, as I think attributionism better handles cases involving moral ignorance.⁵³

The first response strategy is the “quality of will” or “attributionist” approach. On this approach, agents are directly responsible for unwitting wrongful acts when those acts are attributable to them.⁵⁴ This occurs when an act is reasonably regarded as expressing a person’s practical identity or quality of will, which occurs when an act

to know that a person is responsible for a specific act in a specific case. Our judgments lack warrant and suspension of judgment is the epistemically appropriate attitude to take toward judgments of responsibility. We lack an epistemic right to positive judgments of responsibility. Yet, contrary to this, we seem to make warranted judgments of moral responsibility all the time, which is a reason why the skeptical challenge is important to resolve. It conflicts with our practices of holding each other accountable, and such accountability, as a version of a desert-entailing conception of responsibility, should account for the social practices that motivate that conception of responsibility. Another possibility is to be a revisionist about our ordinary practices of holding each other accountable, but that is not the position currently on the table in the debate over the epistemic condition.

⁵³ This is what I end up doing in this dissertation. I modify externalism in response to challenges. Then I embrace attributionism when focusing on cases involving moral ignorance.

⁵⁴ There are different ways of looking at what this means regarding whether the view is an accountability view. It is not a standard (volitional) account of moral responsibility. It allows non-volitional conditions to secure responsibility. Fairly minimal conditions of attributability are enough to secure accountability. Attributability is sufficient for accountability. Does this make it an accountability view? Again, it depends how you look at things. The literature on the nature of moral responsibility is a thicket of competing conceptions of moral responsibility. An alternative way of looking at attributability is that it is not an accountability view because it is not a volitional accountability view. It is something else. Being a hybrid view makes it a new view. Regardless of how things are carved up or characterized, the main thing is that the attributionist account in the literature on the epistemic condition differs from the accountability views that are endorsed by orthodoxy in that literature.

reflects what a person cares about, values, and judges worthwhile as ends to pursue. This view denies step (3) in the skeptical argument. Ignorance of the balance of reasons against performing an action does not always remove original responsibility for the action. When it comes to moral (or normative) ignorance an unwitting wrongful act is an origin of original responsibility when its performance stems from a poor quality of will that reflects a “lack of concern for the features that make [the acts] wrong.”⁵⁵ For instance, the slaveholder that knows all the descriptive facts relevant to her keeping of slaves (e.g., she knows they suffer, and that she could run her business without them) but is ignorant of the moral facts because she believes a false narrative concerning the naturalness of slavery (e.g., that slavery is a fact of nature and is permissible for that reason) is originally responsible for her keeping of slaves because that action reflects a lack of proper concern for what makes slavery wrong, namely that slaves suffer and it deprives them of their autonomy.⁵⁶ On the attributionist approach there is no need to search for an origin of original responsibility in what led the agent to her moral ignorance. There is no need for scouring the casual history for a source of responsibility capable of being the origin of original responsibility for the act. Yet, this view allows that when an agent is ignorant of the balance of reasons against doing

⁵⁵ This is how Wieland (2017a) sketches the right side of the biconditional in the Quality of Will view. Later I will endorse a modified version of this Quality of Will brand of attributionism. Instead of focusing solely on moral concern and how that reflects a person’s quality of will, I will allow rationalistic factors to also count as sufficient for grounding blameworthiness. It is what I call the Quality of Will and Judgment view.

⁵⁶ This view makes room for inverse akrasia, which occurs when one believes that one is doing the wrong thing all-things-considered, but one is actually doing the morally right thing. Huck Finn type cases are often used to capture the phenomenon of inverse akrasia. Huck Finn believes that all-things-considered he ought to turn in runaway slave Jim, but he acts against his better judgment and does the right thing. Jonathan Bennett (1974) and Nomy Arpaly (2003) defend the view that Huck is praiseworthy even though he acts against his better judgment in doing the right thing. For Arpaly, Huck’s action reflects responsiveness to reasons and a proper concern for morality. Being responsible on such a model does not require the agent to possess a belief that he is doing the right thing.

the action, and this ignorance is due to not knowing the circumstantial (non-normative) facts, then tracing is in order. For instance, if the slaveholder is ignorant that her slaves suffer, then we need to trace that ignorance and discover whether such circumstantial ignorance reflects a poor quality of will. Did the agent bury her head in the sand? Is her ignorance of her slaves suffering willful? Or was she, through no fault of her own, kept blissfully ignorant of the suffering of her slaves by the people overseeing her operation? If the circumstantial ignorance is willful or due to negligence, then it is attributable to the agent and provides the source of original responsibility capable of derivatively making the slaveholder responsible for her keeping of slaves.⁵⁷ In chapter 3, I discuss attributability in detail. For now, I set aside the view.

⁵⁷ Thanks to Hanser for urging me to relate this back to the conditions on blameworthiness discussed in section 1.3. The attributionist approach denies the Parity Condition. This means the Transfer Condition, Necessary Condition, and Explanatory Condition does not equally apply to both factual ignorance and moral ignorance. For instance, the Necessary Condition states that an agent is blameworthy for an unwitting wrongful act only an agent is blameworthy for a benighting act. If the agent commits that act in moral ignorance, then the Necessary condition will not hold. The agent can be directly blameworthy for the unwitting wrongful act even if it is false that the agent is blameworthy for a benighting act. The slaveholder falsely believing that slavery is permissible may not be blameworthy for holding that false belief because, say, she was unlucky to be born into a culture morally benighted concerning the wrongness of slavery. The same considerations hold for the Explanatory Condition when it comes to attributionism and situations involving moral ignorance. Yet, when it comes to unwitting wrongful acts committed due to circumstantial ignorance, the attributability theorist can endorse the Necessary Condition and the Explanatory Condition. The attributionist endorses the Transfer Condition in either situations. Yet when it comes to moral ignorance the Transfer Condition is vacuously true, as the agent is not blameworthy for a benighting act. But it is still the case that were the agent blameworthy for a benighting act (of which they are not in cases of pure moral ignorance), then the agent would be blameworthy for the unwitting wrongful act. How does this relate to the main orthodox conditions (I) and (II)? The attributionist will endorse condition (I). A person is blameworthy for her moral ignorance provided that it reflects a poor quality of will, and this is required for blameworthiness for the unwitting wrongful act. Blameworthiness for ignorance that the unwitting wrongful act is wrong is also required in cases of factual ignorance. Yet, the attributionist will not endorse condition (II) in all cases. If the case involves moral ignorance, then (II) will be false. On this view, there is no need for tracing blameworthiness for moral ignorance back to a benighting act for which the agent is blameworthy. Yet, if the case involves factual ignorance, then (II) will hold. Blameworthiness for the ignorance will involve tracing that ignorance back to a blameworthy benighting act.

Another approach to the skeptical challenge is “externalism.” This approach sticks to the standard view of accountability, but it relaxes what is required to secure accountability for an unwitting wrongful action. This view agrees with the attributionist approach in denying step (3) in the skeptical argument. Original responsibility for an unwitting wrongful action does not require that action to be a case of clear-eyed akrasia. Externalists think that acts stemming from vices can be acts for which agents are originally responsible. If so, this halts the regress of blameworthiness. There is no need for further tracing to establish blameworthiness for the act stemming from vice. It is a source of original responsibility. We need not see whether the agent knowingly committed the act of vice against her considered judgment about what she had most reason to do.

Externalism also agrees with attributionism in denying step (3) in the skeptical argument. As mentioned, on this view an act of vice can be the origin of original responsibility even if the agent is ignorant that the balance of reasons does not favor performing the act.

Externalism and attributionism differ when it comes to cases involving moral ignorance.⁵⁸ In such cases, attributionism will regard an agent blameworthy for an unwitting wrongful act even if the agent is not blameworthy for a benighting act that led to the unwitting wrongful act. This holds in cases that involve moral ignorance because the person was unlucky to be born into a morally benighted culture. Even if her moral ignorance does not stem from an act of vice the agent can still be

⁵⁸ Thanks to Hieronymi for urging me to clarify how externalism and attributionism (i.e., a quality of will view) differ.

blameworthy for her ignorance that the unwitting act is wrong. In such cases, externalism traces blameworthiness for ignorance to blameworthiness for a benighting act. But, when there is no blameworthiness for a benighting act (i.e., an act involving vice or clear-eyed akrasia) in the causal history there will not be blameworthiness for the unwitting wrongful act.⁵⁹ This difference between the views leads me to shift from externalism to attributionism in chapter 3, but, to start with, in chapter 2 I articulate externalism in depth and defend the view from objections.

1.7 Conclusion

This chapter detailed the skeptical challenge to moral responsibility from the possibility of wrongful acts done in ignorance. I detailed the key conditions, commitments, and arguments involved in the skeptical challenge. In addition, I discussed two response strategies to the skeptical challenge. In the next chapter, I will defend the externalist response strategy.

⁵⁹ In contrast to what I discussed in a previous footnote concerning attributionism, externalism will embrace orthodox conditions (I) and (II) discussed in section 1.3. Attributionism embraces condition (I) and denies condition (II) in cases involving moral ignorance, but externalism endorses condition (II) in all cases. This also results in differences in the subsidiary conditions, as externalism endorses the Parity Condition whereas attributionism rejects the Parity Condition.

CHAPTER 2: THE EXTERNALIST RESPONSE TO THE SKEPTICAL CHALLENGE

2.0 Introduction

In this chapter I detail externalism's response to the skeptical challenge to moral responsibility. After doing so, I discuss an objection to externalism raised by Neil Levy. Then I consider Philip Robichaud's reply to Levy and Levy's rejoinder to Robichaud. I argue that Levy's rejoinder fails by posing a dilemma for his view and arguing that the view of rationality underlying Levy's rejoinder is implausible. I conclude this chapter by revising and improving externalism's main principle.

2.1 Externalism's Response to the Skeptical Challenge

To better understand externalism about the epistemic condition on moral responsibility it is helpful to compare the view with internalism, a competing view. Externalism and internalism alike endorse the idea that blameworthiness for acts done in circumstantial or moral ignorance must trace back to blameworthiness for that ignorance. The views differ concerning what establishes blameworthiness (or culpability) for such ignorance. Internalism is restrictive. Culpability for ignorance must always trace back to an akratic act—an act whereby one believed overall that one ought to improve one's epistemic position, through reasoning or evidence gathering, but one failed to do so. Externalism is less restrictive. It allows for a disjunctive requirement on culpable ignorance: either the ignorance is the result of an akratic act

or the ignorance is the result of a non-akratic act stemming from the voluntary exercise of vices.⁶⁰ To see this difference, consider the following case:

Dr. Feelbad. A doctor named Dr. Feelbad prescribes her patient Mr. Jones a drug that interacts with another drug Mr. Jones is taking. This magnifies the adverse effects of the drugs and causes Mr. Jones considerable harm to his health. When prescribing the harmful drug, Dr. Feelbad was ignorant that it would adversely interact with a drug Dr. Feelbad knew Mr. Jones was taking. Dr. Feelbad failed to inquire out of laziness into the possibility that the new drug would adversely interact with the old drug. This omission due to laziness was not accompanied with a belief that she should, before prescribing the new drug, investigate possible negative drug-drug interactions.

Dr. Feelbad had a duty to inquire into possible negative interactions of the drug she prescribed Mr. Jones because of her duty not to expose the patients under her care to unnecessary risks of harm. This particular duty of inquiry is derived from the fact that satisfying that duty would enable Dr. Feelbad to see how she could satisfy her duty not to impose unreasonable risk of harm on Mr. Jones, namely by not prescribing him a drug that would interact with his current medications.

Internalism and externalism differ in their diagnosis of the Dr. Feelbad case. According to internalism, the quest continues for the source of original responsibility.

⁶⁰ I follow FitzPatrick (2008) in referring to the exercise of vice as “voluntary.” The qualifier “voluntary” indicates the agent was not, in a sense, coerced by her own desires to act as she acted. The exercise of the vices was not due to compulsions or irresistible desires, for instance. This qualification keeps the focus on the epistemic condition for moral responsibility. Otherwise, we may wonder if the agent acted with the sort of control required for responsibility.

Is there an akratic act that we can identify somewhere in the causal chain that led to the unwitting wrongful act of prescribing Mr. Jones the harmful medication? For externalism, the non-akratic act of failing to satisfy her epistemic obligation due to the exercise of vice renders Dr. Feelbad responsible for the benighting act and the unwitting wrongful act.⁶¹

This divergence between the views is not a benign difference. Unwitting wrongful acts occur often. If responsibility for such acts must trace back to an akratic act, and akratic acts are difficult to identify, this restricts the scope of our warranted judgments of responsibility. Internalism is revisionary of our ordinary practices of holding people responsible. Alternatively, externalism licenses a broader range of responsibility judgments. The view affirms our ordinary practices. Despite this divergence between the views regarding whether they are revisionary of our ordinary practices, the theories agree that the relevant notion of responsibility they are both trying to capture is responsibility as accountability.

Responsibility as accountability focuses on holding people responsible for flouting expectations or demands. However, the demands must be reasonable. The reasonability of the demands is determined by whether it is fair to hold people responsible for flouting them. Why this focus on fairness? Those who embrace responsibility as accountability often think of the reactive attitudes, such as blame, resentment, and indignation, as a form of adverse treatment. They think of the reactive

⁶¹ To further clarify, an agent “exercises a vice” when they freely act on the vice. The phrase is defined negatively. Exercising a vice involves not acting on a vice in a way that subverts satisfaction of the control condition for moral responsibility. For instance, the agent was under external compulsion to act on the vice, and the vice itself was not akin to a freedom-undermining irresistible compulsion.

attitudes as akin to social sanctions. And it is only fair to think people are liable to receive social sanctions when they could have voluntarily complied with those demands or flouting the demands in question was in some sense avoidable for them. This is why responsibility as accountability is regarded as volitional-focused. It is only fair to blame an agent who has rational capacities and the ability to voluntarily exercise those capacities. If complying with a set of demands was not within someone's control or she couldn't help flouting some of the demands communicated to her, it is not fair to then hold her responsible (or blame her) for her failure to comply. (We suppose here that the demands communicated to the agent in question are unreasonable.)

Given that internalism and externalism agree that the nature of responsibility is accountability, the relevant question to consider when it comes to an agent's accountability for her ignorance is:

Reasonability Question: 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?⁶²

How does externalism address the Reasonability Question? To answer this question, consider a case frequently discussed in the literature. The following scenario is based on the movie *It's a Wonderful Life*:

[C]onsider Mr. Potter, a powerful businessman who holds false moral views. He takes certain business practices—such as liquidating Bailey's Building and

⁶² William FitzPatrick (2008: 603) asks this question, and Neil Levy (2009: 732) assents to its importance. FitzPatrick cites Benson (2001: 613-4) as inspiration regarding this question.

Loan and sticking it to the poor families of Bedford Falls—to be “permissibly aggressive,” when in fact they’re “reprehensibly ruthless.” This leads him to do bad things, though he doesn’t understand that he’s acting badly, which means that he’s acting out of a certain kind of ignorance. He’s fully aware of all the circumstances, but he applies flawed normative principles or weightings and comes up with bad decisions. Is he culpable for his bad actions? (FitzPatrick 2008: 600)⁶³

Mr. Potter acts out of moral ignorance. He holds false moral beliefs, and he “applies flawed normative principles or weightings” in his deliberations about what he should do.⁶⁴ Yet Potter is not ignorant of circumstantial facts. He is aware that his business decisions cause the people who are vulnerable in his community to suffer. He is also aware that there are other ways to make money than by “sticking it to the poor families” in his community. Potter does not believe that he is doing anything wrong. He believes his actions are “permissibly aggressive” and shrewd, not “reprehensibly ruthless”. Yet, Potter has an obligation not to impose unreasonable risk of harm to the welfare of the people in his community. From this obligation is derived his procedural epistemic obligation to self-scrutinize his moral beliefs. Had Potter complied with his procedural epistemic obligation he could have seen how to comply with his obligation not to impose risk of harm on others. He could have used accurate normative principles based on true moral beliefs and arrived at decisions that did not risk

⁶³ FitzPatrick (2008) introduces this case into the literature, but the case is analogous to a case found in Rosen (2004).

⁶⁴ The movie portrays Mr. Potter as a Social Darwinist, but he might equally be conceptualized as a narcissist or egoist in the sense that his normative principles disproportionately favor his own interests, while downplaying the exposure to risk of harm to others.

harming others. How does this connect to the Reasonability Question? Why is it fair to expect or demand of Potter that he should have corrected his false moral views?

Externalist William FitzPatrick provides more details in the Potter case in answering the Reasonability Question:

Let us imagine that, while Potter may have been raised with an impoverished moral outlook, he has received a good general education and has since been exposed to lots of other views and challenges to the outlook he has grown comfortable with and that he is capable of understanding them as such. He has, after all, lived right in the same town with Peter and George Bailey for years, which means that the opportunity for improved normative understanding was clearly present in his social context (even if viewed by him mostly as a source of annoyance). The question, then, is why such opportunities were not taken. Suppose that the answer is that they were not taken because of a cavalier dismissal of the relevant considerations, viewpoints, and arguments—an easy labeling and dismissing of them as “socialist” or “liberal,” without honest reflection open to the possibility that they may point to blind spots in his own views. (FitzPatrick 2008: 603-4)

Given the above considerations, FitzPatrick’s answer to the Reasonability Question is based on two claims concerning Potter.

Context/Capabilities Claim: There were 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.

Epistemic Vice Claim: The failure of adequate reflection and information gathering was instead the result of voluntary exercises of vices such as overconfidence, arrogance, dismissiveness, laziness, dogmaticism, incuriosity, self-indulgence, contempt, and so on. (2008: 605)⁶⁵

From these claims FitzPatrick responds to the Reasonability Question: “[Potter] could thus reasonably have been expected to take steps that would have eliminated that ignorance, by refraining from exercising those vices and instead taking advantage of the epistemically relevant opportunities available to him” (605). Given that the externalist position relies on the notion of a vice, I need to clarify the notion. What is a vice? What is it to exercise a vice? And, what is it to voluntarily do so?

A vice stands in contradistinction to a virtue. What is a virtue? A virtue is a quality (i.e., a capacity or character trait) that makes a person an excellent person.⁶⁶ Some virtues are hard-wired cognitive capacities like having a good memory or good perception. Virtues can include skills like the ability to perform logical inferences. Virtues can also include moral traits or intellectual traits like being benevolent or

⁶⁵ One might wonder: why does FitzPatrick include the Context/Capabilities Claim as part of his answer to the Reasonability Question? Isn't it enough that the Epistemic Vice Claim is true? FitzPatrick, as an accountability theorist that countenances volition (or the possibility of voluntary exercise of capabilities) must include the Context/Capabilities Claim to indicate that nothing in the environment was impeding adequate reflection and the gathering of evidence. If social context or capabilities made such actions sufficiently difficult or impossible, then it would be unreasonable to expect him to rectify his ignorance, and that he failed to do so and exercised vice would not be sufficient to make him culpable for his ignorance, as, on FitzPatrick's account, he did not have reasonable opportunities to take advantage of. The exercise of vice should not be held against him.

⁶⁶ For a helpful survey of the nature of virtue and vice see Battaly (2015).

open-minded.⁶⁷ A vice is a capacity or character trait that makes a person a worse person. Vices include having a bad memory, lacking the ability to make logical inferences, being dogmatic or closed-minded, or being unjust. Vices are qualities that make a person lacking in some way. For instance, if a person has the intellectual vice of being closed-minded, they will fail to secure good ends like true beliefs or not be properly motivated to seek truth.⁶⁸ Not being open to the opinions of others shelters one's opinions from critical scrutiny. It forfeits opportunities for moral and epistemic change and improvement.

What is it to exercise a vice? Vices are either hard-wired capacities, acquired skills, or character traits (i.e., dispositions). To exercise a vice is for that vice to causally impact a decision, action, or the formation of a belief. This can occur directly, as when a person exercises the vice of viciousness by hitting an unsuspecting innocent person with their fist. Or, it can occur indirectly as when a person is dismissive of the opinions of others, as in the Potter case, and this results in them not forming a belief they otherwise would have formed had they charitably entertained others' opinions. Voluntarily exercising a vice was previously discussed in a footnote. This is not

⁶⁷ There are moral or epistemic virtues. Likewise, there are moral or epistemic vices. In this dissertation, I do not make hard distinctions between moral and intellectual vices. Most agents in the cases at issue, such as the Potter case, possess a mixture of moral and intellectual vices.

⁶⁸ This is two ways of understanding vices that connect to two ways of understanding virtues. I am intentionally pluralistic between these two readings. Thinking of virtues as making people better because they tend to secure good ends, such as well-being and true beliefs can lead to a reliabilist conceptions of the virtues, which taken alone neglects motivational components of virtue. Likewise, thinking of virtue as strictly a matter of proper motivation can lead to a responsibilitist conception of the virtues as being valuable because one is properly motivated to secure well-being and true beliefs. These ways of conceiving of the virtues can come apart. If motivation is what makes virtues valuable, then one can have proper motivation yet consistently miss securing good ends and still possess the virtue. Alternatively, one can reliably secure good ends yet not be properly motivated and still count as virtuous. I will not take a stand on the true nature of virtue, as that extends beyond the scope of this project. I will allow both readings as admissible. The corresponding vices fall out of those two readings of virtues.

understood in a robust sense in terms of possessing regulative control over whether one acts on the vice.⁶⁹ Such an understanding of voluntariness would imply that we cannot voluntarily act on vices that are hard-wired capacities, such as forgetting something due to a poor memory. Instead, a person voluntarily exercises a vice when there are no freedom-disqualifying conditions present, such as acting on irresistible compulsions or being forced to act in a vicious way due to an external force. With this general understanding of the nature of virtue and vice in place let's return to the regress.

How does externalism halt the regress of blameworthiness? How does the voluntary exercise of vices generate original responsibility for blameworthiness? Why not think introducing vice into the picture generates a question about whether Potter is ignorant that possessing his vices is bad? If so, we can keep the regress going. We can ask whether ignorance of the badness of his vices is culpable. On behalf of the internalist, FitzPatrick considers a rejoinder to his vice-theoretic externalism along these lines:⁷⁰

Internalist Rejoinder: Unless Potter was fully akratic in making the vice-related, epistemically debilitating choices that increased and reinforced his moral ignorance, he must have thought he was behaving well—which is to say, he must already have been ignorant about the status of his character traits and choices. And this brings us back to steps [(3)-(5)] of the [skeptical] argument: he is not originally responsible for those bad choices and is derivatively

⁶⁹ Regulative control (or managerial control) is based on the garden of forking paths metaphor of future decisions or actions. On this model one has the right sort of control when one could, through an act of the will alone, actualize different paths into the future.

⁷⁰ FitzPatrick imagines this internalist objection as possibly being put forward by Rosen.

responsible for them only if he is originally responsible for what led to his ignorance about his character traits and choices. (2008: 605-6)

FitzPatrick responds to the Internalist Rejoinder by arguing that Potter is not completely oblivious that his vices are problematic. Though Potter does not think of his vices as vices Potter does recognize at a general level that his behavior is epistemically problematic.⁷¹ FitzPatrick contrasts the internalist's strong knowledge requirement involving clear-eyed akrasia with a weaker knowledge condition:

Weaker Knowledge Condition: It is reasonable to expect one to rectify one's ignorance...if one knew the general epistemic importance of subjecting one's beliefs to critical scrutiny, so it could be reasonably expected to occur to one that one's behavior with regard to one's moral beliefs was epistemically irresponsible. (FitzPatrick 2008: 607; additions and modifications mine)

Potter satisfies this Weaker Knowledge Condition, as FitzPatrick argues,

It would be disingenuous, for example, to claim on his behalf that, given his views, he couldn't reasonably be expected to appreciate the importance of broader, critical engagement with people outside of his narrow, elite social sphere. As an educated businessman, he would surely have understood the importance of genuinely critical reflection in other spheres, such as analysis of stock market and interest rate trends, and the epistemic dangers of relying for such information on only one firm or group of analysts with vested interests

⁷¹ The general recognition is epistemic, not moral, as Potter is not generally aware his behavior is wrong, but he recognizes that lack of critical thinking is problematic when it comes to other business practices, as FitzPatrick explains in the quote explaining the Weaker Knowledge Requirement.

without taking opposing argument seriously. Presumably, then, he would have insisted on high standards of critical reflection in the sphere of his work, and yet he did not make the obvious extension to the moral sphere—the recognition that one is unlikely to arrive at an accurate moral view by uncritically accepting the received opinion of a narrow group of people with vested interests in a certain set of answers and insulating oneself from open, critical discussion incorporating different points of view and forms of experience. (2008: 607)

Satisfaction of the Weaker Knowledge Condition does not entail satisfaction of the stronger knowledge requirement involving akrasia. Although Potter knew the general importance of subjecting his beliefs to critical evaluation, he did not believe he was doing anything wrong. Yet one may still wonder if Potter satisfying this Weaker Knowledge Condition is sufficient for defusing the Internalist Rejoinder?

An internalist might reply that although Potter did not conceive of his vices as vices and he knew that, in general, shielding his opinions from critical scrutiny does not tend to produce accurate opinions, Potter still made his “epistemically debilitating choices” in ignorance that his vices were bad. To establish culpability for his ignorance Potter needed to know that acting on his vices is something he ought not do all-things-considered, yet he decided to act on his vices anyway. The externalist has not eliminated the need for tracing culpability to an akratic episode.

In response to the internalist we can couple the Weaker Knowledge Condition with a claim about how vices generate original culpability without the need for conceiving of the vices as vices (i.e., as bad character traits worthy of being rectified). Notice that in the Epistemic Vice Claim Potter’s vices are identified as,

“overconfidence, arrogance, dismissiveness, laziness, dogmaticism, incuriosity, self-indulgence, contempt, and so on” (FitzPatrick 2008: 605). Vices are often related to their corresponding virtues. The most prominent intellectual virtues that correspond to the vices Potter exhibits in his deliberations is intellectual humility and open-mindedness. Potter’s vices express a lack of intellectual humility and a lack of open-mindedness. To focus on open-mindedness, vices that demonstrate lack of open-mindedness can generate original culpability without needing to conceive of open-mindedness as good and the absence of open-mindedness, such as the vices of dogmaticism and dismissiveness, as bad. As externalist James Montmarquet argues, it is reasonable to expect and demand of people that they are open-minded because

a certain quality of *openness* to truth- and value-related considerations is expected of persons and that this expectation is *fundamental*, at least in the following regard. The expectation is not derivative of or dependent upon one’s (at the moment in question) judging such openness as appropriate (good, required, etc.)—just the opposite: it would include a requirement that one be open to the need to be open, and if one is not open to this, one may be blameworthy precisely for that failure. (Montmarquet 1999: 845)

On this picture, Potter is directly blameworthy for his lack of openness to the truth concerning his moral views. Given his context and capabilities Potter could have engaged in critical reflection and evidence gathering concerning his moral opinions. However, he failed to satisfy his procedural epistemic obligations in this regard due to the voluntary exercise of vices. Though he did not conceive of those vices as vices and think himself doing something wrong by failing to scrutinize his moral beliefs, Potter

did satisfy a weaker knowledge condition in that he recognized the general importance of critically scrutinizing one's beliefs. Though he never applied this general knowledge to the domain of his moral beliefs, it is reasonable for us to expect and demand that he should have recognized the importance of subjecting his moral beliefs to the critical pressure of opposing arguments and viewpoints. Additionally, we can imagine that the vices that prevented him from subjecting his moral beliefs to such scrutiny evidence the absence of the virtue of open-mindedness. Our reasonable expectations concerning open-mindedness are non-derivative. We need not demand that Potter recognize open-mindedness as good or a worthwhile trait in order to ground culpability for his failure to be open-minded; instead, it is reasonable to hold and communicate expectations that, "include a requirement that one be open to the need to be open, and if one is not open to this, one may be blameworthy precisely for that failure" (Montmarquet 1999: 845). Potter is directly culpable for his failure to be open-minded because his vices violated the non-derivative expectation that he be open-minded, which is reasonable despite its fundamentality.⁷² These strands are brought together into externalism's main strategy for halting the regress of culpability in the case involving Potter:

It is enough for culpability if Potter made his epistemically debilitating choices—cavalierly dismissing opposing arguments, insulating himself from open, critical discussion or relevant sources of information, and so on—out of indulgence of vices, in a context where he could reasonably have been expected

⁷² Given this, the Origination Thesis is false. Similarly, a restricted version of the Origination Thesis, which I dub the Awareness Thesis, is false. The Awareness Thesis holds, "One is directly culpable for being in a given mental state only if one is aware of the wrongness of being in that state" (Montmarquet 1999: 844). Rosen (2004), Zimmerman (2008), and Levy (2011) embrace the Awareness Thesis.

to know better and to do a better job of informing himself morally, given his capabilities and culturally available opportunities. (FitzPatrick 2008: 606)

Generalizing the claims above results in externalism's Culpable Ignorance Principle:

Culpable Ignorance Principle: Ignorance, whether circumstantial or normative, is culpable if the agent could reasonably have been expected to take measures that would have corrected or avoided it, given his or her capabilities and the opportunities provided by the social context, but failed to do so either due to akrasia or due to the culpable, nonakratic exercise of such vices as overconfidence, arrogance, dismissiveness, laziness, dogmatism, incuriosity, self-indulgence, contempt, and so on. (FitzPatrick 2008: 609)

The Culpable Ignorance Principle indicates that failure to take advantage of opportunities when one has the capacity to do so is culpable given either knowing wrongdoing (akrasia) or unknowing wrongdoing because of the exercise of vices. Potter commits unknowing wrongdoing due to acting on vices. Such a principle responds to the Skeptical Challenge as put in argument form by indicating that step (3) in the argument is false. Recall that step (3) claims:

(3) Act A is an origin of original responsibility (such as case (2c) above) only if

(e) S knows the balance of reasons against doing A; ignorance of this—whether due to circumstantial ignorance or to normative ignorance—removes original responsibility for A.

Potter's exercise of vices is a source of original responsibility despite it being false that Potter knows that, due to the balance of reasons, he should not exercise vices of

overconfidence, laziness, dogmatism, and so on.⁷³ This line of thinking also provides a response to the regress of blameworthiness that plagued internalism. Consider the regress applied to the case involving Potter:

(1) Potter is blameworthy for reprehensible ruthless act A2 only if

(i) Potter believes A2 is wrong, or (ii) Potter is blameworthy for his ignorance that A2 is wrong;

Condition (ii) holds only if

(2) Potter is blameworthy for the past omission A1 that resulted in his ignorance that A2 is wrong;

Condition (2) holds only if

(iii) Potter believes A1 is wrong, or (iv) Potter is blameworthy for his ignorance that A1 is wrong;

Condition (iv) holds only if...

The externalist response strategy denies that (iii) and (iv), as stated, are the only two ways that condition (2) can be satisfied. Instead of making (iii) and (iv) the only conditions individually necessary to establish condition (2), the externalist posits a *sufficient* condition for Potter's blameworthiness for his past omission to engage in critical reflection on his moral beliefs. His blameworthiness for his omission need not trace to his belief that the omission is wrong or blameworthiness for his ignorance that

⁷³ To clarify, the laziness in question is epistemic laziness with respect to the opinion of others concerning the morality of his business practices. The character Potter in the movie is industrious. But Potter doesn't know that the balance of moral reasons favors him investigating the moral opinions of others, or that he ought not act on vices that hinder discovery of alternative moral opinions.

the omission is wrong. Instead, using the Culpable Ignorance Principle, we get a disjunctive sufficiency condition for (2).

Condition (2) holds *if*

(iii*) Past omission A1 is due to akrasia, or (iv*) Past omission A1 could reasonably have been expected to have been avoided, given Potter's capabilities and social context, but Potter failed to avoid A1 due to the non-akratic exercise of vices.

Potter's past omission A1 being due to the non-akratic exercise of vices suffices for Potter being blameworthy for the omission that resulted in his ignorance that A2 is wrong. Potter's past omission is a source of direct (or original) blameworthiness. The regress of blameworthiness is halted.

2.2 Levy's Objection to Externalism

Neil Levy agrees with FitzPatrick that considerations of fairness govern attributions of blameworthiness for ignorance. Liability to blame is liability to receive a social sanction, a burden, or a setback to one's interests. As a result, it needs to be fair that the person in question is blamed for his or her ignorance. Levy embraces the Reasonability Question as the right question to ask.⁷⁴ Levy agrees with FitzPatrick that it is fair to consider a person blameworthy for her ignorance only if the person could have reasonably been expected to correct the ignorance. Yet Levy parts way with FitzPatrick regarding when it is reasonable to expect the rectifying of ignorance.

⁷⁴ As Levy says about the Reasonability (R) Question, "I concur; R indeed asks the right question. I shall argue, however, that if we answer it fairly the right response is nothing: there is nothing that Potter could reasonably have been expected to do to avoid or remedy his normative ignorance" (Levy 2009: 732). Levy embraces (R), but he answers it in the negative as opposed to the affirmative.

Levy's answer to the Reasonability Question marks a significant departure from FitzPatrick's answer. Levy argues that it is only reasonable to expect an agent to rectify her ignorance when that agent possesses the capacity to rectify it. As Levy (2009: 735) argues:

Agents are culpable for actions (inter alia) if it is reasonable to expect them to conform their behavior to the appropriate normative standards, and this expectation is in place only if they have the capacity to behave accordingly. It is a necessary condition of agents' possessing this capacity—and therefore of us reasonably expecting them to behave appropriately—that conforming their behavior to normative standards is something they can do rationally (and not merely by chance or accident). Potter, I argue, could not rationally have taken advantage of the opportunities for moral improvement that were (in some sense) available to him; hence, we cannot reasonably expect him to do so.

When does an agent have such a capacity to rectify her ignorance? An agent has the capacity to rectify her ignorance when there is an explicit reasoning procedure available to her that she could use to rectify her ignorance. Levy thinks, further, that such a reasoning procedure is available to an agent, in the relevant sense, only if she possesses internal reasons to engage in that reasoning. When an agent possesses internal reasons to engage in reasoning that would rectify her ignorance, it is not only rationally permissible for her to engage in such reasoning, it is reasonable for us to expect her to do so and therein comply with her procedural epistemic obligations.

Philip Robichaud (2014: 141) packages Levy's view into two main claims:

Reasonable Expectations Claim: It is reasonable to expect an agent to form her beliefs in accordance with some procedural epistemic obligation only if she has the capacity to do so.

Capacity Claim: An agent has the capacity to form her beliefs in the expected way only if she can conform rationally to that expectation.

Let's return to the Potter example to see how these claims work. We might hold Potter to a procedural epistemic obligation by insisting that he take steps to better inform himself about alternative views of his social obligations: accounts on which they are more expansive than Social Darwinism allows. If he does not educate himself about alternative normative views, we can infer that Potter is not trying to get at the truth regarding what he ought (and ought not) do. In the absence of a positive first-order argument against Social Darwinism we assume here that discharging his epistemic obligations would nevertheless expose Potter to an effective argument against Social Darwinism which would in turn force him to conclude that his business practices are morally wrong, as they have led him to embrace lending practices causing significant and avoidable harm to the welfare of people in his community. By assumption Potter would no longer hold false moral beliefs (in Social Darwinism) were he to "vet" his beliefs by satisfying his procedural epistemic obligations. But given the Reasonable Expectations Claim, it is reasonable to *expect* Potter to rectify his moral ignorance in situ only if he has the capacity to do so. Potter has the capacity to rectify his moral ignorance only if it is rational for him to do so. Can Potter rationally comply with his procedural epistemic obligation? Levy does not think so. But, this is because of how Levy is conceiving of rationality.

Levy endorses an internalist conception of rationality and practical agency, “what an agent can do rationally...is a function of what she takes her reasons to be” (Levy 2011: 127).⁷⁵ Being rational is a matter of being guided by internal reasons. Levy endorses a specific view of what an internal reason is:

Internal Reason: An agent has a reason to ϕ if she actually believes that ϕ -ing will (tend to) bring about a state of affairs that furthers the satisfaction of her goals. (Levy 2009: 735)

Two further normative commitments underlie Levy’s conception of rationality:

Mentalism Norm: What it is rational for me to do—that is, what I can decide to do as the result of engaging in reasoning—is settled by my actual mental state (Levy 2009: 735); and,

Awareness Norm: We can only reasonably be expected to do what we can do by an explicit reasoning procedure, a procedure we choose to engage in, and when we engage in explicit reasoning we cannot deliberately guide our behavior by reasons of which we are unaware, precisely because we are unaware of them (Levy 2009: 736 n.16).

Given the above norms and commitments, Levy endorses a strongly internalist view of rationality and practical action. Rationally doing things to rectify one’s ignorance only occurs when actions aimed at rectifying one’s ignorance are guided by internal, instrumental reasons that represent such actions as likely to promote the satisfaction

⁷⁵ Levy’s book *Hard Luck* includes a slightly revised version of his original response essay to FitzPatrick (i.e., Levy’s 2009 article in *Ethics*). I quote from both Levy sources in getting his view on the table. I also refer to Levy’s further defense of his internalist view as found in Levy (2016).

of one's goals.⁷⁶ Further, an agent needs to be aware of such internalist reasons. An awareness of such reasons makes it possible for an agent to intentionally rectify her ignorance, as she can choose to engage in an explicit reasoning procedure to rectify her ignorance from her awareness of reasons to do so. When an agent satisfies such conditions, the agent has the capacity to rationally rectify her ignorance, and it is reasonable and fair for us to expect her to do so (or to expect her to have done so).

An agent does not have the capacity to rectify her ignorance intentionally when it is merely "possible that the agent remedies her ignorance by mere chance or through some glitch in her epistemic agency" (Levy 2016: 2). An agent engages in a reasoning procedure by chance when she does not believe that she ought to improve her epistemic position through reasoning but she engages in such reasoning anyway. From the agent's perspective, engaging in the reasoning is irrational, as she possesses no (internal) reason to engage in it. The reasonability of our expectation that the agent rectifies her ignorance tracks the agent's possession of such a rational capacity for reasoning. Levy connects this to the case involving Potter because:

It is often, perhaps usually, the case that we cannot reasonably demand of agents that they do not act in ways that express their epistemic vices. So acting is...not something that agents can do rationally and it is unreasonable for us to ask people to behave in ways that are not rational for them. (Levy 2011: 126-7)

It is not fair to expect Potter to correct his ignorance. Given his vices, Potter does not have internalist reasons to improve his epistemic position through scrutinizing his

⁷⁶ Where such reasons consist in belief-desire pairs.

moral beliefs. As such, it is not rational for Potter to scrutinize his moral beliefs through an explicit reasoning procedure. And, it is not reasonable for us to require Potter to do something that it is not rational for him to do. For Levy, Potter possesses no internal reasons to revise his opinions:

Because Potter's outlook is constituted by his moral views, the opportunities for moral improvement he encounters—say, editorials by Paul Krugman in the *New York Times*—are seen by him as the whining of socialists or sissies. It is true that these judgments express Potter's epistemic vices, and that his complacent dismissal of these views is inconsistent with the relatively high standards he sets himself in other areas of inquiry. But *by his lights*, Potter governs his normative views adequately. He gives competing views the attention he takes them to deserve....But if Potter does not see that he is managing his moral views badly, he has no (internal) reason to manage them any differently. Potter exhibits epistemic vices aplenty, but because he does not conceive of them *as vices*, he has no internal reason to refrain from so doing. (Levy 2011: 127)

Given that Potter does not recognize his vices as vices, he does not possess a mental state with the content that he has poorly managed his moral opinions. Levy notes:

FitzPatrick concedes as much himself, saying that *because* Potter fails to apply the knowledge concerning effective epistemic management he has learned in the course of his business activities to the regulation of his moral views, he 'does not see that he is acting badly with respect to the management of his moral opinion' [FitzPatrick (2008: 127)]. Indeed, FitzPatrick must concede the point,

since he is concerned exclusively with cases that do not involve benighting acts; having Potter manage his opinion in ways that are not adequate by his own lights would immediately reintroduce such an act. (Levy 2011: 127)

Given Levy's claims and norms, Potter does not have the capacity to comply with his procedural epistemic duty to take steps to rectify his ignorance.⁷⁷ He cannot rationally do so. It is thus unreasonable for us to expect Potter to act to rectify his moral ignorance. To connect this back to the regress and skeptical argument, Potter is not directly blameworthy for the past omission that resulted in his ignorance that his unwitting wrongful act is wrong. The search continues for the source of culpability for Potter's ignorance.

As mentioned, Levy is an internalist. He thinks his claims and norms entail the idea that *culpable* ignorance always traces back to a prior akratic act. As he summarizes:

Agents can only correct their ignorance through a reasoning procedure...only when they can rationally embark on that procedure. But they can rationally embark on that procedure only when they recognize that they ought to take an available opportunity to improve their epistemic situation. When an agent recognizes that they ought to take an opportunity, but fails to do so (and their

⁷⁷ To be charitable, Levy may have opted for his strongly internalist capacity claim in response to FitzPatrick's externalism to avoid a response from FitzPatrick. If Levy argued that Potter has the capacity to consider Krugman's views seriously, but because he has no internal reason to exercise the capacity it would be unreasonable for us (or those considering blaming him) to expect him to exercise the capacity, then FitzPatrick could respond that this culpability still terminates in Potter's exercise of vices. Potter lacks internal reasons that rationalize his exercising his general capacity because he exercises moral and epistemic vices. It is reasonable for us to expect Potter not to exercise such vices, so it is reasonable for us to expect him to possess internal reasons that rationalize using his general capacity for critical reflection. So, culpability for ignorance still bottoms out in the exercise of vices. Levy embracing strong Internalism avoids this response, as Potter lacks the requisite capacity.

failure is not itself excused by compulsion, coercion or what have you), they act akratically. So culpable ignorance is traceable back to akratic management of one's epistemic agency. (Levy 2016: 2)

Before making plausible the externalist position by filling in details of the view I will argue against Levy's objection to externalism by defending a reply to Levy.

2.3 Robichaud's Two-pronged Reply to Levy's Objection

Philip Robichaud (2014) offers a two-pronged reply to Levy's objection to externalism. The first prong argues culpable ignorance is not always traceable to an akratic act, even given Levy's norms and commitments from internalist rationality. An agent can have the capacity to rectify their ignorance through an explicit reasoning procedure, fail to exercise that capacity, not act akratically, and yet be culpably ignorant. According to Robichaud's second prong there is ambiguity in Levy's internalist rendering of what it means to have the capacity to rationally correct one's ignorance. Levy's capacity notion is disambiguated in four ways. The first three ways will not establish Levy's claims, and the fourth way is independently implausible.

Robichaud's first prong targets Levy's claim that an agent must believe that she should improve her epistemic position for her to possess the capacity to rationally improve her epistemic position. Robichaud thinks Levy has neglected an important class of internalist reasons, namely motivating reasons that are sufficient yet nondecisive. Motivating reasons are reasons that explain why an agent did something.⁷⁸ This contrasts with normative reasons. Normative reasons are reasons

⁷⁸ For simplicity I am combining motivating and explanatory reasons. Motivating reasons are not always explanatory reasons but this distinction need not concern us for this discussion.

that favor or justify performing an action. Normative reasons are often tied to the existence of facts and depend less on what an agent thinks about those facts.⁷⁹ Motivating reasons are tied more closely to the agent's mental states. Imagine that unbeknownst to me my friend is a member of the Irish Mob. I might (falsely) believe I have a reason to enter into business dealings with him. But, in fact, I may not have a reason to enter into business dealings with him because he is corrupt and will take me for all that I am worth. In such a case, I have a motivating reason to go into business with him, but I do not have a normative reason.⁸⁰ Motivating reasons are reasons why I think something good or worthwhile to pursue. They are decisive when I take those reasons to have enough strength to decisively support an action. Using the example of quitting smoking cigarettes Robichaud combines these considerations into his claims regarding motivating reasons that are sufficient yet nondecisive:

[A]n agent has sufficient, nondecisive motivating reasons to quit smoking when she takes herself to have reasons that are strong enough to make quitting a rational option, but not strong enough as to decisively support quitting. Such an agent might simultaneously take herself to have reasons to continue smoking that she (wrongly, let's say) takes to be just as strong as her reasons to quit. These other sufficient though nondecisive reasons would make it rational for the agent to continue with her smoking habit....it follows both that an

⁷⁹ Of course, whether normative reasons are facts or non-factive mental states will depend on your view of normative reasons. I follow Robichaud here and think of normative reasons in terms of facts and motivating reasons in terms of non-factive mental states (e.g., desires, beliefs, etc.). For a helpful overview of the nature of reasons see Wiland (2012).

⁸⁰ This is a simplification of the literature on reasons. There is a position called "perspectivism" according to which normative reasons track the perspective of the agent. Given this view, I would have a normative reason to enter into business dealings with my friend. After all, he's my friend and seems trustworthy.

agent's internalist reasons can simultaneously rationalize two distinct actions and that she can rationally perform either action. (Robichaud 2014: 142-3)

For Robichaud, a similar situation to the smoking situation can occur regarding having motivating reasons to improve one's epistemic position. An agent can have sufficient reasons to improve her epistemic position, yet those reasons can be nondecisive in that there are competing reasons of equal force indicating that she should not improve her epistemic position. When this happens an agent possesses the capacity to rationally improve her epistemic position. There are reasons that can guide the agent to a reasoning procedure that would lead to the improvement of her epistemic position. Yet, there are alternative reasons sufficient to justify her failing to improve her epistemic position. When the agent decides not to improve her epistemic position, she is not acting contrary to an all-things-considered judgment about what she has most reason to do, as the reasons are nondecisive in this regard. Yet, even if we accept Levy's internalist account of the relevant norms, it is still reasonable to expect the agent to rectify her ignorance because she had internalist reasons that enabled her to possess the capacity to rationally rectify her ignorance. Were she to rectify her ignorance through an explicit reasoning procedure, this would not necessarily involve the provenance of chance or luck.

We have seen how sufficient yet nondecisive reasons to improve one's epistemic outlook work in general. Now we need to see how this applies to Potter-like cases. Robichaud introduces the following variation on the case involving Potter:

CEO Topper: Topper is just like Potter in most respects. She owns a company. She has epistemic vices. She chooses not to scrutinize her moral beliefs, and

this results in her ignorance about the moral permissibility of her ruthless business practices. However, at several points in her moral development, Topper believed that she had sufficient, nondecisive reasons to engage in the kind of moral self-scrutiny we expect of her. Imagine that she took her reasons to scrutinize her beliefs about her business practices and her reasons to do whatever else to be of comparable strength....Topper is not in a position to claim that our expectation that she should have scrutinized her moral views is unfair on the grounds that such scrutinization was not a rational option for her. Because Topper took herself to have sufficient reasons to scrutinize her beliefs, she has internal reasons that ground her capacity rationally to scrutinize. Of course, like Potter, Topper did not believe she was under any kind of obligation to reflect on her business practices, but...this fact actually does not entail that she lacks the capacity rationally to conform to such an obligation. This capacity can be grounded in her belief that scrutinization is an action that has sufficient support given what she already believes...since Topper never believed that she was obligated to scrutinize her beliefs, she did not act akratically when she chose to do something other than self-scrutinize. (Robichaud 2014: 144)

The CEO Topper case is like the case involving Potter in key respects. Both Topper and Potter had an epistemic obligation to scrutinize their beliefs concerning the morality of their business practices. Both agents did not believe they possessed such an epistemic obligation, and both agents failed to conform to their respective epistemic obligation.⁸¹ As a result, both agents committed morally reprehensible acts in

⁸¹ It is an assumption of the cases that they have such an epistemic obligation.

ignorance that they were doing anything wrong. Yet, Topper and Potter are importantly different. Topper satisfied Levy's constraints for possessing the capacity to rationally comply with her epistemic obligations. She falsely believed she was not obligated to scrutinize her beliefs, but she possessed sufficient reasons that, had she taken them as decisive, would have guided her to employ a reasoning procedure that would have led to the improvement of her beliefs concerning the ethics of her business practices. For Levy, given that Topper did not believe she ought to (all-things-considered) scrutinize her beliefs, Topper lacked even the *capacity* to rationally comply with her epistemic obligations. Robichaud thinks such a result is absurd. Clearly Topper, as opposed to Potter, possessed internal reasons sufficient to support scrutiny of her moral beliefs. Any attempt by Topper to escape accountability for failing to meet her epistemic obligations that involves a plea from lack of rational capacities is surely misguided. The response to such a plea would be that the agent had the capacity to comply, as she had reasons that would enable her to do so rationally. She simply failed to act on those reasons. She chose not to improve her epistemic position. However, because she did not possess an all-things-considered judgment that she should improve her epistemic position, she never acted akratically by failing to comply with her epistemic obligation. Her false belief that she was not obligated to vet her beliefs does not get her off the hook. It does not justify the plea from lack of capacity.

If successful, the first prong of Robichaud's argument against Levy expands the scope of culpable ignorance. Agents can be culpable for their ignorance in the absence of a belief that they are obligated to comply with epistemic obligations. Agents can be culpable for their ignorance when they possess sufficient, though nondecisive, reasons to comply with their epistemic obligations. The agents may falsely believe they are not

obligated to comply with their epistemic obligations, but if they possess reasons that would enable them to rationally comply with such obligations, they possess the capacity to comply with such obligations. In many such cases, it will be reasonable for the affected parties to expect them to so comply, and when the agent fails to so comply, it is fair for harmed parties and their surrogates to hold the agent accountable for that failure, even though the failure does not involve akrasia or the agent knowingly going against her better judgment.

The second prong of Robichaud's attack on Levy's objection to externalism no longer uncritically embraces Levy's notion of being able to exercise one's capacity for critical reflection via possession of internal reasons. Robichaud assumes for the sake of argument that the first prong of his attack fails and aims to demonstrate that, "Levy's argument is problematic even on the assumption that possessing the rational capacity to perform an expected action requires believing that one ought to perform that action" (Robichaud 2014: 146 n.14). Under what conditions do one's internal reasons rationalize exercising one's capacity for complying with one's procedural epistemic obligations?

Robichaud disambiguates the situations under which internalist rationality licenses exercising one's capacities. Two distinctions are relevant in this regard. One can rationally vet one's beliefs via reasoning that features one's actual representations or mental states in a couple of ways. The first distinction is between conscious or unconscious mental states. Beliefs can be conscious or unconscious. I might consciously believe that I should investigate research that is relevant to my beliefs concerning whether it is ethical to eat meat; or, I might unconsciously, and perhaps

dispositionally, believe that I should investigate such research. The second distinction relates to how beliefs rationalize action. Beliefs can rationalize action directly or indirectly. A belief directly rationalizes an action when the agent holds the belief and the belief indicates that an agent ought to perform an action to improve her epistemic position. A belief indirectly rationalizes an action when the agent does not believe that she should perform an action to improve her epistemic position, but she does have beliefs that entail she should perform such an action. These two distinctions generate four possibilities when Levy's internalist rationality requirement might be satisfied:

(1) *Conscious + Direct*: An agent rationally can ϕ given that she consciously believes that she should ϕ .

(2) *Conscious + Indirect*: An agent rationally can ϕ given that she both consciously holds beliefs the contents of which entail that she should ϕ and does not believe consciously or otherwise that she should ϕ .

(3) *Unconscious + Direct*: An agent rationally can ϕ given that she unconsciously believes that she should ϕ .

(4) *Unconscious + Indirect*: An agent rationally can ϕ given that she both unconsciously holds beliefs the contents of which entail that she should ϕ and does not believe consciously or otherwise that she should ϕ .⁸²

Levy would reject indirect rationality requirements (2) and (4). On Levy's view, a successful characterization of the internalist requirement must make sense of two

⁸² Robichaud (2014: 146) does not number the four possible combinations. He also presents them in a slightly different order. I follow the numbering and order presented in Levy (2016). I do this for consistency between Robichaud's reply to Levy and Levy's rejoinder to Robichaud.

things. It must make sense of why it is reasonable to expect agents to bring their beliefs into alignment with their epistemic obligations, and it must make sense of how failing to do this involves an agent acting akratically. Situations (2) and (4) do not establish this latter requirement for Levy. An agent can have beliefs that indirectly rationalize undertaking a reasoning procedure, but, in such a scenario, the agent need not actually believe that she should undertake a reasoning procedure. When the agent fails to do the thing her beliefs entail concerning her procedural epistemic obligations the agent does not act akratically. She does not act contrary to something she actually believes she should do. Thus, as Robichaud argues, given his internalist conception of rationality, Levy would reject disambiguation (2) and (4).

What about disambiguation (3)—Unconscious + Direct? According to (3) an agent rationally can comply with an epistemic obligation given that she unconsciously believes that she should comply with that obligation.⁸³ Robichaud considers a parent who unconsciously believes that she should check the references of a babysitter prior to hiring the babysitter to watch her children. Imagine that the parent fails to do so. If so, the parent's negligent failure is akratic. The parent knew better, but she acted contrary to her (unconscious) belief that she should have checked the babysitter's references. She is culpable regarding her ignorance of the babysitter's lack of qualifications and reliability. This disambiguation gives Levy the internalist

⁸³ Robichaud does not define what he means by "unconscious belief." But, he contrasts unconscious belief with conscious belief. Perhaps a better terminology is "dispositional" versus "occurrent" beliefs. To avoid confusion, though, I continue to use the terminology adopted by Robichaud. Unconscious belief are beliefs not in a person's field of present awareness, whereas conscious beliefs are presently before a person's field of awareness.

commitment to culpable ignorance tracing back to akrasia. But, Robichaud argues that this case involving the babysitter shows that Unconscious + Direct is problematic:

Since the negligent parent's belief that she should check references remains unconscious, it does not rationalize the action from within her first-person deliberative perspective. If this unconscious belief never enters the conscious stage, then the negligent parent seems unable to guide her actions by it. Indeed, if challenged, the parent could claim with some plausibility that the fact that checking the references slipped her mind is inconsistent with her genuinely being able rationally to check them. If she actually found herself checking references, it would have only been through a kind of agential glitch of the sort that [internalists] are especially wary of mandating. (Robichaud 2014: 149)

One reason Levy cannot embrace Unconscious + Direct is his commitment to a norm like the Awareness Norm, which indicates that, "when we engage in explicit reasoning we cannot deliberately guide our behavior by reasons of which we are unaware, precisely because we are unaware of them" (Levy 2009: 736 n.16). This means that we cannot reasonably be expected to do what we are unable to do, "by an explicit reasoning procedure, a procedure we choose to engage in" (Levy 2009: 736 n.16). Thus, Robichaud is correct, and Levy (2016) concedes the point, that Unconscious + Direct will not work as a disambiguation of Levy's Internalism.

Characterization of the internalist requirement Conscious + Direct is the remaining option consistent with Levy's view. Robichaud thinks this is a possible disambiguation that captures an internalist understanding of rational action and an

akrasia requirement. However, Robichaud thinks option (1) is implausible because it is overly restrictive concerning what we can hold agents accountable for.

Understanding the notion of internalist reasons in terms of (1) makes us unable to reasonably expect that people will perform actions that are not the result of explicit, conscious reasoning. This eliminates from the scope of our reasonable expectations, as Robichaud (2014: 150) indicates, “automatic actions...that are not the immediate causal upshots of conscious deliberation.” Robichaud argues against this strong restriction of the scope of reasonable expectations by way of examples involving automatic actions:

I take it that it seems plainly reasonable to expect drivers to check their mirrors before changing lanes, and the reasonability of this expectation on a given occasion does not depend on whether the agent consciously believes on that occasion that she should check them. However, Conscious + Direct implies that the expectation that a driver check her mirrors is only reasonable in cases where she consciously believes that she should check them, for it is only then that she would have an internalist reason to check them. An agent who lacked this conscious belief and who failed to check her mirrors would not be morally responsible for harm caused by colliding with a car she did not see in the adjacent lane. Similar problems arise in cases where agents conform to their duties via automatic actions. If a parent automatically attends to her child’s immediate needs, then she lacks internalist reasons that ground both her capacity to attend rationally to her child and the fairness and reasonableness of our expectation that she do so. But, if this expectation were unfair or

unreasonable, then it would be at best odd to think this attentive parent is praiseworthy or reliable because she conforms to this expectation. The Conscious + Direct reading of internalism rules out such expectations as unfair or otherwise inappropriate. Given that agents probably only very rarely act while holding conscious beliefs about what they morally should or shouldn't do, the Conscious + Direct view entails both that we are only rarely capable of acting rationally and that it is only rarely fair to expect people to do what they should. (Robichaud 2014: 150)

The conclusion of Robichaud's second prong of attack on Levy's internalist reasons is that it makes rational actions and reasonable expectations concerning people acting in accordance with their obligations exceedingly rare. Many automatic actions are the result of cultivating dispositions to avoid unreasonably risking harming others. It is reasonable for us to expect people to drive safely by checking their mirrors prior to changing lanes, even when drivers do not consciously believe or judge that they ought to do so. When agents fail to perform such actions, and unreasonably risk harming others, their failure to consciously consider these procedural epistemic obligations does not exempt them from culpability for whatever ignorance results from their failures. Additionally, as Robichaud indicates, people are no less praiseworthy for actions that automatically lead to compliance with their procedural epistemic obligations. In fact, automatic actions may indicate that the person possesses a virtue that promotes the welfare of others or avoids risking harm.

2.4 Levy's Rejoinder to Robichaud's First Prong

In this section and the next, I will criticize Levy's rejoinders to Robichaud's two-pronged attack. Levy's view is overly restrictive and unable to sustain or justify our social practices of holding each other accountable for many of those acts and omissions for which we do in fact hold people responsible.⁸⁴

Recall that Robichaud's first prong involves accepting Levy's internalist conception of rationality, while arguing that internalism does not entail epistemically deficient agents are akratic. We can hold agents accountable for failing to avail themselves of opportunities for the improvement of their beliefs even if they do not believe all-things-considered they ought to avail themselves of such opportunities. This will occur when an agent has sufficient reasons to take advantage of an opportunity to meet her epistemic obligations, yet she has sufficient reasons of equal force that underwrite not taking advantage of that opportunity. Limiting ourselves to two choices, the agent will either take advantage of the opportunity or not. Either way, she will do something rationally (i.e., supported by reasons). Yet, if the agent fails to take advantage of the opportunity we can still hold her accountable because, given her sufficient internalist reasons, she had the capacity to take advantage of the situation for improving her epistemic position but failed to do so. When this results in

⁸⁴ I thank Hanser for requesting clarification on the nature of my criticism. Should the objection be that any view that fails to vindicate our ordinary practices is defective, which is something Levy would just deny? Or, is there something more? In reply, the objection is that responsibility as accountability, which is the conception of responsibility at issue, is a social concept that accords with our practices of holding people accountable for their wrongdoings. Any view that does not capture such ordinary practices is revisionary of the concept of responsibility at issue. There are theorists that endorse this approach to moral responsibility, such as Manuel Vargas (2012). But, that is not what Levy is up to in embracing an internalist version of responsibility as accountability.

ignorance—moral or circumstantial—the agent is culpable for that ignorance even though the ignorance itself is not the result of an akratic act.

Levy (2016) begins his response to Robichaud’s first prong by reiterating the meaning of ‘capacity’ that he is working with. Recall Levy’s chain of nested requirements: a requirement for it being reasonable to expect an agent to remedy her ignorance is that she has the capacity to correct her ignorance, an agent has a capacity to correct her ignorance only if a reasoning procedure is available to the agent that would enable her to remedy her ignorance, and this reasoning procedure is available to the agent only if she possesses internalist reasons that rationalize using such a procedure. For Levy, this ultimately means that, “it is reasonable to expect someone to do something only if they have the skills, and the opportunity to intentionally do it (or do something that entails doing it) via a reasoning procedure; where to do something via a reasoning procedure requires that the agent does not do it by chance or through a glitch in their agency” (Levy 2016). Levy argues that the case Robichaud has in mind of possessing sufficient but nondecisive reasons results in the agent improving her epistemic position by way of chance or glitchy agency. When, for Levy, is an event chancy?

Inspired by the work on luck by Duncan Pritchard (2005) and E. J. Coffman (2007), Levy (2011) develops a modal understanding of when an event is chancy:

Chance: Event E is chancy if it occurs in the actual world at t_1 , but it fails to occur in a large enough proportion of possible worlds obtainable by making no

more than a small change to the actual world at t_0 ; and the agent lacks direct control over E's occurrence. (Levy 2011: 19)⁸⁵

In figuring out whether an event E in the actual world is the result of chance or a glitch in one's agency we construct nearby possible worlds that are reachable by only making a small tweak in the conditions just prior to E happening. We press play on the causal sequence of those nearby possible worlds and see if an event other than E occurs in those worlds. If a different event occurs in a large enough proportion of those possible worlds, then event E in the actual world is a chancy event. What about the qualifier "large enough"? What constitutes a failure of E to happen in a *large enough* number of nearby possible worlds?

Levy's view of chance is an improvement over Pritchard and Coffman's view. Pritchard and Coffman think luck requires both significance and chance. The notion of luck is agent-relative. It is not purely a matter of statistics and the frequency of occurrence of events. An event must have significance to an agent for that event to be lucky relative to the agent. If I do not care whether a specific leaf falls into my backyard, then, when it falls into my backyard, it would be odd to say that it is lucky that the leaf fell into my backyard. Even if the event is chancy in that in most worlds reachable by making a minor tweak in the physical conditions (i.e., leaf position, wind direction and speed, and so on) the leaf does not fall into my backyard, the lack of significance prevents the event from being lucky when it occurs in the actual world. Levy improves on Pritchard and Coffman's view of luck by integrating the significance

⁸⁵ Quoted with the time dimension (t) signified with subscript numbers instead of symbols. I will largely bracket the issue of control in the Chance principle. The thought behind the control qualification in Chance is that an event is not chancy if the subject possesses direct voluntarily control over the event. Events subject to chance are typically outside this form of agential control.

requirement into the chance requirement. Levy makes “large enough” co-vary with the significance of the event for the agent. To see how this works, consider this case involving Russian roulette:

Russian Roulette: Samuel plays Russian Roulette with a revolver that holds six bullets. He loads a single bullet in the revolver, spins the revolver’s bullet chamber, points the gun at his head, pulls the trigger, and lives.⁸⁶

For Levy, Samuel living through playing Russian roulette is a chancy event. Even though Samuel survives the game in a high proportion (5/6) of possible worlds, this ratio is not *high enough* given what is at stake, namely Samuel’s life. Coffman considers an event chancy so long as it fails to occur in at least half of the nearby possible worlds reachable by making a small tweak in actual conditions. For Coffman, Samuel is not lucky to survive playing Russian roulette, as the event is not chancy. Instead, Samuel is fortunate. Against Coffman’s position, Levy argues:

Though I think that Coffman may be right in distinguishing between luck and fortune by reference to the chanciness condition, I doubt we can establish even a rough threshold, below which an event is not chancy enough to count as lucky, which will hold for all possible cases. It is more plausible to hold that the degree of chanciness necessary for an event to count as lucky is sensitive to the significance of that event for the agent. Surely if Samuel is lucky to survive playing a round of Russian Roulette (as Coffman must concede) he is also lucky to survive a single round with a gun which has a 49.5 percent probability of firing....On the other hand, there is a threshold, even for an event as significant

⁸⁶ This specific Russian roulette case originates in Rescher (1995). Levy (2011: 16) mentions it.

for the agent as the death Samuel risks, below which surviving is merely fortunate and not lucky (suppose that the probability of the gun's firing was 0.00001%); conversely, for relatively trivial events, the probability of the event's occurrence might have to be well below 50 per cent for it to count as lucky (suppose that the gun would give Samuel a mild and brief headache; in that case the probability of its *failing* to fire might have to be quite remote for this failure to count as lucky for the agent). (Levy 2011: 17)

Levy's view of luck improves on Coffman's view because it does not set a hard threshold above which an event counts as chancy. Even if, as in Levy's case above, Samuel plays Russian roulette and the gun is slightly more likely than not to fail to fire, Samuel living through playing the game is still lucky. Luck is involved even though he survives in (slightly) more of the nearby possible worlds than he dies (and hence is not chancy for Coffman). Alternatively, for an event of low significance for the agent to count as lucky the event may need to have a low chance of occurring (i.e., a probability well south of 50%). We can imagine a game of Russian roulette played with a toy revolver gun and bullets that if fired at one's head at close range gives one a minor headache. Imagine that 5 toy bullets are loaded in the gun. We can look at this case and a case of standard Russian roulette to illustrate the relation between significance, chance and luck for Levy. Consider the cases in this table:

	Gun Fires	Gun Fails to Fire	Levy's Verdict
Russian Roulette: 1 Real Bullet	Death Odds: 1/6 (16.7%)	Survival Odds: 5/6 (83.3%)	Lucky to survive, yet odds favor this outcome.
Russian Roulette: 5 Toy Bullets	Minor Headache Odds: 5/6 (83.3%)	No Headache Odds: 1/6 (16.7%)	Lucky no headache, yet odds against this outcome.

The odds in the chart are not hard and fast thresholds, but they illustrate how Levy thinks of luck. An event that is hugely significant for the agent, such as death, must have vanishingly small odds of occurring for the failure of its occurrence (e.g., surviving) to count as fortunate, not lucky (as when there's only a 0.00001% chance of the gun firing). When a very significant event has non-insubstantial odds of occurring, such as the 16.7% chance of death in the standard Russian roulette case, then the failure of its occurrence—which is largely more probable than not—can still qualify as a lucky event for the agent. Conversely, when an insignificant event has substantial odds of occurring, such as the 83.3% chance of a minor headache in the gun with five toy bullets case, then the failure of its occurrence—which is largely less probable than not—can qualify as a lucky event for the agent.

With this background in mind, let's consider Levy's rejoinder to Robichaud's first prong of attack on Levy's internalism. Recall, we're considering a case in which an agent has sufficient, non-decisive reasons that equally favor one of two options: either taking steps to improve one's epistemic position or not taking such steps. Robichaud claims that an agent can "rest on her epistemic laurels," be blameworthy for doing so, yet fail to act akratically. Levy replies that Robichaud's analysis only works if it is reasonable to expect the agent to improve her epistemic position. But in the case imagined by Robichaud the agent does not have the capacity to comply with such demands. Why not? The agent complying with such demands is an event that is chancy, and we cannot reasonably expect agents to do things the occurrence of which is a matter of luck.

That an agent with sufficient, nondecisive reasons for complying with her procedural epistemic obligations complies with those obligations is a chancy event in the actual world. This is because the agent's reasons are equally sufficient for rational compliance and non-compliance. Yet, this assumes the factors are insignificant. If the agent decides one way instead of the other, and significant or stable factors determine this choice, the event will not be chancy on this view.⁸⁷ Insignificant factors, such as "situational and internal primes," can generate cross-world instability for that event. The event can fail to occur in a large proportion of possible worlds reachable by making small tweaks in the state of the actual world just prior to the event occurring, which would make the event in the actual world chancy. Levy explains how such cross-world instability between the event in the actual world and events in nearby possible worlds can occur:

Perhaps the decision will be caused by a stochastic brain mechanism that cuts short deliberation between options; in that case, the agent will just plump for one option or the other. Perhaps it will be situational influences, of the kind that social psychologists have studied, that will be decisive. Or perhaps it will be the force with which a consideration strikes her, or the order in which they occur to her. All of these are, or are the product of, trivial aspects of the environment or of the agent herself. It follows that in a large proportion of nearby possible worlds, the agent will choose differently; thus her actual choice is chancy. (Levy 2016)⁸⁸

⁸⁷ Thanks to Greco for bringing to my attention the need for this qualification.

⁸⁸ I adopt Levy's use of the phrase "actual choice," but this phrase is ambiguous between choice in the actual world versus how the agent in fact chooses in the actual world. Chanciness is not determined by how the agent in fact chooses, it is determined (in this context) by whether a choice in the actual world

Situational and internal primes can cause a person to choose one way as opposed to another when her reasons between the choices are equally sufficient and forceful in supporting the respective choices. Levy takes this point a step further:

More generally, when it is genuinely the case that an agent has sufficient but not decisive reasons to choose between two or more conflicting options, chancy factors will play a decisive role in how she chooses. It is only when an agent has decisive reasons to choose a particular option that her choice will be resistant to the influence of chance events⁸⁹....In such cases, how the agent decides is subject to chance and it is not reasonable to expect (in the relevant sense of 'expect') that someone be subject to chance. (Levy 2016)

There are several ways of calling into question Levy's response to the first prong of Robichaud's objection to Levy's reply to the externalist about culpable ignorance. First, it is unclear that "trivial aspects of the environment or of the agent herself" *will* prove decisive in choosing between options that are each supported by reasons the agent takes strong enough to support two distinct courses of action. It matters why the agent takes herself to have strong reasons supporting both choices. The internal reasons that rationalize the distinct actions may involve deep aspects of the person's psychology and values. To see how this works, let's borrow an example that Robichaud uses to illustrate when a subject possesses sufficient but nondecisive reasons. This case involves someone deciding to quit smoking or keep smoking:

to meet her epistemic duty *would* fail to obtain in a significant number of nearby possible worlds. It is important to keep in mind that "actual choice" is referring to whether a choice in the actual world subject to modal evaluation is chancy, which, again, may not be how the agent actually chooses.

⁸⁹ This is because, "the stronger her reasons for a particular choice, the larger the proportion of possible worlds that differ only trivially from the actual world in which she will choose accordingly, such that if she makes that choice in the actual world, her choice is not chancy" (Levy 2016).

an agent has sufficient, nondecisive motivating reasons to quit smoking when she takes herself to have reasons that are strong enough to make quitting a rational option, but not strong enough as to decisively support quitting. Such an agent might simultaneously take herself to have reasons to continue smoking that she (wrongly, let's say) takes to be just as strong as her reasons to quit. These other sufficient though nondecisive reasons would make it rational for this agent to continue with her smoking habit. (2014: 142-3)

We can fill in the case a bit. This is one way a person might take herself to have internal reasons that simultaneously rationalize quitting smoking and keeping smoking. An agent might have motivating reasons she takes as strong enough to make quitting smoking rational because smoking is a somewhat nasty habit and she values beauty and good aesthetics. For instance, she values pleasant smells and foods that taste good. Smoking makes her not smell good and it dulls her taste buds. Yet, an equally deep value for the agent might be independence and autonomy. The agent might value not following the crowd and smoking is largely frowned upon in her culture. For her smoking is a way to not conform to societal norms. The values of aesthetics and personal autonomy might be both deeply held and equally strong.⁹⁰ They will not be overridden (i.e., the tie will not be broken) between the sufficiently strong reasons favoring acts that reflect those values by random brain glitches, the order of presentation of considerations, or other situational influences. Such chancy factors likely will not make a difference in deciding between two options, given that the

⁹⁰ Notice I am not claiming that the aesthetic value and the personal autonomy value are objectively strong values. The context of discussion is motivating reasons, which are reasons the agent takes to be sufficiently strong. Objectively speaking the reasons may not be sufficiently strong but that is not what is at issue here.

reasons are sufficiently strong for the agent because tied to things she deeply values. It is reasonable to think chancy factors will not necessarily break the tie and this counters Levy's claim otherwise when he states, "when it is genuinely the case that an agent has sufficient but not decisive reasons to choose between two or more conflicting options, chancy factors *will* play a decisive role in how she chooses" (2016).⁹¹ However, if situational and internal primes are truly insignificant features of the agent, then they will not necessarily make one option decisive over another equally weighty and sufficiently supported option, especially when the two options are tied to deeply held values.

Secondly, for the sake of argument, we can allow that the situational and agential primes generate choices in nearby possible worlds that differ from the choice in the actual world but argue that such primes do not entail the claim that the agent's choices will differ in a "large proportion of nearby possible worlds." Again, one could support this contention by way of claims regarding how some of the primes will not secure different choices in enough nearby possible worlds. The nearby possible worlds are reachable by making minor changes to the initial conditions right before the moment of choice, and such minor changes, given the sufficiency of the reasons supporting each option, fail to make a decisive option emerge in a high proportion of those possible worlds. I can strengthen this point.

⁹¹ Emphasis mine. I thank Greco for pointing out that Levy should say that "chancy factors may play" a decisive role in how she chooses. But Levy does not say this, and he is committed to the claim that they *will* play a decisive role in choice. Countering him on this point involves showing they need not play a decisive role, or, more strongly, they *will not* play a decisive role in deciding between options equally rationalized by one's internal reasons.

Even if Levy demonstrates that trivial environmental and agential primes can make the agent's choices differ in possible worlds compared to the actual world in a "large proportion of nearby possible worlds," one can hold Levy's feet to the fire and argue that that large proportion is *not large enough* in nearby possible worlds to make the choice in the actual world chancy. This holds Levy to his Chance principle. It is possible that, as in the Potter case, the choice to comply or not to comply with one's procedural epistemic obligations will have moral consequences that are, for the agent, insignificant. They will be more like a "minor headache" for the agent. As such, the event in the actual world will only qualify as chancy given that the alternative option (i.e., the minor headache⁹²) involved in failing to satisfy one's epistemic obligations and rectify one's moral ignorance has a sufficiently high probability such that in the actual world the agent would be lucky to comply with his epistemic obligations and avoid the minor headache. However, given that the minor tweaks in the form of trivial features of the agent and environment generate odds of different outcomes in nearby possible worlds, even if such tweaks result in different choices in a large proportion of such possible worlds, it is reasonable that that proportion will fail to be large enough.⁹³ It fails to make the event, which is insignificant to the agent, obtain in a *high enough*

⁹² Such as, in Potter's case, the suffering and whining of "sissies and socialists."

⁹³ Greco notes that it is difficult to establish what is reasonable in this case without saying more about what determines the relevant space of worlds. What is the similarity ordering that establishes the space of relevantly close possible worlds? And won't that be, partly at least, a function of what counts as "substantial features of the environment and agent." I agree with the thrust of Greco's comment. Given that Levy relies on a possible world understanding of whether choice in the actual world is chancy, he needs to say more about what fixes the relative space of possible worlds. I am simply raising various ways factors that Levy considers trivial can be called into question and not support what he claims, namely that only decisive reasons are resistant to chance. If motivating reasons are sufficiently strong to the agent, then, even if they are nondecisive in the actual world, it does not necessarily follow that choice in the actual world is chancy. That might depend, as Greco's question illuminates, how the possible worlds are ordered in terms of relative location to the actual world. If triviality is determining this ordering, then Levy's view might simply beg the question against his opponent—assuming as trivial what his opponent takes as non-trivial.

proportion of the worlds to make it lucky that in the actual world he would choose to comply with his epistemic obligations and avoid the headache.

Thirdly, one can argue that the primes Levy mentions are substantial features of the environment and agent. The worlds formed by making such changes will not be nearby possible worlds. Such worlds that are distant from the actual world fail to tell us whether the choice in the actual world is chancy. For instance, it is possible to question whether the “internal primes” Levy mentioned are trivial aspects of the agent. The force with which a consideration strikes an agent may reveal how valuable the consideration is in relation to the agent’s goals and plans. It may reveal her evaluative judgments or what matters to her. Additionally, the order of occurrence of considerations may reveal the importance of those considerations. We often readily remember things that matter to us. Forgetting something can reveal that we do not actually value that thing as much as we thought we did, as when we forget a friend’s birthday, and this reflects a lack of care and concern.⁹⁴ Readily remembering something need not always reflect care and concern, but it does not follow that order of remembrance or the act of remembrance is, as Levy states, “trivial aspects...of the agent herself.” Such internal primes can be nontrivial features of the agent.

We can bring together the ways of resisting Levy’s rejoinder to Robichaud’s first prong by positing a dilemma for Levy’s position:

1. Either the situational and internal primes are trivial features of the environment and agent or they are not.
2. If situational and internal primes *are* trivial features of the environment and agent, they will fail to show that the agent’s choice in the actual world is a

⁹⁴ This is a case discussed by Angela Smith (2005).

- chancy event, as they will fail to show that the worlds reachable by making minor changes in the actual world just prior to the choice will fail to make the agent's choice in the actual world occur in a *large enough* proportion of such nearby possible worlds.
3. If situational and internal primes *are not* trivial features of the environment and agent, they will fail to show that the agent's choice in the actual world is a chancy event, as such nontrivial primes will not occur in the worlds reachable by making minor changes in the actual world prior to the moment of choice.
 4. Thus, situational and internal primes will fail to show that the agent's choice in the actual world is a chancy event.

The upshot of the dilemma is that Levy has failed to impugn Robichaud's counterexample to Levy's strong internalism—a form of internalism that traces all culpable ignorance back to akratic acts.⁹⁵ Instead, given Levy's rejoinder, and pending his further defense and development of that rejoinder in response to the dilemma, it is reasonable for us to expect an agent with sufficient but nondecisive reasons for complying with her procedural epistemic obligations to take advantage of her opportunities to do so. Her deciding to do so is not equivalent to asking her to do something that is chancy or something that can only be done by way of a glitch in her agency. Rather, she has the capacity to comply with such an expectation given her possession of sufficient reasons for compliance.

In closing this section, a less restrictive form of internalism than the form embraced by Levy is capable of handling Robichaud's first prong of attack on Levy's internalism. Additionally, a form of externalism that allows internal reasons to factor into assessments of responsibility but doesn't limit internal reasons to directly

⁹⁵ To clarify where this puts us in the dialectic, given that the argument shows that Levy has failed to counter Robichaud's counterexample, the counterexample is something Levy must still contend with. This places the burden of proof back on Levy. Robichaud's case against Levy still stands.

accessible and conscious reasons, can also handle the attack.⁹⁶ In the CEO Topper case, we have someone who does not believe that she is obligated to vet her beliefs about the ethics of her business practices. She fails to scrutinize her beliefs and this (we're assuming) results in her remaining ignorant of the moral impermissibility of her standard practices. Like Potter, her business activities are reprehensibility ruthless. Unlike Potter, she possesses sufficient, yet nondecisive, reasons to comply with her procedural epistemic obligations to evaluate her economic/political ideology. Even though she falsely believes she is not obligated to comply with these epistemic obligations, she *can* rationally comply with them. And since Topper can rationally comply with her obligations, we can argue that it remains reasonable for us to demand her to form her beliefs in alignment with those obligations. As Robichaud indicates, "Topper is not in a position to claim that our expectation that she should have scrutinized her moral views is unfair on the grounds that such scrutinization was not a rational option for her. Because Topper took herself to have sufficient reasons to scrutinize her beliefs, she has internal reasons that ground her capacity rationally to scrutinize" (2014: 144). When Topper fails to comply with her obligations she does not act akratically, yet she is accountable for failing to exercise the rational capacity she possessed to comply with those obligations. This failure licenses our assigning culpability to her ignorance of the immorality of her business practices.

⁹⁶ Externalism and this broader form of internalism will align in many cases, but they will diverge in cases where internal reasons are not possessed (consciously or unconsciously) by the agent that rationalizes (directly or indirectly) the agent performing the act in question. As in the original Potter case, Potter does not hold beliefs (consciously or unconsciously) the contents of which entail that he should perform the act in question (or refrain from performing that act). The externalist thinks it is still reasonable to expect Potter to rectify his ignorance, but the broad internalist thinks it is not reasonable to expect Potter to rectify his ignorance, as he does not possess internal reasons that ground a capacity to comply with his normative obligations (i.e., what he should or should not do).

In addition to the Reasonable Expectations Claim and the Capacity Claim I need to spell out what it means to be able to rationally comply with one's procedural epistemic obligation. This is to expand on the consequent in the Capacity Claim. I can make this explicit by adding to the other claims a Rational Conformity claim:

Rational Conformity Claim: An agent can conform rationally to the expectation that she form her beliefs in accordance with some procedural epistemic obligation if the agent believes that she has sufficient, if nondecisive, reason to conform to the expectation.

Rightly understood, externalism embraces the Rational Conformity Claim. Externalism does not require culpable ignorance to always bottom out in akrasia. However, the Culpable Ignorance Principle put forward by Fitzpatrick needs to be refined in light of the Rational Conformity Claim.

Culpable Ignorance Principle:* Ignorance, whether circumstantial or normative, is culpable if the agent could reasonably have been expected to take measures that would have corrected or avoided it, given his or her capabilities and the opportunities provided by the social context, but failed to do so due to one of the following: an act of clear-eyed akrasia, a decision to not act in accordance with one's belief that one has sufficient, if nondecisive, reason to perform actions that would correct one's ignorance, or the culpable, nonakratic exercise of vices.

The Culpable Ignorance Principle* is like FitzPatrick's original externalist principle, but it allows for an additional non-akratic avenue for grounding culpable ignorance. To see how this works, let's bring back the details of Robichaud's Topper case, which was a variant of FitzPatrick's Potter case.

Topper is just like Potter in most respects. She owns a company. She has epistemic vices. She chooses not to scrutinize her moral beliefs, and this results in her ignorance about the moral permissibility of her ruthless business practices. However, at several points in her moral development, Topper believed that she had sufficient, nondecisive reasons to engage in the kind of moral self-scrutiny we expect of her. Imagine that she took her reasons to scrutinize her beliefs about her business practices and her reasons to do whatever else to be of comparable strength...Because Topper took herself to have sufficient reasons to scrutinize her beliefs, she has internal reasons that ground her capacity rationally to scrutinize. Of course, like Potter, Topper did not believe she was under any kind of obligation to reflect on her business practices, but...this fact actually does not entail that she lacks the capacity rationally to conform to such an obligation. This capacity can be grounded in her belief that scrutinization is an action that has sufficient support given what she already believes...since Topper never believed that she was obligated to scrutinize her beliefs, she did not act akratically when she chose to do something other than self-scrutinize. (Robichaud 2014: 144)

Recall that Levy grounds the capacity to fulfill a procedural epistemic obligation in having internal reasons that rationalize conformance with that obligation. And, Levy thinks an agent can only satisfy this requirement by believing that she needs to meet that requirement. When an agent like Topper has sufficient but nondecisive reasons to undertake such self-scrutiny, then she does not believe that she is obligated to perform such a self-evaluation of her moral beliefs. When she fails to inventory and scrutinize her moral beliefs, she does not act counter to what she believed she had

most reason to do. But, in fact, Topper had a duty of inquiry concerning self-scrutiny of her moral beliefs because, as even Rosen believes, such a duty is a general duty to take steps to better inform oneself as a business person because doing so decreases the threat of harm that one's business practices might otherwise pose. When Topper fails to act in accordance with her sufficient, though nondecisive, reasons to self-scrutinize she unwittingly falls afoul of that duty of inquiry. Her ignorance is culpable. So, her ignorance of the ruthlessness of her business practices fails to excuse her from moral responsibility for the harm she causes through those practices. Despite her ignorance that such practices are wrong, her ignorance does not excuse her from being blameworthy for her wrongful actions.⁹⁷

2.5 Levy's Rejoinder to Robichaud's Second Prong

Recall that Robichaud's second prong targets Levy's internalist conception of when it is reasonable to expect an agent to perform an action. For Levy this occurs when an agent has internal reasons (i.e. actual mental representations) capable of rationalizing performing an action by way of an explicit reasoning procedure, such as the act of self-scrutinizing one's opinions in light of relevant information. Levy thinks one must believe that one ought to perform such an act to make it reasonable for us to expect performance of it. Robichaud notes that there are two ways of possessing the sort of internal reasons Levy's account requires and two ways of acting from these reasons. The reasons can be conscious or unconscious, and they can directly or indirectly

⁹⁷ You might wonder if this line commits externalism to letting off the hook non-akratic slaveholders or genocidal killers, given that they did not have opportunities to have their views challenged. In chapter 3, I consider agents living in such benighted cultures. I argue that they are not let off the hook. I thus widen externalism to allow for greater grounds of culpability—moral ignorance is not an excuse condition regarding responsibility, even if the agents lacked substantial reasons to suspect that their ideologies are subject to coherent challenges.

rationalize action. Recall that this generates four ways of disambiguating Levy's internalist view of rational action:

(1) *Conscious + Direct*: An agent rationally can ϕ given that she consciously believes that she should ϕ .

(2) *Conscious + Indirect*: An agent rationally can ϕ given that she both consciously holds beliefs the contents of which entail that she should ϕ and does not believe consciously or otherwise that she should ϕ .

(3) *Unconscious + Direct*: An agent rationally can ϕ given that she unconsciously believes that she should ϕ .

(4) *Unconscious + Indirect*: An agent rationally can ϕ given that she both unconsciously holds beliefs the contents of which entail that she should ϕ and does not believe consciously or otherwise that she should ϕ .

Robichaud argues that (1)-(4) are capable of grounding our reasonable expectations regarding an agent performing an action by way of an explicit reasoning procedure. However, as Levy points out, "none of them entails that agents can reasonably be expected to perform the action only when it would be akratic of them not to perform it" (Levy 2016). Cases involving indirect rationalization of action, such as (2) and (4), make it reasonable to expect an agent to perform an action given that the contents of her beliefs entail that she ought to perform that action; however, when the agent fails to perform the action in these conditions she needn't be acting akratically as she needn't believe that she should then be executing the course of action she then fails to execute. Levy (2016) agrees with Robichaud about cases (2) and (4) involving indirect

rationalization of action. It is unreasonable to expect an agent to perform an action she cannot perform. Levy's strongly internalistic definition of "capacity" in this context implies that an agent has the capacity to perform the action only if she believes that she ought to perform the action. In cases involving indirect rationalization agents lack such a belief. So, Levy is committed to denying that agents have the capacity to perform the target actions in these contexts. Relatedly, Levy (2016) states, "While it is (sometimes) true of such agents that they can engage in reasoning such that, were they to do so, they would come to be conscious of the fact that they ought to ϕ , that fact certainly doesn't entail that they have the capacity to ϕ ." Thus, Levy grants Robichaud's claims regarding (2) and (4), as these theses are incompatible with Levy's internalism.

Levy also concurs with Robichaud's claims about (3)—Unconscious + Direct.⁹⁸ Levy agrees that when an agent only unconsciously believes that she should perform an action it is not reasonable for us to expect her to perform the action. Levy's only quibble with Robichaud on this score regards Levy's Reasonable Expectations Claim. Levy reiterates that it is unreasonable to expect an agent to perform an action when she does not have the capacity to perform that action.

This leaves option (1)—Conscious + Direct—as the disambiguation of rational action on which Levy takes his stand, as he claims that, "understanding the capacity to improve one's epistemic situation as requiring a conscious and direct relationship with one's beliefs and the action gives me what I want—the conjunction of reasonable

⁹⁸ Levy is not concurring with Robichaud that cases (2) and (4) are cases of direct responsibility. Rather, Levy concurs with Robichaud's interpretation of his position. Levy does not concur with what Robichaud ends up saying about (2) and (4), namely that we are responsible in these cases. Thanks to Hanser for pointing the need for the clarification.

expectation and akrasia—without counterintuitive restrictiveness” (Levy 2016: 269). In what follows, I present Levy’s response to Robichaud’s counterexamples to option (1), argue that Levy’s response does not work, and I present further considerations that highlight the “counterintuitive restrictiveness” of Levy’s position.

Levy responds to Robichaud by doubling down on option (1). He argues that Robichaud’s counterexamples to (1) do not involve cases wherein an agent has reasons that are consciously held and directly rationalize action. In Robichaud’s cases, agents are at best *indirectly* responsible for their actions. Recall that the regress of blameworthiness halts when we trace the ignorance back to an act or omission the agent is directly responsible for. Yet, according to Levy (2016: 269-70), in Robichaud’s cases:

The person who fails to check their mirrors when driving may be blameworthy, but they are blameworthy indirectly (in virtue of their failure to inculcate in themselves the habit of checking their mirrors). Similarly, the parent who automatically attends to his children’s need is praiseworthy for doing so (if he is), in virtue of inculcating a disposition to do so in himself, through occasional attention to his reactions and patterns of response. Once we make the direct/indirect distinction, Robichaud’s counterexamples are easily dealt with.

Levy’s response reinforces the counterintuitive restrictiveness of his view and commits him to counterintuitive bullet biting. Regarding the latter, Levy bites the following bullet, “Some may balk at the suggestion that parents are praiseworthy for being non-deliberatively attentive to their children only indirectly; due to their having inculcated in themselves appropriate habits and dispositions” (Levy 2016: n. 11). A similar

consideration applies to the case of automatic action involving omitting to check one's mirrors while driving. It is counterintuitive to say that an agent is only blameworthy for failing to check their mirrors given that they are blameworthy for cultivating the disposition to fail to check their mirrors. At best, introducing the direct/indirect distinction concerning moral responsibility is misguided. At worst, introducing the distinction relies on the falsity of the very thing at issue, namely externalism and whether the voluntary exercise of vices can ground direct responsibility and halt the regress of blameworthiness. Let me say a bit more about each of these suggestions.

Regarding Levy's response to Robichaud being misguided, consider a father who is driving with his young daughter in the passenger seat. A dog runs out in front of the car. The father slams on the brakes and automatically puts his arm across his daughter's body to prevent her from slamming her head on the dashboard. The father did not form the conscious belief that he ought to use his arm to prevent his daughter from hitting her head on the dash. Levy must claim that the absence of an internal reason in the Conscious + Direct sense means that the father lacked the capacity to rationally do as he ought and protect his daughter from harm. Yet, the father clearly had this capacity, and Levy's "rational compliance" restriction is overly restrictive in terms of what it takes to possess and to be able to exercise such a capacity. Going this route means the father is not responsible for protecting his daughter from harm, which seems misguided. Alternatively, Levy can take the route he does take and argue that the father is responsible and praiseworthy in such a case for preventing his daughter from being harmed, but the father is only indirectly responsible and praiseworthy. However, this response also seems misguided. Consider the following imagined

philosophical dialogue between a person who observed the father's action and the father:

Passerby: You're responsible and praiseworthy for protecting your daughter!

Father: Thank you, but if you don't mind me asking, why am I praiseworthy? Is it because I did what it would be reasonable to expect a father to do?

Passerby: Not quite. You're praiseworthy for protecting your daughter because you're praiseworthy for inculcating in yourself the disposition to care for your daughter due to consciously paying attention to the needs of your child.⁹⁹

Father: I'm confused. I just responded appropriately. I'm not sure you can trace my responsibility back to inculcating in myself the disposition to respond appropriately in non-deliberative situations. I'm actually somewhat negligent in consciously heeding the needs of my daughter.

Passerby: Well, then, we need to either keep tracing your praiseworthiness backwards or accept that you're actually not responsible and praiseworthy for preventing your child from hitting her head on the dashboard.

Father: Wait, what? I did do that. I saved my child from getting hurt. I had the capacity to do so and I did so.

Passerby: Yes, but again, if you're praiseworthy for that act it is only indirectly, and we haven't found a source that can ground such praise.

⁹⁹ Levy tries to make his bullet biting more palatable along these lines. He claims that you can consciously cultivate the disposition to do as you ought to do in non-deliberative situations without intentionally aiming to cultivate such a disposition or without, "intending to bring about nonconscious patterns of response" (Levy 2016: n. 11).

Father: That seems like an unnecessarily complicated and incorrect way of thinking about what just happened. Speaking of neglecting caring for my daughter, I've got to go.

A less counterintuitive thing to say about the sorts of cases that are problematic for Levy's view is that they involve direct responsibility and direct blameworthiness or praiseworthiness. They show that people can rationally do things, and act in accordance with what rationality and morality requires, without them having the capacity to do so hinge on them possessing a conscious belief that they ought to do so. We get responsibility (and praise and blame) without what Levy's view requires, namely the truth of (1)—Conscious + Direct rationalization of actions. Further, lack of compliance and the blameworthiness it entails is not traceable to an akratic act.

Lastly, Levy's response to Robichaud's second prong opens Levy to the charge that he is begging the question against Robichaud. Levy's response to Robichaud on this score assumes the falsity of an important thing at issue in the dialectic, namely whether the voluntary exercise of vices (and virtues) can generate direct responsibility for one's actions. Externalists argue that this is the case. Levy holds that unless one conceives of one's vices as vices they cannot generate an internal reason that rationalizes acting on them to rectify one's ignorance. Levy argues against the externalist by way of his claims about internalist rationality. However, now that we are "waist deep" in the dialectic the viability of Levy's internalist rationality is under scrutiny. In defending that internalist rationality as underwriting his brand of internalism Levy cannot defend its truth by way of assuming its truth. Yet this is what Levy is doing in his response to Robichaud's second prong. Defending option (1),

namely Conscious + Direct, from Robichaud's examples involving non-deliberative situations Levy assumes that the dispositions at issue in such situations cannot directly ground responsibility. The grounding must be indirect; it must be traceable to responsibility for inculcating the dispositions. But, what supports the claim that dispositions that manifest (or constitute) virtues and vices are incapable of grounding direct responsibility, and direct praiseworthiness and blameworthiness? Answer: the truth of Levy's internalist rationality. However, the truth of Levy's internalist rationality is the very thing at stake in Robichaud's second prong. Thus, Levy's rejoinder to this prong is rightly regarded as question begging.

2.6 Weakening the Internalist Commitments

In the prior two sections of this chapter I presented Levy's defense of internalism. I found his rejoinders to a two-pronged attack on his position by Robichaud to be unsuccessful. Pending further defense of his view it is reasonable to set aside Levy's internalism. Now the question becomes whether there is a view that handles Robichaud's objections and can account for his cases? Recall that a view will qualify as externalist only if it does not exclusively ground culpable ignorance in akrasia.¹⁰⁰ The view must also endorse responsibility as accountability. For the externalist the

¹⁰⁰ This specifies a necessary condition for a view qualifying as externalist. But it is not sufficient for a view to qualify as externalist that it does not require direct responsibility to originate in acts of clear-eyed akrasia. To qualify as externalist the view must also countenance external reasons (i.e., motivating reasons that do not directly or indirectly, consciously or unconsciously rationalize action). When a subject has no internal reason to not act on their vices, but they have an external reason to not act on their vices because (totally unbeknownst to them) so acting violates a procedural epistemic obligation, then the externalist still thinks the agent has a reason to rectify her ignorance. An internalist, even one broadly permissive about what qualifies as an internal reason, will jump off the ship. At his point, the internalist thinks there are no reasons sufficiently capable of guiding the agent to rectify his ignorance. It is unreasonable to expect an agent to do so. The internalist requires some conception of the vice as a vice, however remote or indirect, to rationalize not acting on the vice and instead to act to better one's epistemic position (or at least not worsen it further).

voluntary exercise of vices is a source of original responsibility. This is how the externalist can handle Robichaud's second prong against strong internalism. Further, the externalist can take on some internalist commitments to successfully handling Robichaud's first prong against Levy's internalism. The key is to relax the strength of the internalist commitments. One way to do this is to not insist on *direct* rationalization of action by way of beliefs. In fact, Robichaud suggests this as a possible way of handling the case involving Potter and arguing that Potter has internalist reasons that rationalize him engaging in reasoning to rectify his ignorance:

Interestingly, the Conscious + Indirect and Unconscious + Indirect accounts of internalist reasons might imply that Potter as described by FitzPatrick actually has internalist reasons to scrutinize his moral beliefs after all and, thus, that it is indeed reasonable to expect him to manage his epistemic affairs differently. If Potter believes both that awareness of disagreement ought to give one pause, and that people disagree with him about how his employees should be treated, then he has beliefs from which he can reason to the conclusion that he should scrutinize his beliefs about what he owes to his employees. This possibility seems not to be ruled out by the story as it is told, but, even if it is, it remains the case that certain epistemically vicious agents can reason their way to true beliefs about what they should investigate and how. When these agents successfully engage in this reasoning, they reveal that they indeed held beliefs that indirectly rationalized our expectation that they self-scrutinize. The Conscious + Indirect and Unconscious + Indirect versions of internalism establish that such agents—and maybe even Potter—have the rational capacity to engage in self-scrutiny even in the absence of their engaging in reasoning

that would make the entailments of their beliefs explicit. It follows from what I've called the Reasonable Expectations Claim and the Capacity Claim that there may be nothing unfair or unreasonable about maintaining that certain epistemically vicious agents are culpably ignorant, even if their epistemic mismanagement is plainly nonakratic. (Robichaud 2014: 147-8)

I think Robichaud's comments are suggestive.¹⁰¹ I will argue that Potter has internalist reasons that indirectly rationalize self-scrutiny of his moral beliefs. Remember that indirect rationalization of action comes in two forms:

(2) *Conscious + Indirect*: An agent rationally can ϕ given that she both consciously holds beliefs the contents of which entail that she should ϕ and does not believe consciously or otherwise that she should ϕ .

(4) *Unconscious + Indirect*: An agent rationally can ϕ given that she both unconsciously holds beliefs the contents of which entail that she should ϕ and does not believe consciously or otherwise that she should ϕ .

Indirect rationalization of action is indirect because one does not hold beliefs that directly indicate that one ought to perform an action, but one does hold beliefs—either consciously or unconsciously—that indirectly indicate that one ought to perform an

¹⁰¹ I think Robichaud's comments are suggestive in outline but not in exact detail. Instead of claiming, as Robichaud does, that indirect rationalization may occur in the Potter case given Potter possessing a belief about disagreement, I focus on Potter's belief concerning the general importance of critical reflection on one's opinions. Why go this route? Given what we know concerning Potter it is highly unlikely that, "Potter believes both that awareness of disagreement ought to give one pause, and that people disagree with him about how his employees should be treated," as suggested by Robichaud. Potter is unlikely to believe that disagreement with the nature of his business practices is grounds for vetting his beliefs concerning those practices. He writes off all such disagreement as the "whining of socialists and sissies." Potter does not regard such people as epistemic peers. And beliefs concerning the presence of reasonable disagreement among peers is the sort of thing that could indirectly rationalize scrutinizing or altering one's opinions and their related practices.

action. They indirectly rationalize action because they entail that one ought to perform the action, but the entailments have not been explicitly inferred. The entailments are not part of one's belief set. Although there is a way that one could infer those entailments, namely by reasoning about the contents of one's beliefs and competently deducing a belief indicating that one ought to perform a specific action. Though Levy's overly-restrictive internalist commitments cannot admit as acceptable the indirect rationalization of action, as akrasia requires the direct rationalization of action,¹⁰² I can account for the cases involved in Robichaud's second prong of attack by allowing for indirect rationalization of action.¹⁰³ One of the cases Robichaud raised against Levy's view is the case involving a negligent parent who unwittingly leaves her kids with an unreliable babysitter. This is because the parent neglected to check the sitter's references. Though Levy's view will not be able to handle such a case, as disambiguation of internalist reasons in terms of options (2) and (4) will fail to secure an akrasia requirement, which is what makes internalism the view that it is, externalism can embrace (2) and (4) and account for Robichaud's case. Robichaud explains how this might go:

according to Conscious + Indirect and Unconscious + Indirect, this agent can rationally ϕ —she can reason her way to compliance with the obligation to ϕ given that she holds beliefs that can function as premises in an episode of

¹⁰² Namely, akrasia requires belief that one ought to all-things-considered perform an action yet one acts contrary to what one believes one ought to do.

¹⁰³ To clarify, Levy's view cannot secure the result that the agents are directly responsible. What is distinctive of my response is that it shows that an externalist can, still using internalist materials, account for the direct responsibility of an agent. Though a broad internalist can account for these cases in a similar way, the broad internalist will diverge from the externalist in cases like the original Potter case, as explained in a previous footnote. What my perspective shows is that an externalist can handle these cases in a way that makes use of internalist ideas, namely by relaxing what we mean by internal reasons.

reasoning that leads her to believe she should ϕ . That she has this rational capacity seems to clear the way for someone to argue that it would be reasonable to expect her to ϕ . But, this just demonstrates that both the Conscious + Indirect and Unconscious + Indirect interpretations of internalist reasons fail to establish the akrasia requirement of reasonable expectations, which constitutes the heart of the [internalist's] take on culpable ignorance and their reason for thinking that ignorant agents are only very rarely morally responsible. If a parent, say, holds beliefs that entail that she should check a babysitter's references—perhaps she can quickly infer from beliefs she holds to the belief that she should call the references—it follows from Conscious + Indirect and Unconscious + Indirect that she can rationally check them. Given this rational capacity, it seems not unreasonable to expect her to check the references. However, if she never actually engages in the reasoning that would make her aware of this obligation, her failure to check the references would be nonakratic. This parent's ignorance about the babysitter's experience and reliability would not trace to akratic belief mismanagement. It appears that any species of internalism about reasons for action that allows for indirect rationalization of potential moral duties will fail to secure the akrasia requirement. (Robichaud 2014: 147)

Given that Conscious + Indirect and Unconscious + Indirect can account for a babysitter type of case in a way that does not ground culpability in akrasia, the externalist can embrace an account of why in such a case the agent is directly responsible and do so in terms of internal reasons. The externalist allows that culpable

ignorance might in some cases trace back to akrasia, but the externalist does not require all culpable ignorance to originate in akratic episodes. Allowing for indirect rationalization of action explains how one can possess internal reasons that enable one to rationally comply with reasonable expectations and failure to so comply can still ground culpability for ignorance in the absence of an akratic act.

Now that I have handled Robichaud's cases that featured in the second prong of his attack on Levy's internalism I turn to relating this to externalism and addressing the Potter case. Why does Potter have conscious or unconscious beliefs that indirectly rationalize him complying with his procedural epistemic obligations to do things to scrutinize his moral beliefs? Remember that in section 2.1 I elaborated on externalism's response to the skeptical challenge to moral responsibility by showing how FitzPatrick contrasts the internalist's strong knowledge requirement involving clear-eyed akrasia with a weaker knowledge condition:

Weaker Knowledge Condition: It is reasonable to expect one to rectify one's ignorance if one knew the general epistemic importance of subjecting one's beliefs to critical scrutiny, so it could be reasonably expected to occur to one that one's behavior with regard to one's moral beliefs was epistemically irresponsible. (FitzPatrick 2008: 607; additions and modifications mine)

FitzPatrick argued that Potter satisfied this Weaker Knowledge Condition. Given the skills and knowledge possessed by Potter as a successful businessperson,

he would surely have understood the importance of genuinely critical reflection in other spheres, such as analysis of stock market and interest rate trends, and

the epistemic dangers of relying for such information on only one firm or group of analysts with vested interests without taking opposing argument seriously. Presumably, then, he would have insisted on high standards of critical reflection in the sphere of his work, and yet he did not make the obvious extension to the moral sphere—the recognition that one is unlikely to arrive at an accurate moral view by uncritically accepting the received opinion of a narrow group of people with vested interests in a certain set of answers and insulating oneself from open, critical discussion incorporating different points of view and forms of experience. (FitzPatrick 2008: 607)

Potter possessed an internal reason that indirectly rationalized taking steps to scrutinize his moral opinions. Potter possessed knowledge concerning the general epistemic importance of subjecting his beliefs to critical scrutiny. Though Potter did not infer an important entailment of this belief, as he “did not make the obvious extension to the moral sphere,” Potter nonetheless consciously or unconsciously held a belief that entailed that he should scrutinize his moral beliefs. If Potter had reasoned appropriately, he could have possessed the belief that he should subject his moral opinions to critical scrutiny. Potter could rationally comply with the expectation that he should satisfy his procedural epistemic obligations with regard to his moral beliefs; he could reason from the justified true belief he possessed concerning the epistemic value of subjecting one’s views to critical pressure to the entailment that he should subject his moral beliefs to critical pressure. Thus, there was a way of reasoning from beliefs Potter possessed to their “extension to the moral sphere” or what they entail in that regard; there was available to Potter a way of reasoning in compliance with the obligation that he subject his moral beliefs to such scrutiny to arrive at true moral

beliefs and avoid false ones concerning the moral status of his business practices and their impact on the wellbeing of people living in his community. As such, it is reasonable for us to expect that he should have done so, and Potter is culpable for his ignorance due to his failure to rationally comply with this expectation.

The considerations above suggest the following modification to the Rational Conformance Claim:

*Rational Conformance Claim**: An agent can conform rationally to the expectation that she form her beliefs in accordance with some procedural epistemic obligation if the agent believes that she has sufficient, if nondecisive, reason to conform to the expectation, or she has a justified true belief that indirectly rationalizes conformity with that expectation.¹⁰⁴

Rational conformity to an expectation that an agent form her beliefs in alignment with a procedural epistemic obligation now includes the following alternative sufficient condition: an agent has a consciously or unconsciously held justified true belief the contents of which entail that the agent should comply with the expectation. This Rational Conformance Claim* suggests a modification to the Culpable Ignorance Principle*, but I will hold off on modifying the externalist's main principle until the final modification to the Rational Conformance Claim has been made.

¹⁰⁴ Though justified true belief is not equivalent to knowledge, as attested by Gettier cases, I do not intend to take a stand on the proper amendment of JTB that is equivalent to knowledge. JTB secures knowledge in standard cases whereby there is no epistemic funny business present. The cases at issue in this literature do not involve epistemic funny business whereby what justifies the belief comes apart from what makes the belief true.

2.7 Externalism's Akrasia Condition

The Culpable Ignorance Principle* includes in one of its disjuncts the possibility of akrasia grounding culpable ignorance.¹⁰⁵ Now that I have relaxed the internalist requirements on rationality it is worth closing this chapter by discussing the akrasia condition that externalism allows to ground culpable ignorance. Externalism should not *solely* substantiate the possibility of akrasia in terms of the *Conscious* + Direct rationalization of action. Doing so would result in falling back into Levy's overly restrictive brand of internalist rationality. Instead, externalism should account for the possibility that akrasia may result from the *Unconscious* + Direct rationalization of action. However, there is an initial worry related to going this route.

Considering Unconscious + Direct as an interpretation of internalist reasons capable of grounding an akrasia condition appears to neglect the motivation behind internalism.¹⁰⁶ The motivation for insisting that culpable ignorance be traceable back to beliefs that directly rationalize the actions that the agent neglected to perform to correct her ignorance is that the agent was capable of being guided by such reasons. She actually possessed the belief that she ought to perform the actions she failed to perform. This makes it appropriate to demand that, given the first-person perspective of the agent, she should have acted based on the internalist reasons she possessed. However, when the beliefs that directly rationalize such actions are unconscious this seems to defeat the purpose of allowing internalist commitments into one's theory of culpable ignorance. How can one be guided by reasons that one only possesses

¹⁰⁵ Culpable Ignorance Principle* was presented at the end of section 2.4.

¹⁰⁶ This is similar to the original worry, namely whether there can be direct responsibility when reasons are not conscious.

unconsciously?¹⁰⁷ Robichaud brings out this worry when discussing the Unconscious + Direct disambiguation of internalist reasons. He does this in the context of considering the viability of Levy's restrictive view of internal reasons. Robichaud raises this worry by way of the case involving the parent that neglects to check the references of the babysitter:

Since the negligent parent's belief that she should check references remains unconscious, it does not rationalize the action from within her first-person deliberative perspective. If this unconscious belief never enters the conscious stage, then the negligent parent seems unable to guide her actions by it. Indeed, if challenged, the parent could claim with some plausibility that the fact that checking the references slipped her mind is inconsistent with her genuinely being able rationally to check them. If she actually found herself checking references, it would have only been through a kind of agential glitch of the sort that [internalists] are especially wary of mandating. (Robichaud 2014: 149)

In section 2.4 I argued against Levy's claims regarding chanciness made in his rejoinder to Robichaud's first prong of attack. I did so by using Levy's own view of luck as a way of undermining his response to Robichaud. Returning to that view of luck we

¹⁰⁷ Thanks to Greco for requesting clarification on the status of this worry. Greco rightly wonders, "Isn't it uncontroversial that unconscious beliefs can drive inference and action, function as reasons?" It is uncontroversial that people are in fact moved to do things and say things by unconscious motivating reasons. But, to clarify the context of this worry, the question is whether it is reasonable to expect agents to act based on such reasons, such that were they to fail to so act (and violate a duty of inquiry) that it would be permissible, right, or fair to hold that against them? Would it be fair to judge that their ignorance is culpable? One reason that might underwrite the claim that it is not fair to hold them accountable for failing to act based on unconscious motivating reasons is the fact that such reasons only unreliably motivate people to act in accordance with their procedural epistemic obligations. I do not necessarily endorse this line, but it does lend initial plausibility to the worry that there is a disconnect between what we expect of an agent and how things (reliably) appear in her first-person deliberative perspective.

can ask the following: Would the negligent parent in possession of unconscious beliefs that directly rationalize checking the babysitter's references only do so by chance or some sort of agential glitch?

Recall that, for Levy, lucky events are a matter of chance and significance. More specifically, significance variably sets the threshold at which an event counts as lucky. In the standard Russian roulette case, though the odds favor the gun failing to fire (83.3%) Samuel is still lucky to survive. This is because survival is significant for Samuel. When an event is insignificant for an agent the threshold at which an event is lucky is much lower regarding chance. In the Russian roulette case involving a toy gun the odds are against the outcome of getting no headache (16.7%). If one fires the gun at one's head and happened to get the one chamber out of six not containing a toy bullet, then one is lucky that one gets no headache. In deciding whether the event of the parent checking the babysitter's references is lucky, despite only possessing an unconscious belief that directly rationalizes doing so, we need to know the odds that the event would occur and the significance of the event occurring for the parent. If the parent is *laissez-faire* and not concerned with checking babysitter references, then if she finds herself checking the references, and this event is somewhat less likely than not to occur, then the event would not qualify as lucky.¹⁰⁸ Because the event is

¹⁰⁸ Hanser raises questions regarding Levy's use of the notion of significance. He asks whether Levy is an internalist about significance and whether the luck should be interpreted subjectively or objectively? Hanser also asks about whether Levy endorses a distinction between direct vs. indirect significance? To my knowledge, Levy does not explicitly endorse these distinctions or a line on them, but it is an interesting suggestion that Levy pursue these distinctions in potentially providing a response to myself and Robichaud. For instance, if there is indirect significance, then the parent cares about the welfare of her child. They are not *laissez fair* about how they parent, they just do not connect the dots between their caring and checking the baby sitter's references. If so, then there would be indirect significance of checking the references and doing so would be indirectly subjectively significant for the parent. This would then count as a case where it is significant (so not lucky) that she checks the references. This contrasts with a parent that is more thoroughgoing in how *laissez-faire* they are about the degree to

insignificant to the parent the odds of it occurring would have to be very low—the event would have to be largely improbable—in order for her checking the references to be a lucky occurrence as a result of glitchy agency. The parent would have to not really care whether she checks the references and be very forgetful or unlikely to bring to consciousness beliefs that are unconscious. What this shows is that how things appear from a first-person perspective is not the determining factor regarding whether an action is the result of an agential glitch. Thus, how things appear from the first-person deliberative standpoint does not determine whether the agent’s internalist reasons rationalize her actions.

As you may recall, Levy limits internal reasons to beliefs that can consciously and explicitly guide an agent by way of a reasoning procedure. This is why, for Levy, internalist reasons that are Unconscious + Direct cannot factor into an agent having the capacity to comply with her procedural epistemic obligations. However, I relaxed these assumptions. Altering the example above, imagine that a parent really cares about checking a babysitter’s references, yet she is unlikely to do so. Suppose that she is motivated by an unconscious belief to check the candidate’s references, and the unconscious nature of this belief renders her action “lucky” in the intended sense. The externalist can accommodate this case by likening it to more traditional cases of moral luck. Even if remembering to check the sitter’s references is not directly in the agent’s control—and there is low antecedent probability that she will exercise due diligence—it is still fair to evaluate her based on reasons that rationalize her complying with the epistemic obligations in place. The reasons for adopting this stance resemble those

which they are protective of the welfare of their child. If so, then checking the references for such a parent would be lucky.

invoked to justify holding people accountable for events that are subject to moral luck.¹⁰⁹ The agent's protest that looking into the sitter's references slipped her mind would not be an exculpating plea. Given that remembering to check the references was significant for the parent, failing to do so, even when doing so would be subject to chance, remains something she is rightly held accountable for failing to do. The only relevant necessary condition on the propriety of holding her responsible is her having the capacity to do what we require of her. And, *ex hypothesi*, she had this capacity.

Another consideration that dampens the worry that allowing unconscious reasons into the picture of what is capable of grounding an akrasia requirement on culpable ignorance is going against the motivation behind internalism is a helpful contrast discussed by Robichaud:

There is an important difference between an agent who unconsciously believes that she should ϕ and an agent who lacks this belief entirely but is in the presence of (even overwhelming) evidence that she should ϕ . According to Unconscious + Direct, ϕ -ing is internalistically rational only for the former agent—she need only engage in a bit of reflection and introspection in order to bring her belief that she should ϕ to consciousness. By contrast, the latter agent must be both receptive and reactive to the evidence she has if she is to come to believe she should ϕ . Absent such evidence acquisition and processing, her reasons to ϕ will be unavailable to her as she deliberates about what to do. (Robichaud 2014: 149)

¹⁰⁹ For a recent defense of moral luck see Hanna (2014).

Unconscious beliefs that are dispositionally held are typically capable of being brought to mind and used in a reasoning process.¹¹⁰ This differs from subconscious or repressed beliefs.¹¹¹ Such beliefs are typically inaccessible upon reflection. It is a less demanding expectation, and perhaps a more reasonable one, that agents should access and utilize the beliefs they already possess as opposed to the expectation that agents acquire and process evidence and beliefs they do not already have. I think this point is debatable. But, there is something to the contrast that Robichaud raises. It does seem that an agent in an evidence-rich environment must have the cognitive capacity for reasons-responsiveness in a more robust form than an agent that already has the evidence and just needs to be responsive to it. However, we can imagine that the evidence in the evidence-rich environment is not possessed and the belief that one should gather evidence remains unconscious. When this happens, as Robichaud concedes, “from the deliberative perspective, an unconscious belief that one should ϕ that never comes into the agent’s conscious awareness is for the internalist effectively as unavailable as the unprocessed evidence that one should ϕ ” (Robichaud 2014: 149). This will be a problem for a strict internalist that rationalizes action only by way of an explicit reasoning process the agent can engage in. This is what leads Robichaud to

¹¹⁰ As previously mentioned, the distinction between conscious and unconscious beliefs is most often aligned with the distinction between occurrent and dispositional beliefs (i.e., beliefs one is presently aware of versus beliefs that one is not presently aware of but is disposed to act in accordance with). Yet, the conscious/unconscious distinction cuts across the occurrent/dispositional distinction. This is why I do not equate unconscious beliefs with dispositional beliefs, and I fully specify the target mental states as “unconscious beliefs that are dispositionally held.”

¹¹¹ I am using the term “subconscious” to refer to beliefs that are not “unconsciously” held yet easily accessed upon proper prompting—whether through introspection or environmental priming. Beliefs in the subconscious, as I am using the term, are not easily accessed or brought into awareness. This is why I also use the term “repressed” beliefs, though that term is somewhat misleading, as it may imply willful repression. Though repressed beliefs are often buried in the subconscious due to traumatic events. Such repression is not willful by the person but is a self-protective mechanism largely operating outside manual cognitive control.

think Unconscious + Direct, though capable of grounding an akrasia requirement, is unavailable as an option for Levy's account. You may also recall that Levy agrees with this point. However, as Levy admits, all views must bite some bullets. So, my less restricted form of internalist rationality can allow for Unconscious + Direct to ground an akrasia requirement. If it remains counterintuitive that this way of disambiguating internalist reasons captures the spirit of internalist motivations, then I am prepared to accept that counterintuitiveness. In fact, I am prepared to embrace that counterintuitiveness and take things a step further. I am prepared to claim that the fact that Unconscious + Direct reasons can rationalize action, and yet such reasons are as unavailable to the agent as reasons that are not possessed but are ripe for the taking in the environment, makes it plausible that external reasons can rationalize action. Let me expand on this point.

In this chapter I have made sense of externalism as a contrast to internalism. I have made sense of the Potter case by way of him possessing internal reasons that indirectly rationalize him complying with the epistemic obligations that he violates. I have stayed in the sandbox of internalist rationality. But, my aspirations are greater. That is, I want to articulate and defend a theory of culpable ignorance that is truly externalist. Externalist theories do not ground all culpable ignorance exclusively in an akrasia requirement. Strong forms of externalism go further. I want to allow for the sort of situation that Robichaud mentioned whereby an agent is in an evidence-rich environment and the reasons in that environment rationalize actions the agent might perform whether the agent in question "possesses" these external reasons. Suppose that an agent is in an evidence-rich situation and does not acquire the evidence she could have acquired, and her reasons for complying with her obligations are not

cognitively accessible to her as she decides what she will do because she has not “acquired” those reasons. The failure of such an agent to do what the external reasons suggest is capable of grounding any blame we might level at her concerning her ignorance. There is a sense in which the agent should have acquired that evidence and is culpable for that failure. Establishing this claim would help strengthen the Context/Capabilities claim, as Potter, for instance, was in an evidence-rich environment. As the claim indicates, “There were 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” (FitzPatrick 2008: 605). Thus, there were externalist reasons that underwrite reasonable expectations concerning Potter complying with his procedural epistemic obligations.

Before proceeding to the next chapter and summarizing this chapter let me bring together the considerations above by revising the Rational Conformity Claim* to account for the akrasia requirement:

*Rational Conformity Claim***: An agent can conform rationally to the expectation that she form her beliefs in accordance with some procedural epistemic obligation if the agent believes that she has sufficient, if nondecisive, reason to conform to the expectation, or she has a justified true belief consciously or unconsciously held that indirectly rationalizes conformity to the expectation, or she has a belief consciously or unconsciously held that directly rationalizes conforming to the expectation.

Rational Conformity Claim** entails modification to the Culpable Ignorance Principle*:¹¹²

*Culpable Ignorance Principle***: Ignorance, whether circumstantial or normative, is culpable if the agent could reasonably have been expected to take measures that would have corrected or avoided it, given his or her capabilities and the opportunities provided by the social context, but failed to do so due to one of the following: an act of clear-eyed akrasia accompanied by beliefs held consciously or unconsciously that directly rationalize conforming with one's epistemic obligations, a decision to not act in accordance with one's belief that one has sufficient, if nondecisive, reason to perform actions that would correct one's ignorance, the culpable (nonakratic) exercise of vices, or not acting in accordance with a justified true belief consciously or unconsciously held that indirectly rationalizes conforming with one's epistemic obligations.

The Culpable Ignorance Principle** is a disjunctive principle that allows culpable ignorance to arise from, and be grounded in, several different processes. It avoids an overly narrow and restrictive construal of internalist rationality. It accommodates all four ways of disambiguating what it means for an agent to have internal reasons that rationalize actions in conformance with procedural epistemic obligations. It makes sense of why agents are culpable for their ignorance, and, as a result, it captures why they are culpable for any wrongful acts they performed from a state of circumstantial or moral ignorance.

¹¹² For ease of reference, Culpable Ignorance Principle* was presented at the end of section 2.4.

2.8 Conclusion

This chapter involved canvassing externalism's response to the skeptical challenge to moral responsibility. I considered an objection to externalism raised by Neil Levy and a response to Levy by Philip Robichaud. I explained Levy's rejoinder to Robichaud, and I argued that Levy's rejoinder to Robichaud failed. I did this by positing a dilemma for Levy's response. I also argued that the view of rationality underlying Levy's rejoinder was implausible. I concluded this chapter by bolstering the main principle endorsed by externalism about the epistemic condition on moral responsibility.

CHAPTER 3: MORAL IGNORANCE IS NEVER EXCULPATORY

Responsibility is...one aspect of the identity of character and conduct. We are responsible for our conduct because that conduct is ourselves objectified in actions.

–John Dewey “*Outlines of a Critical Theory of Ethics*”

3.0 Introduction

Blameless ignorance about non-moral facts excuses one from responsibility for wrongdoing. If through no fault of my own I don't know that a light switch, when flipped, will ignite a bomb in my neighbor's house, I'm not blameworthy for flipping the switch and harming my neighbor. Had I known of the connection between the bomb and the switch I never would have flipped it. What about *moral* ignorance? Does moral ignorance ever excuse a person from responsibility for wrongdoing?

Ignorance of moral facts is intuitively less exculpatory. If I know that my neighbor has a torture chamber in his house and sometimes lures innocent people in and tortures them, but I don't know that torturing innocent people is wrong, my moral ignorance doesn't excuse me from responsibility for failing to notify the authorities.

Does moral ignorance ever excuse? Some theorists argue that moral ignorance is never exculpatory.¹¹³ They embrace this thesis:

¹¹³ The main proponent of this position is Elizabeth Harman (2011, 2014). Yet the view is endorsed with various qualifications by Moody-Adams (1994), Arpaly (2003), FitzPatrick (2008), Guerrero (2007), and Talbert (2013). Aristotle is a historical source of inspiration for support of this position (*Nicomachean Ethics*, 1110b-1111a).

*No Excuse. Moral Ignorance never excuses.*¹¹⁴

A couple of qualifications of the No Excuse thesis to keep in mind. Implicit in the literature on No Excuse is that it is possible to discover the moral truth. The discussion of No Excuse involves looking at cases where the moral truth is difficult, but not impossible, to discover. In this paper, I set aside the question of whether moral ignorance fails to excuse in cases where the moral truth is impossible to discover. What type of possibility is at issue? The possibility at issue is not physical or conceptual possibility simpliciter. For instance, the Korowai are hunter-gatherers that live in the remote jungle of New Guinea. They practice cannibalism as part of their system of retributive justice. They believe that people can be overtaken by an evil demon called the khakhua. If so, the tribe needs to kill and eat the khakhua to exact justice on the evil demon that took over the person. While it is conceptually possible simpliciter that a Korowai could realize that such ritual killings and eating of people is morally wrong, it is not possible in the sense at issue. Given the tribe's remoteness and lack of defeaters in their cultural environment, it is not possible in the relevant sense to recognize that cannibalism as justice is wrong.¹¹⁵ The possibility at issue concerns conceptual possibility tied to the presence of defeaters of moral falsehoods within the cultural context. In the cases at issue in this literature, the culture largely embraces moral falsehoods and misleading evidence abounds. But, it is conceptually possible within the cultural context to discover the moral truth by way of dissenting opinions.

¹¹⁴ Wieland (2017b: 147) dubs this the "No Excuse" thesis.

¹¹⁵ Interestingly, there has been some contact with the tribe from Western visitors. But the contact has been very minimal. People who visit the tribe fear being cannibalized, though this would not likely happen. The tribe highly esteems their system of justice and the practice cannibalism that exacts that justice. It is unlikely that the few outsiders that have contacted the tribe have morally questioned the practice to the tribespeople. Thus, defeaters of the moral permissibility of the practice is non-existent within the cultural context.

It is just difficult to do so because the prevailing winds favor the moral permissibility of the bad behavior at issue. So, the sense of possibility at issue is conceptual possibility a result of cultural possibility.

One might wonder: Does the No Excuse thesis overreach? Aren't there cases where moral ignorance excuses an agent from moral responsibility?

Cases that readily test whether No Excuse is true involve cultures in which a moral truth was not widely known. In such a time and place, were an agent to conduct a thorough inquiry motivated by a desire to know the truth, they would not necessarily arrive at knowledge of the relevant moral truth. Slavery in ancient Greece or sexism in America in the 1950s are examples used to illustrate this point.¹¹⁶ Slavery and sexism in those cultures were not widely regarded as wrong. As such, moral ignorance of the wrongfulness of slavery and sexism by people living in those cultures might excuse them from blameworthiness for keeping slaves or acting in a sexist manner. Thus, moral ignorance sometimes excuses, and the No Excuse thesis is false.

Contrary to the previous reasoning, I will argue that the No Excuse thesis is true. I provide a new defense of No Excuse by paying attention to the nature of moral responsibility and blameworthiness. Underlying the most detailed defense of No Excuse is a conception of moral responsibility. Elizabeth Harman (2011, 2014) defends No Excuse while endorsing a narrow conception of responsibility, which is a version of attributability put forward by Nomy Arpaly (2003). On this view, lacking proper desire or care for what is in fact morally significant makes one blameworthy. And failing to believe the moral truth is often a way of not caring properly about what

¹¹⁶ For such examples, see Gideon Rosen (2003, 2004).

matters morally. Grounding blameworthiness in a care-centered model of responsibility undermines Harman's defense of No Excuse. To avoid this result, I broaden the notion of responsibility to include attributability understood as judgement-sensitivity or answerability.¹¹⁷ Moral ignorance is never exculpatory because such ignorance never excuses an agent from being answerable for their actions and omissions.

What hangs on whether No Excuse is true? The truth value of the No Excuse thesis is relevant to the present-day. Sometimes it's hard to attain moral knowledge. There's no guarantee that thinking hard about what we owe to each other, to animals, and to the environment will yield knowledge. Thoughtful people who care about morality sometimes arrive at false moral views. Contemporary moral issues are highly complex and misleading evidence abounds. Issues involving business ethics, environmental ethics, animal rights, privacy rights, abortion, and immigration exhibit this pattern. A person might think hard about what morality demands and end up with a false moral view. If No Excuse is true, as Harman states, "being caught in the grip of a false moral view is not exculpatory."¹¹⁸ We are on the hook morally for more than we realize. Our moral ignorance will not excuse us from blameworthiness for wrongful acts done from such ignorance.

¹¹⁷ This is a broadly rationalist approach, as opposed to a more narrowly affect-based approach. Proponents of the rationalist approach include Hieronymi (2014), Scanlon (1998), Smith (2012), Strabbing (2016), and Talbert (2012). Judgement-sensitivity or answerability is more amenable to rationalism than sentimentalism because cognitive judgments and rational relations between those judgments is the focus. The views just mentioned do not make emotions and desires the central feature of attributability and answerability. The views are more rationalistic in this regard.

¹¹⁸ Harman (2014: 1).

3.1 Obligations and Blameworthiness

In this section, I get on the table a defense of the No Excuse thesis by Harman (2011, 2014).¹¹⁹ To better understand Harman's view it is useful to contrast her view with orthodox commitments endorsed by most views.

Most theorists think one is blameworthy for one's wrongful act done from ignorance that the act is wrong *only if* one is blameworthy for the ignorance from which one acted.¹²⁰ Absence of blameworthiness for one's ignorance excuses one from blameworthiness for the unwitting wrongful act. When is someone blameworthy for their ignorance? To use a term put forward by Holly Smith (1983) in a hallmark paper on this topic, ignorance that is blameworthy traces back to a 'benighting act'—a failure to improve one's epistemic situation or a positive impairment of one's epistemic situation. A benighting act can occur when one fails to gather evidence relevant to one's action, and this can occur when one exhibits intellectual vices of laziness, arrogance, dismissiveness, and closed-mindedness.¹²¹ In such cases, blameworthiness for one's unwitting wrongful act traces back to a benighting act or poor management of one's beliefs, such as failure to gather evidence relevant to updating them.

¹¹⁹ Harman's view is a response to the view of William FitzPatrick (2008), which finds its root in James Montmarquet (1999). For FitzPatrick, one's ignorance need not bottom out in an act of clear-eyed akrasia, which is knowing wrongdoing in full view of the relevant facts. One can be directly responsible for one's ignorance in virtue of belief mismanagement, namely the voluntary, non-akratic exercise of vices.

¹²⁰ Wieland (2017a: 10) calls this "Claim 2" that is endorsed by orthodoxy. Blameworthiness for ignorance (factual or moral) is necessary for blameworthiness for unwitting wrongful action.

¹²¹ It can also trace back to an act of clear-eyed akrasia. For instance, one might judge that all-things-considered, given the facts, one ought to research the moral dimension of one's business dealings, but one acts contrary to one's better judgment and omits to gather the relevant information via research or conversation. As previously discussed, internalists require all culpability to trace back to acts of clear-eyed akrasia. According to theorists such as Zimmerman (1997) and Rosen (2004), this view leads to epistemic skepticism about judgments concerning responsibility.

Blameworthiness traces to violations of procedural epistemic obligations—obligations to gather information relevant to our actions.¹²²

Harman's view overcomes a deficiency in views that ground blameworthiness for one's ignorance in belief mismanagement.¹²³ What happens when a person doesn't violate any procedural epistemic obligations to manage their beliefs well, yet they fall into moral ignorance and perform a wrongful act as a result? Orthodox views will declare the person off the hook for their unwitting wrongful act because they are not on the hook for their ignorance.¹²⁴ In certain cases, this may strike you as the wrong result. This can happen when a person investigated the moral dimensions of their practices but ended up with false moral beliefs due to the absence of moral truths in their environment. Consider a mobster that asks whether it is morally wrong to kill people who fail to pay protection money. Everyone the mobster asks, perhaps out of fear, says it is morally permissible.¹²⁵ The mobster might consider such a person

¹²² Why think procedural obligations govern the management of attitudes and actions? As Gideon Rosen (2004: 301) explains:

As you move through the world you are required to take certain steps to inform yourself about matters that might bear on the permissibility of your conduct. You are obliged to keep your eyes on the road while driving, to seek advice before launching a war and to think seriously about the advice you're given; to see to it that dangerous substances are clearly labeled, and so on....these procedural obligations are always 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 direction, etc. The procedural obligation is not itself an obligation to know or believe this or that. It is an obligation to take steps to ensure that when the time comes to act, one will know what one ought to know.

¹²³ A view summarized by the Culpable Ignorance (CI) principle. See FitzPatrick (2008: 609).

¹²⁴ Principles that underwrite this orthodox commitment are, "(i) S is blameworthy for [her unwitting wrongful act] only if S is blameworthy for her ignorance that [the unwitting wrongful act] is wrong; and (ii) S is blameworthy for her ignorance that [the unwitting wrongful act] is wrong iff S is blameworthy for a benighting act...that led to [the unwitting wrongful act]" (Wieland 2017a: 9). Also, it is important to remember that the act's wrongness is an objective fact. As mentioned, this is a stipulation of the literature on the epistemic condition to isolate when responsibility for objectively wrongful acts is instantiated and when it is not. Responsibility and blameworthiness come in tandem.

¹²⁵ Something similar can happen when people are surrounded by so-called "yes men." Moral degradation can persist due to echo chambers that reinforce moral falsehoods. Instead, one often needs

innocent in terms of not deserving death given the person only failed to pay money, but the mobster might believe it is permissible all-things-considered to kill the innocent person because they need to deter others from failing to pay the mob. Killing the innocent person is morally permissible overall because the ultimate moral code guiding the mobster's actions is ultimate loyalty to the mob. Such a person knew that they were killing an innocent person, but they were morally ignorant, and their act reflected values that they endorsed, such as utmost loyalty to and protection of their social clan. Does the fact that the person did not violate any procedural epistemic obligations excuse the person from responsibility and blameworthiness for killing an innocent person? Harman's view can account for the blameworthiness of such a person. Even though the person did not mismanage their opinions, the person is still blameworthy for their unwitting wrongful *action* because that action is attributable to the person in a way that reflects something morally objectionable about the value and significance the person attaches to others, such that it is permissible to knowingly kill innocent people.

At this point, you might wonder whether claiming that moral ignorance never exculpates has a strictly theoretical payoff, i.e., it allows us to not let off the hook, to adequately classify, people whose bad behavior is attributable to them in a manner showing that there's something morally problematic about their stance toward others? But, in addition to this payoff, the truth of the No Excuse thesis has practical, everyday implications.

moral disagreement or friction to discover moral truth. As Scripture says, "As iron sharpens iron, so one person sharpens another" (Proverbs 27:17).

Every culture in history was morally ignorant about some things. It's a form of cultural arrogance to claim that we, as contemporary citizens, escape moral ignorance. Though we might be morally enlightened about matters morally benighted to our ancestors, we are morally blind concerning matters that future generations will see clearly. We might think hard about the ethics of artificial intelligence, animal rights, reparations for slaveholding and extermination of indigenous people groups, the ethics of business practices, environmental ethics, and so on, and due to the absence of reliable evidence and the abundance of misleading evidence, we may fall into moral falsehood and perform wrongful acts as a result.¹²⁶ This suggests a posture of humility. If No Excuse is true, we must concede that we're on the hook morally for more than we realize. This can prompt vigilance in performing a searching inventory of our moral values and how those values reflect what we think about persons, animals, and the environment.¹²⁷ Liability to social sanctions and moral criticism can motivate such inquiry. We want to not do the wrong thing, which can prompt moral vigilance yet often is not enough to do so. But we also want to avoid being subject to the resentment, indignation, and moral anger of others. If we are on the hook for unwittingly performing morally wrong actions, as moral ignorance is no excuse, then we are legitimate targets of such moral sanctions. Such liability to adverse treatment can

¹²⁶ This is because, as Harman (2014) entitled a paper defending the No Excuse thesis, "Ethics is Hard! What follows?" In this chapter, I largely set aside the issue of whether intellectual difficulty is exculpatory. If it's hard to discover the moral truth, does that mitigate blameworthiness? Bradford (2017) and Guerrero (2017) argue that difficulty *per se* is not mitigating.

¹²⁷ Recall that it is possible to discover the moral truth, not impossible. Even if it is difficult, given one's cultural context, to discover the moral truth it is possible to discover it through reflection or moral exemplars that are ahead of their time from a moral perspective.

motivate vigilance and encourage us to not uncritically accept the moral norms that are prevalent in our culture today.

With this background in mind, let's get a statement of Harman's view on the table. Harman's view proposes a symmetry between actions and attitudes:

We are morally obligated to believe the moral truths relevant to our actions (and thus not to believe false moral claims relevant to our actions), and we are often blameworthy for failing to meet these moral obligations, even if we have not been guilty of mismanagement of our beliefs, and even if our ignorance is not motivated. Wrong actions that result from false moral beliefs are not thereby blameless; indeed, they may be loci of *original* responsibility. While both the beliefs and the actions are blameworthy, the actions are not blameworthy *because* the beliefs are blameworthy. Rather, the actions and the beliefs are blameworthy for similar reasons. (Harman 2011: 459)¹²⁸

For Harman there is an obligation to believe the moral facts and an obligation to act in accordance with the moral facts. Failure to possess the moral facts and failure to act in accordance with the moral facts are both sources of non-derivative responsibility. If a person holds a false moral belief, we need not scour the causal history of the belief in search of instances of belief mismanagement from which we can derive responsibility for the false belief. It is enough that the person doesn't believe the moral

¹²⁸ Recall that one is directly or originally responsible for something only if one is not responsible for it in virtue of being responsible for something else. In contrast, indirect or derivative responsibility requires that one is responsible for that thing by way of something else that led to it. It is important to discover a source of original responsibility to halt the regress of blameworthiness regarding wrongful acts done from ignorance. See Wieland (2017a) for more on this regress.

truth and the false belief is attributable to the person because it reflects their values.¹²⁹ Similarly, when a person fails to act in accordance with the moral truth, we need not search for belief mismanagement in the etiology of the action. It is enough that the action is attributable to the agent. This grounds direct blameworthiness. To illustrate this view, consider a twist on the original Mr. Potter case:¹³⁰

Enlightened Mr. Potter. A businessman...has thought seriously about the ethics of business and...concluded that what are in fact 'reprehensibly ruthless' practices are really 'permissibly aggressive.' This businessman is bothered by the fact that some people criticize his practices. He isn't surprised that those who are hurt by the practices say they are ruthless, but it bothers him that his brother says so. He discusses it at length with his brother and also with others. He comes to the view that it is better for the economy overall if companies operate to maximize profits, and that economic progress always has casualties. He believes that by pursuing what is in the best interests of his shareholders, he is playing a valuable role in the economy. (Harman 2011: 456-7)

Is enlightened Mr. Potter blameworthy for his reprehensibly ruthless actions done in ignorance that they are wrong? A view that requires blameworthiness for unwitting wrongful action to trace to blameworthiness for ignorance, and thinks of blameworthiness for ignorance as strictly involving belief mismanagement, motivated ignorance, or akratic acts, will declare Potter excused from moral responsibility and

¹²⁹ As Harman explains about beliefs, "[O]ne can be blameworthy for having false moral beliefs although one has not been in any way irresponsible in the management of one's opinion: one may not have acted in any way procedurally badly. Rather, one has violated some moral norms that apply to beliefs themselves, not to the management of one's beliefs" (Harman 2011: 459).

¹³⁰ Recall that the original Mr. Potter case is modeled after the character by that name in the movie *It's a Wonderful Life*. The case is discussed at length in FitzPatrick (2008).

blameworthiness for his ruthless business practices.¹³¹ By contrast, Harman thinks Potter is directly responsible for his ignorance *and* is directly responsible for his actions. Potter is not responsible for his actions *because* he is responsible for his ignorance. Rather, Potter's ruthless actions are attributable to him. They reflect his values. They suggest something morally objectionable in his stance toward persons, as his all-things-considered judgment that his actions are morally permissible indicates a morally deficient weighting of economic interests versus the harm caused to people and communities in pursuing those interests.¹³²

Though I have explained *how* Harman's view captures the verdict that enlightened Mr. Potter is directly responsible for his reprehensibly ruthless actions and the false beliefs underlying those actions, I have not yet explained *why* Potter is blameworthy. On this view, violating norms that apply to actions and attitudes results in a failure of caring properly for morality. Believing the moral truth, and acting in accordance with it, is a way to properly care for morally significant facts. Flouting such

¹³¹ Enlightened Mr. Potter was not guilty of belief mismanagement. It bothered him that his brother and others criticized his practices as being ruthless, so he talked to them at length. Despite thinking hard about the permissibility of his practices he ended up with a false belief that his actions were all-things-considered morally permissible. Yet, intuitively, just because Potter discussed his practices with his brother and others this doesn't excuse him from moral responsibility and blameworthiness for his reprehensibly ruthless business practices.

¹³² Thanks to Aaron Zimmerman for asking a question about manipulation. Mr. Potter might have improperly weighted economic interests versus harm were he brainwashed or directly programmed with bad moral belief and values. Would Potter be blameworthy in such a case? One thing to say is that it depends on whether Potter was brainwashed or directly programmed with the beliefs and values such that he had time to take responsibility for the implanted mechanism of action. If so, then Potter is responsible for the action emanating from a mechanism that produces a deficient weighting of economic interests versus harm caused. Potter owns the mechanism. However, if the implantation is directly before the moment of action, and there was not time for the implanted mechanism to be reflectively endorsed or become part of Potter's moral character, then Potter is not responsible. In the cases at issue it is explicit or reasonable to infer that the person reflectively endorses their bad values. Also, the bad values were not implanted by nefarious manipulators in these cases.

norms evidences a lack of moral concern. Harman takes this line by embracing a Quality of Will view associated with Arpaly (2003).¹³³

The Quality of Will approach to moral responsibility is a broadly Strawsonian view of responsibility.¹³⁴ Framing this approach as a principle concerning blameworthiness:

Quality of Will. An agent S is blameworthy for one of her actions (or omissions) X iff S does X with a lack of good will. (Wieland 2017: 152)

When does a person perform an action or commit an omission while lacking a good will? For Arpaly, a person performs an action while lacking a good will when they fail to care *de re* about the aspects of their actions that, in fact, matter morally. *De re* caring about morality is relativized to a person's desires and emotions tracking the features of their actions that are in fact morally significant. This contrasts with a person caring *de dicto* about morality, which concerns, from a first-personal perspective, whether the person desires to act morally. *De re* and *de dicto* caring can come apart, as in the case involving Huck Finn. Regarding whether he should turn in his friend Jim who is a runaway slave, Huck does not care *de dicto* about morality, but he does care *de re* about morality. Huck does not want to act morally because he's caught in the grip of a false view about what morality demands. Huck believes the morally right thing to do is to turn in Jim.¹³⁵ But, Huck does care about the features of his action that in fact

¹³³ Harman (2011) tentatively embraces Arpaly's view. Then Harman doubles down on Arpaly's view in Harman (2014) and affirms that commitment in Harman (2017). A related view to Arpaly's view is proposed by Julia Markovits (2010).

¹³⁴ Witness the claim from Strawson (1962) that, "The reactive attitudes...are essentially reactions to the quality of others' wills towards us, as manifested in their behaviour: to their good or ill will or indifference or lack of concern."

¹³⁵ And, as things progress, Huck feels increasingly guilty about not turning in Jim.

matter morally.¹³⁶ Huck's act of not turning in Jim shows sensitivity to Jim's autonomy and humanity, which would be badly compromised by turning in Jim.¹³⁷ What matters for praiseworthiness is caring *de re* about morality. On Arpaly's view, Huck is praiseworthy for not turning in Jim. His action was performed with a good quality of will. By contrast, a person performs an action while lacking a good will when they fail to care *de re* about the morally significant features of their actions. Harman relates this back to being blameworthy for actions and attitudes that flout the norms to believe the moral truth and to act in accordance with the moral truth:

Behaving in a certain way is blameworthy just in case (and to the degree that) the behavior results from the agent's caring inadequately about what is morally significant—where this is not a matter of *de dicto* caring about morality but *de re* caring about what is in fact morally significant...Beliefs (and failures to believe) are blameworthy if they involve inadequately caring about what is morally significant. Believing a certain kind of behavior is wrong on the basis of a certain consideration is a way of caring about that consideration. (Harman 2014: 13)

Embracing this Arpalyian view, Harman thinks that actions and attitudes are blameworthy for a similar reason, namely because they stem from a deficiency of *de re* caring for what in fact matters. Such actions and attitudes stem from the absence

¹³⁶ It is important to note that *de re* caring should also require caring "in the right way" for what in fact matters morally. People can care for what in fact matters morally but do so for the wrong reasons. For instance, a person might care about the fact that stabbing an innocent person is wrong because they want to do something wrong. They care about what is in fact morally significant, but they care about it for the wrong reason—for a reason that does not count in favor, from the moral perspective, of doing the action. Thanks to Hieronymi for suggesting this important qualification of *de re* caring.

¹³⁷ Huck *de re* cares for Jim's humanity. Huck's actions demonstrate that he cares for Jim's autonomy as a person—that Jim be treated as a person and not as a mere means, or a mere piece of property, serving the slaveowner's ends.

of a good will (i.e., absence of enough moral concern). To illustrate Harman's Arpalyian view, let's consider two cases from Harman (2014: 10):

Gang Avenger. A gang member kills a member of an opposing gang who killed his friend (and fellow gang member). He believes he is doing the right thing.¹³⁸

Mob Enforcer. A mafia "family" member kills a shop owner who refuses to pay a weekly extortion fee. (The mafia demands a weekly payment for "protection"; otherwise they threaten violence.) He believes he is doing the right thing.¹³⁹

The agents are blameworthy for their actions. They knowingly kill innocent people, yet they do not believe that they're doing anything wrong. Harman connects this to the absence of moral desire and care in the Mob Enforcer case:

The mafia family member knows he is killing an innocent person in order to secure financial gain for his family. He is inadequately moved by the fact that his action kills an innocent person. He acts as he does because he does not have a strong enough desire not to kill innocent persons—he does not care enough about not

¹³⁸ It's important to stipulate that the opposing gang member that gets killed is not actively posing a threat to anyone at the time of his murder. At the time of the revenge hit, the opposing gang member is innocent in that he has not done anything to remove his right not to be killed. A similar stipulation applies to the shop owner. At the time of being killed, the shop owner has not done anything to vacate his right not to be killed, such as attacking the mobster with intent to kill. Both agents are innocent in this regard. The killings are not justified via self-defense.

¹³⁹ Harman indicates that the mafia family member imbibed their moral beliefs. This is a case involving, "someone raised in a Mafia family who goes into the family business and believes in an ethics of deep loyalty to the family business group and no moral obligations to those beyond it; this man kills a store owner who won't 'pay for protection'" (Harman 2011: 457). One might wonder whether the fact that the Mafia family member was raised to hold the beliefs in question is blameworthiness mitigating? However, on the care-based attributionist model, the history of how the mobster came to possess their moral beliefs is not freedom-undermining. At some point in the stage of development the young mobster came to care about morality on the basis of bad moral values. This may have occurred after the child could recognize false beliefs and engage in counterfactual reasoning (e.g., if only I hadn't hit my sister, I wouldn't have hurt her and got in trouble"). At that point, likely early in childhood development, the child became a moral agent. The mobster was able to care about morality. The fact that he imbibed the false moral beliefs from his mafia upbringing does not absolve him of responsibility for actions done from such beliefs.

killing innocent persons. This is why he is blameworthy for the killing. But on my view, the mafia family member is also blameworthy for his moral beliefs. His belief that it is morally right to kill the shop owner is a way of having a morally objectionable attitude to the shop owner—holding it is to hold that the shop owner’s life is cheap and can permissibly be sacrificed to his family’s own goals. This attitude itself is blameworthy, on my view. (2014: 14)

One might object: How have such agents thought carefully about morality and not violated any procedural obligations? In response to such an objection, Harman elaborates on the agents’ perspectives:

The gang member understands that he himself might be killed one day by a rival gang member; he does not think that the gang member would be acting morally wrongly. The gang member and the mafia family member, unlike the ancient slaveholder, know that there are people who think their behavior is morally wrong. But they think that others have been suckered into a false touchy-feely moral view of loving everyone; and that others do not adequately appreciate each person’s moral duties to take care of her own. (2014: 11)

The gang avenger and mob enforcer are not inconsistently applying moral standards. They apply moral standards consistently to themselves and others. They think others fail to appreciate what morality demands—utmost fidelity to those directly under one’s care and protection. They do not live in a benighted culture (i.e., a desert regarding moral truth), and they have considered the counter-opinions of others and think that such people embrace a false moral view. They do not exhibit intellectual vice or

motivated ignorance.¹⁴⁰ Yet, their acts and attitudes express a failure of concern for what in fact matters morally. They lack a good quality of will because they are inadequately sensitive to the fact that their actions result in the killing of innocent lives.

3.2 A Challenge to Harman's Defense of No Excuse

Jan Willem Wieland (2017b) introduces a counterexample to Harman's defense of the No Excuse thesis. It is a counterexample to Harman's use of Arpaly's account of blameworthiness to try to support the No Excuse thesis. The case involves knowingly enslaving someone without the presence of ill will or the absence of good will. Harman's account renders the slaveholder morally blameless and their moral ignorance exculpatory, which counters Harman's view that moral ignorance never exculpates. Here's the case:

Sympathetic Slaveholder. 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

¹⁴⁰ For more on motivated ignorance see Moody-Adams (1994) and Wieland (2017c).

the limited social context, and concluded that it was permissible. (Wieland 2017b: 156-7)¹⁴¹

Wieland provides an illuminating analysis of why Harman's view fails regarding this case. Wieland agrees with Harman's positive account that, "to believe that keeping slaves is wrong because of their suffering (and so not being ignorant) is a way of caring about their suffering," but Wieland doesn't think it follows from this that, "failures to believe the truth always involve a lack of good will (or an inadequate concern for what is morally significant)" (Wieland 2017b: 158). In the Sympathetic Slaveholder case, Cleo cares about the slaves' suffering, but Cleo doesn't believe that enslaving people is wrong because of their suffering. Thus, it is false that, "if you care about Y, then you believe that X is wrong because of Y."¹⁴² Yet the truth of this necessity claim is needed for moral ignorance to never exculpate, given Harman's account of blameworthiness. Otherwise, as in this case, it is possible for *de re* sensitivity to morally significant facts to float free from believing that something is morally wrong because of those facts. Moral ignorance can exculpate, which is contrary to Harman's view. So, the Sympathetic Slaveholder case shows that Harman's own theory of blameworthiness doesn't support her defense of No Excuse.¹⁴³

¹⁴¹ Wieland labels this Case 3**.

¹⁴² Wieland (2017b: 158) calls this the Necessity claim. Implicit in this principle is the connection between X and Y, namely that X is wrong because of Y.

¹⁴³ A proponent of the Harman-Arpaly line might recall that I added the qualification to *de re* caring that it involves caring "in the right way" for what is in fact morally significant. Perhaps Cleo does not care in the right way for the well-being of her slaves. This might stem from the fact that she continues to keep them. I thank Hieronymi for raising this potential worry. In response, I think it matters why Cleo continues to keep slaves. She cares about the fact that they suffer. She cares about what in fact matters morally, and she does not care about what in fact matters morally because she wants to inflict suffering on her slaves or do something morally wrong. She does not care about what in fact matters morally for the wrong reasons (from a moral perspective). Instead, her continuing to keep slaves is the result of a failure of inference. She fails to infer the wrongness of keeping slaves from the fact that they

I affirm Wieland's analysis. A person can *de re* care about the morally relevant facts, i.e., the suffering of slaves, and care about morality *enough*, without believing that keeping slaves is morally wrong because of those facts. Failing to believe the moral truth does not always implicate a poor quality of will in the narrow sense endorsed by Harman.¹⁴⁴ Equating the lack of a good quality of will with not caring enough for morally significant facts is overly-restrictive. Instead, I will broaden the notion of attributability that grounds blameworthiness.¹⁴⁵

3.3 New Attributionism

The Arpalyian view of moral responsibility is a version of an attributability view concerning the nature of moral responsibility. Attitudes and actions are attributable to agents as grounds for responsibility and blameworthiness on the basis of affect-

suffer, and this is due to her limited social context. Thus, the mere fact that Cleo keeps slaves does not show that she fails to care about them in the right way, for the right reasons.

¹⁴⁴ Even though Cleo cares adequately, she still has a false moral belief.

¹⁴⁵ It is worth mentioning here the case of Ishmael Beah. He was a child soldier in Sierra Leone during the civil war in the country during the 1990s. Beah was brainwashed and given drugs by those who forced him to fight. One might object that it is cruel to punish him for what he did as a child soldier, and the No Excuse thesis is problematic on these grounds. I thank Zimmerman for raising this concern. As an initial response, I should mention that the Beah case is a somewhat straightforward case of coercion. Unlike the Potters, slaveholders, etc. Beah was subject to freedom undermining manipulation. He lacks the requisite control over his actions. So, the question concerning the epistemic condition is not as readily applicable to this case. In the mobster case, Potter cases, etc. the agents are not subject to such coercion. They freely commit their wrongful acts. Assuming the control condition is satisfied in these cases allows us to focus on the epistemic condition. Secondly, assuming the control condition is satisfied in the Beah case, someone like Beah is blameworthy for his childhood crimes on my account, but whether he should be jailed and punished for the murders focuses on criminal responsibility and whether actual applications of blame are justified. Beah is blameworthy in the desert-based sense of being responsible, but actually blaming him for his childhood crimes may be unjustified for reasons that go beyond backward-looking desert or are forward-looking from the time of action. Beah repented of his wrongs, turned away from the life of crime, and now helps kids coming out of being childhood soldiers. Given these factors it would be cruel to punish him for what he did as a child. Yet, as Harman and I contend, he is still blameworthy for those crimes. Though, I disagree with Harman that they do not always evidence a lack of care for morality. Beah may have cared about morality but been unable to act on that care given that running away from the unit would have meant death. I think, though, that his actions still reflected an objectionable quality of judgment concerning what it is permissible to do to people. On my view, he deserves blame for murdering all those people, but it may not be appropriate to actually blame him given his repentance and reparations for his wrongs.

oriented components, such as emotions, desires, and cares. A more encompassing conception of attributability allows more rationalistic components, such as normative reasons and evaluative judgments, to ground blameworthiness as well. To overcome the objection that caused trouble for Harman's defense of the No Excuse thesis I propose being ecumenical about attributability. I will focus on the main idea that unites attributability views while at the same time honoring the differences between the views. I will embrace a view that allows affective considerations to count as sufficient for blameworthiness and rationalistic considerations to count as sufficient as well.

New Attributionism consists in a cluster of views about the nature of moral responsibility and blameworthiness.¹⁴⁶ The views are united by their embrace of a central idea, namely that it is, "sufficient for blameworthiness that a person's bad behavior is attributable to her in a way that indicates something morally objectionable about her and her orientation toward other people."¹⁴⁷ To deepen our understanding of New Attributionism I will contrast it with an influential conception of attributionism, and I will show how different attributionist views flesh out the main idea behind attributionism.

New Attributionism contrasts with the brand of attributionism discussed by Gary Watson (1996, 2004). Watson thinks of responsibility as Janus-faced. One face is the attributability face, and one face is the accountability face. Attributability concerns appraisals of a person's moral character. When a person is responsible in the

¹⁴⁶ New Attributionists include Hieronymi (2008, 2014), Scanlon (1998, 2008), Sher (2006, 2009), Smith (2005, 2015), and Talbert (Harman 2017; 2013).

¹⁴⁷ Talbert (2016: 52). The term 'New Attributionism' comes from Talbert (2016).

attributability sense for her behavior that behavior reflects her real self or her moral identity. When a person is responsible for her behavior in the accountability sense it would be appropriate to hold her accountable for her behavior; she qualifies as an appropriate target of reactive attitudes that impose burdens on her.

The two faces of responsibility can come apart. When a person had a childhood that consisted in moral deprivation, physical and mental abuse, and so on, that person's behavior may reflect who she really is, but it would not be appropriate to hold her accountable for her behavior given that it reflects the natural outpouring of her poor formative circumstances, which she did not choose or control. According to Watson (2004: 281), "Facts about [her] formative years give rise to the thought that the individual has already suffered too much and that we too would probably have been morally ruined by such a childhood." In such a case, the person is responsible for her actions in the attributability sense but not in the accountability sense.

New Attributionism combines Watson's two faces of responsibility. It accounts for the way agent's behavior reflects moral faults and who the person really is, and it accounts for the suitability of responding to the agent in ways characteristic of holding the agent accountable for her actions. New Attributionism captures the responses licensed by accountability blame, but, as Matthew Talbert (2016: 53) explains, "contemporary attributionists believe that wrongdoers can be open to these responses simply because bad behavior is attributable to them...attributability is enough for accountability." This means that New Attributionism is not focused on considerations involving etiology and formative circumstances. The focus in attributionism is whether the person's present behavior is attributable to them. It is not concerned with

whether the person could have done something different given their history. With this background in mind, I turn to further clarifying the commitments associated with New Attributionism by looking at specific views.

A way that theorists cash out the main idea behind New Attributionism is in terms of answerability.¹⁴⁸ These answerability theorists embrace the basic idea that one is answerable for such attitudes and actions when they bear a rational connection to one's evaluative judgments or the normative reasons one took to justify one's attitudes and actions. When this obtains it is right to request of the agent that she answer for her actions and attitudes—it is appropriate to request of the agent her reasons for action and belief. Such reasons reveal what the agent values and thinks of as true and good. To better understand answerability as a version of New Attributionism, we will look at two iterations of the view.

Pamela Hieronymi characterizes answerability along Anscombean lines. Answerability is a matter of a why-question being rightly asked of an agent, “Why did you ϕ ?”¹⁴⁹ This request for an agent's reasons is appropriate given that the agent

¹⁴⁸ Such theorists include Pamela Hieronymi (2008, 2014) and Angela Smith (2005, 2012, 2015). There is a live debate in the literature on the nature of moral responsibility concerning pluralism or monism about responsibility. A pluralist like David Shoemaker (2011) conceives of attributability, answerability, and accountability as separate senses of responsibility, which each license different moral reactions. He might regard New Attributionism as blurring important distinctions between separable accounts of responsibility. On the other hand, a responsibility monist like Angela Smith (2012) might not welcome her view being grouped with other views competing for the one true view of the nature of moral responsibility. Here, though, I cannot wade into the debate over the genuine nature of moral responsibility.

¹⁴⁹ G. E. M. Anscombe (1957) regards intentional action as being a matter of the applicability of a why-question within the context of action. A person is intentionally performing an action under a certain description if the person would not reject a question about why they are performing that action. Hieronymi embraces the applicability of a why-question regarding responsibility. But, Hieronymi does not think that the fact that you did not know you were ϕ -ing refuses the question: why are you ϕ -ing? As Hieronymi explains, what refuses the why-question is the fact that you have not settled the whether-question (whether to ϕ). If you have not settled the whether-question, then you will not have practical knowledge of what you are doing (indeed, you will not be doing anything intentionally). But Hieronymi

settled the question of whether to ϕ by taking certain reasons to count in favor of ϕ -ing. To see how this works, consider a verbal fight between people:

Suppose...that you intentionally end the fight. We know, then, that you settled for yourself (positively) the question of whether to end it. If we know a little about the context of the fight, and a little bit about your particular epistemic situation, knowing that you decided to end the fight tells us something of how you think about the world and your place in it. We will react in ways that reveal that we find your decision reasonable or unreasonable, justified or unjustified. If we further think you decided to end it for certain more-or-less elaborated reasons, we may form certain further, more-or-less elaborate opinions about you: we might think you have been disloyal, spineless, magnanimous, mature, or conniving. Such assessments are typically thought to license certain corresponding sorts of reactions: resentment, contempt, regard, admiration, or distrust. (Hieronymi 2014: 15-6)

Various reactive attitudes are fairly targeted at you based on our evaluation of the reasons for which you settled the question of whether to ϕ . Such reasons reveal, as Hieronymi (2014: 15) says, “a certain stretch of your mind.”

Answerability not only applies to actions, it also applies to attitudes. Angela Smith (2005) defends an answerability view she calls the “rational relations view.” Smith’s view contrasts with a volitional view of responsibility according to which we

does think (pace Anscombe) that you might have settled the whether-question and yet be self-deceived about what you are doing intentionally (and so sincerely reply “I did not know I was ϕ -ing”). Thanks to Hieronymi for clarifying her approach in relation to Anscombe.

are only responsible for actions and attitudes traceable to voluntary choices.¹⁵⁰ The rational relations view accounts for attitudes that are not the result of conscious decisions. It accounts for spontaneous emotions, what things occur to us and what we forget, and what things we notice and neglect to notice. Such patterns of mental states are not within our direct voluntary control.¹⁵¹ They reveal our evaluative commitments. As Smith argues:

When we praise or criticize someone for an attitude it seems we are responding to something about the content of that attitude and not to facts about its origin in a person's prior voluntary choices, or to facts about its susceptibility to influence through a person's future voluntary choices....it seems we are responding to certain judgments of the person which we take to be implicit in that attitude, judgments for which we consider her to be directly morally answerable. If this is correct, then it is a mistake to try to account for a person's responsibility for her own attitudes in terms of their connection to her prior or future voluntary choices, because that obscures the special nature of our relation to our own attitudes: we are not merely producers of our attitudes, or even guardians over them; we are, first and foremost, inhabitants of them. They are a direct reflection of what we judge to be of value, importance, or significance...it is in virtue of their rational

¹⁵⁰ The absence of focus on direct voluntary control or choice is a hallmark of New Attributionism. For more on this nonvolitional aspect of New Attributionism, see Robert Merrihew Adams (1985) and George Sher (2006, 2009).

¹⁵¹ Hieronymi (2008) makes an important point in this regard. It is precisely because such attitudes constitute our commitments that they are outside our direct voluntary control (i.e., not under managerial or manipulative control). Our attitudes represent our take on the matter, and, as such, "they are not the sort of thing one can form or revise or maintain for any reason one takes to show it worth doing" (Hieronymi 2008: 371). This is what makes such "commitment-constituted attitudes" proper objects of moral evaluation, as they reflect our take on the world (i.e., what we value and think of as true and good).

connection to our evaluative judgments that they are the kinds of states for which reasons or justifications can appropriately be requested. (2005: 251)

For Smith, we can be responsible for things that we do not directly voluntary control or choose, such as our attitudes and the things we notice or fail to notice. If such attitudes and mental states are reflective of, “what we judge to be of value, importance, or significance” regarding persons and other things, then they are things we can be responsible for. Such things reflect our judgments about what is of value, and, as Talbert (2016: 56) explains, “insofar as our attitudes reflect objectionable judgments, they may be the source of the sort of moral offense that reasonably elicits the responses involved in blame.”

The last New Attributionist view we will look at is a paradigm attributionist view. T.M. Scanlon (1998, 2008) regards responsibility as focusing on whether an action is attributable to an agent such that it can license moral evaluation. Moral appraisal of an agent indicating that they are blameworthy involves situations such that the agent has, “governed him- or herself in a manner that cannot be justified in the way morality requires” (Scanlon 1998: 272). Proper evaluation of an agent targets things that allow for self-governance, such as judgment-sensitive states concerning one’s reasons for belief and action. Failures of self-governance for states that are judgment-sensitive open a person to blameworthiness and blame. As Scanlon indicates:

to claim that a person is *blameworthy* for an action is to claim that the action shows something about the agent’s attitudes toward others that impairs the relations that others can have with him or her. To *blame* a person is to judge him

or her to be blameworthy and to take your relationship with him or her to be modified in a way that this judgment of impaired relations holds to be appropriate.

(Scanlon 2008: 128-9)

For Scanlon, failures of self-governance occur when an agent has governed him- or herself, “in a way that shows a lack of concern with the justifiability of his or her actions, or an indifference to considerations that justifiable standards of conduct require one to attend to” (2008: 141). States such as “fear, anger, admiration, [and] respect” are responsive to and reflective of judgments concerning reasons for action. As judgment-sensitive states, when objectionable, they can serve as a proper basis for blame in the attributability sense. But not all failures of self-governance are moral failures. Think of a poor move in a game of soccer. This may open the agent to a form of criticism, as the bad move may reflect a judgment about what she had reason to do, but such criticism is not moral criticism. Instead, failures of *moral* self-governance involve morally significant impairment of interpersonal relations. For instance, if a person acts in a way that shows indifference to the justifiable standard to do no unnecessary harm to innocent people this will damage relations that that person can have with those harmed by his actions. The innocent victim will rightly resent and not trust the person that harmed them. Though, Scanlon’s view is subject to criticism.

Scanlon’s emphasis on the impairment of relationships is not without its detractors.¹⁵² R. Jay Wallace, in Wallace, Kumar, and Freeman (2011), argues that Scanlon’s view, “leaves the blame out of blame.”¹⁵³ Wolf also thinks that blame can be

¹⁵² I thank Zimmerman for pressing me consider carefully this feature of Scanlon’s view.

¹⁵³ In this volume on Scanlon’s work, Wolf also objects to this aspect of Scanlon’s account.

deserved in the absence of damaged relations. Think of a family that ultimately loves each other, and their relationships are not damaged, but they are fair targets of blame by family members on certain occasions. Or, think about a victim of rape by a stranger. There is no relationship to be damaged between the perpetrator and victim, but the perpetrator is blameworthy for violating the victim's rights and causing harm. Scanlon responds to this sort of criticism by claiming that there is baseline relation that all rational agents have. He calls this the "moral relationship." For my purposes, most of the cases in the literature do involve relationships that are damaged or must be modified because of the blameworthy actions done from ignorance. Yet, if Scanlon's emphasis on impaired relations is problematic, I think it is fine to jettison this feature in a broadly New Attributionist perspective. The key, then, is that stating that a person is blameworthy shows something about the agent's attitudes toward others, such that it is permissible to treat persons in a morally objectionable way. Whether that attitude reflects a stance that damages relations that others can have with the person is secondary. It may show up in some cases but not others. In the case involving a rape committed by a stranger the perpetrator's action reflects a morally objectionable attitude toward others, such as thinking it permissible to violate a person's autonomy and rights in the name of sexual pleasure.

The key for a broadly attributionist perspective is the focus on evaluative judgments and judgment-sensitive states reflected in actions. Damaged relationships need not serve as the foundation of blameworthiness. An expansive brand of New Attributionism accounts for quality of will and quality of judgment, as embodied in this principle:

Quality of Will and Judgment. An agent S is blameworthy for one of her actions (or omissions) X iff S does X with a lack of good will or S does X with a lack of good judgment.¹⁵⁴

3.4 New Attributionism Applied

Now I will apply New Attributionism and the Quality of Will and Judgement principle by focusing on the quality of judgment portion of the principle.¹⁵⁵ I will analyze the Sympathetic Slaveholder case that caused a problem for the pure Quality of Will view Harman embraced in defending the No Excuse thesis.

According to New Attributionism, Cleo is blameworthy for keeping slaves because Cleo's keeping of slaves is accompanied by a lack of good judgment. Cleo judged that all-things-considered slavery is permissible. Cleo's attitudes show the all-things-considered judgment objectionable. Cleo possessed knowledge of the slaves suffering and affective mental states as of sympathy for their suffering. That Cleo knew and felt such things yet judged slavery permissible is the source of moral offense that makes reasonable the judgment that Cleo is blameworthy for keeping slaves.¹⁵⁶ But,

¹⁵⁴ Thanks to Greco for prompting me to think about whether the view can factor in both will and judgment. It can as long as the right-side of the biconditional is formulated disjunctively. For an attributability view that focuses on judgment-responsiveness but is not committed to judgments necessarily revealing the moral identity (i.e. values, cares, concerns) of the agent see Strabbing (2016).

¹⁵⁵ This focus reflects the fact that there is an absence of poor quality of will in the case. If No Excuse is supported by attributionism, then the second disjunct (or sufficiency condition) in the principle must capture why the Sympathetic Slaveholder is blameworthy and her moral ignorance is no excuse.

¹⁵⁶ Cleo's all things consider judgment is not reached based on mere self-interest, such as the financial benefits she gets from keeping slaves. If this were the case, this would not be a clear case of attributability. In such a case Cleo would be blameworthy for being partial to her own economic self-interests over the suffering of the slaves. Presumably, such a valuing of narrow self-interest would have been avoidable, which would make the case an accountability case whereby history and avoidability of faulty reasoning are doing the work in generating blameworthiness. Instead, on an attributability reading of the case, Cleo's all-things-considered judgment is reached on basis of slavery promoting ethical goods, such as the it freeing slaveholders to participate in civic life and improve the well-being of citizens. Further, Cleo's reasoning is not epistemically objectionable given the evidence available in her cultural context, as slavery was widely practiced.

isn't her cultural context the reason she failed to conclude slavery impermissible?¹⁵⁷
Doesn't such vacancy of moral truth mitigate blameworthiness?

That Cleo judged slavery permissible all-things-considered due to her limited cultural context doesn't mitigate her blameworthiness. Instead, it suggests the source of her objectionable judgment. Either Cleo's conclusion that slavery is permissible stems from fallacious reasoning or it stems from a morally deficient weighing of reasons.

Cleo's lack of good judgment may have been the result of poor reasoning. Recall that Cleo possessed reasons—in the form of propositional attitudes and affect-based mental states—favoring the immorality of slavery. She knew they suffered and felt sympathy for their suffering. But, these reasons favoring the immorality of slavery did not prove decisive. This might have been because the reasoning used to reach her all-things-considered judgment did not properly account for the good reasons against the permissibility of slavery that she possessed. She might have engaged in fallacious reasoning along these lines, “people have been trying for centuries to provide decisive evidence for the immorality of slavery, but no one has ever succeeded, so it must not be true that slavery is immoral.”¹⁵⁸ Such fallacious reasoning might have been accompanied with commitment to the mores of her culture. Cleo might have thought, “given that most people in my culture think slavery is an acceptable practice, it must be permissible overall, despite my misgivings or reasons to think otherwise.” Some such reasoning and commitments might have generated reasons sufficient to

¹⁵⁷ After all, we are told that Cleo, “did not succeed in drawing the inference that slavery is wrong because of the limited social context” (Wieland 2017b: 157).

¹⁵⁸ Such reasoning is a version of an appeal to ignorance (*Argumentum ad Ignorantiam*).

outweigh the good reasons that she possessed, which favored the truth of the immorality of slavery.¹⁵⁹ However, such reasoning and commitments are fallacious—based on an appeal to ignorance or popularity to justify a moral conclusion.

Cleo's conclusion that slavery is permissible may stem from a morally deficient weighing of reasons. Cleo's limited cultural context may have involved some people questioning the permissibility of slavery, but such people may have attached more value to prudential and economic reasons in favor of the permissibility of slavery.¹⁶⁰ Using ancient Greece as an example, here's an explanation of how testimony against a moral practice can fail to generate a decisive reason to think that practice immoral:

[C]ertain people in ancient Greece did question the practice [of slavery], and in that respect they could provide others with reasons to consider the permissibility of slavery themselves...that there were such debates does not imply that many people in fact concluded that the practice was wrong. Moreover, even if some thought slavery is wrong from a moral perspective, they still thought it is all-things-considered permissible since the greatness of Athens depended on it. Without slavery they might not have been able to realize other values such as democracy. For in that case Athens' citizens would not have been free to engage in a political life. (Wieland 2017b: 160)

¹⁵⁹ We also need to exercise charity regarding the target of the case. Recall that Wieland is considering a problem for Harman's view, and Harman's view is a response to FitzPatrick's view, which grounds blameworthiness for ignorance in belief mismanagement. If Cleo is guilty of fallacious reasoning (e.g., an appeal to ignorance) stemming from epistemic vices, this is a case that FitzPatrick's view handles. Harman grants that there are such cases. She thinks this is not the whole story, though. Often enough, epistemically responsible people to try to discover the moral truth yet fall into error. And Harman considers such people blameworthy.

¹⁶⁰ For subtle analysis of moral progress and the abolition of slavery see Elizabeth Anderson (2014).

This suggests a plausible reading of the case involving Cleo. Given Cleo's cultural context, Cleo was committed to certain values (e.g., commerce, democracy, etc.), and those values trumped, or failed to make decisive, the other reasons Cleo had to think slavery immoral. The weight Cleo attached to the moral reasons she possessed was insufficient and failed to yield an all-things-considered judgment that slavery is impermissible. Such insufficient weighing of her moral reasons reflects what Cleo judges most significant and of value, namely economic, political, or prudential concerns. The weighing of reasons that generated Cleo's all-things-considered judgment reflects her evaluative judgments. This established a connection between her actions of keeping slaves and her judgment that slavery is permissible.

Additionally, it is important to see that Cleo's incorrect weighing of reasons amounts to failing to give her slaves the proper respect or due regard they are owed as persons. Not all incorrect weighing of reasons will amount to disrespecting or disregarding her slaves. For instance, she might have a difficult choice to make between two options that both showed her slaves due regard. But she might get the weighing of reasons wrong and perform the action that is less important in terms of her slaves' well-being. For instance, she may take it that she has a duty to provide drinking water to her slaves and be torn between pumping water to get them water and working on building a new well. She may judge that pumping water is what she ought to do, but in fact working on the new well was of more importance because the old well soon became tainted and threatened to leave the slaves without drinkable water.¹⁶¹ Yet, in the Ancient Greece case involving the weighing of the interests of her

¹⁶¹ This example is inspired by an example given to me by Hieronymi in comments. Hieronymi also brought to my attention the importance between making the connection between error and disregard.

slaves against the interests of the participation in democracy of the elite, this error in weighing amounts to failing to show her slaves due regard. The error in weighing reasons amounts to disregarding the moral interests of her slaves.

Given the previous considerations, Cleo's actions are attributable to her and she is answerable for them. It is reasonable to request of Cleo the reasons she took to justify keeping slaves. Her judgment of the permissibility of slavery evidenced a lack of due regard for the humanity of the slaves, and it impaired relations with them.¹⁶² Whether Cleo's conclusion that slavery is permissible stemmed from fallacious reasoning or from a morally deficient weighing of reasons, Cleo's all-things-considered judgment that slavery is permissible is a flawed moral judgment. Cleo's keeping of slaves is accompanied by a lack of good judgment. As a result, according to the New Attributionist, Cleo's limited cultural context doesn't mitigate her blameworthiness. Her ignorance that slavery is wrong doesn't excuse her from responsibility and blameworthiness for keeping slaves.¹⁶³

Thus, the shift to New Attributionism, as grounding blameworthiness in quality of will and judgment, overcame a counterexample to Harman's defense of No Excuse.

¹⁶² The impairment of this relationship may have licensed resentment and indignation from the slaves toward their owner, especially if aware that Cleo knew that they suffered and sympathized with their suffering yet persisted in forcing them to work without pay. Though this point most directly concerns standing to blame, such as whether and on what grounds slaves could justly blame their slaveholders. It is also worth mentioning that if slaves were taken in war, then the relationships were already impaired regarding their slaveholders prior to any wrongful actions on slaveholders' part. We can stipulate that Cleo's case does not involve owning slaves as the result of the spoils of war. Thanks to Zimmerman for requesting clarification on this point.

¹⁶³ For the attributionist, similar considerations explain why enlightened Mr. Potter, the gang avenger, and the mob enforcer are not off the hook for moral responsibility and blameworthiness. Though such agents are ignorant that their actions are wrong, their judgments that their actions are permissible stem from fallacious reasoning or morally deficient weighing of reasons. Their actions are accompanied by a poor quality of judgment, as they knowingly inflict harm on individuals and communities, including killing innocent people.

My alternative view returns the result that moral ignorance is no excuse in the Sympathetic Slaveholder case, which supports the thesis that moral ignorance never excuses.

3.5 An Objection to New Attributionism

In response to what I argued in section 3.4, one might object that something is missing from the New Attributionist view.¹⁶⁴ New Attributionism needs to include the requirement that the potential target of the reactive attitudes cannot provide a reasonable response that exculpates; the agent cannot answer the challenge by way of excuse. Regarding the Sympathetic Slaveholder case, even if it is appropriate to challenge Cleo for the reasons she took to justify keeping slaves, and what she did was in fact wrong, it is still possible that Cleo could answer the challenge in a way that exculpates.¹⁶⁵ What sort of reason might Cleo give that suggests that she is not a fair target of the negative reactive attitudes?

Cleo might provide an excuse to the challenge for her reasons justifying her all-things-considered judgment that slavery is permissible by citing accountability-type considerations. She might claim that it would be unfair to impose burdens on her in the form of adverse moral responses. Even though her bad behavior is attributable to her, she did not possess the ability to conform to the norms endorsed by myself and Harman. She could not comply with the obligation to believe the moral truth and act in accordance with that truth.¹⁶⁶ Given her upbringing and lack of opportunities to

¹⁶⁴ I thank Hanser for raising a worry along these lines.

¹⁶⁵ One response Cleo could give is foreclosed. Cleo possessed knowledge of the non-moral facts, so blameless factual ignorance is not an answer that could get Cleo off the hook.

¹⁶⁶ Specifically, these norms are as follows:

Belief Norm. We are morally obligated to believe the moral truths relevant to our actions (and thus not to believe false moral claims relevant to our actions), and we are often blameworthy for failing

question the morality of the practice of slavery, she is excused from blameworthiness for keeping slaves, even though that behavior is attributable to her in that it reflects her evaluative judgments.

The challenge just enumerated is problematic.¹⁶⁷ New Attributionism is committed to the idea that attributability is sufficient for accountability. The worry above applies pressure to this commitment. New Attributionism does not focus on whether an agent could have acted or believed differently than they did. It does not focus on an agent's history, upbringing, or opportunities. It focuses on whether the actions and attitudes are attributable to the agent in a way that reflect what the agent values and thinks true and good.

To respond to this challenge, we need to clarify what the blaming responses embodied by the reactive attitudes amount to. According to attributionists, what matters for the fairness of the reactive attitudes is their evaluative significance concerning due regard for the humanity of persons. If Cleo could marshal a response to our challenge for her reasons for slaveholding in a way that shows that her slaveholding was not accompanied by the absence of a good quality of judgment, Cleo

to meet these moral obligations, even if we have not been guilty of mismanagement of our beliefs, and even if our ignorance is not motivated (Harman 2011: 459);

Action Norm. We are morally obligated to act in accordance with the moral truth and we are often blameworthy for failing to meet this obligation, even if we violate this obligation due to beliefs not the result of belief mismanagement or motivated ignorance.

The Action Norm is in keeping with Harman's claim that, "Wrong actions that result from false moral beliefs are not thereby blameless; indeed, they may be loci of *original* responsibility. While both the beliefs and the actions are blameworthy, the actions are not blameworthy because the beliefs are blameworthy. Rather, the actions and the beliefs are blameworthy for similar reasons" (Harman 2011: 459).

¹⁶⁷ This challenge echoes the challenge to attributability (or Real Self) views put forward by Susan Wolf (1990). Wolf argued that attributability views are too thin. They cannot capture the robust conditions on blame captured by reasonable moral demands and negative responses that are deserved for failure to comply with such demands. In chapter 4 I discuss Wolf's argument against attributionism.

would be an unfair target of the reactive attitudes. However, Cleo citing her difficulty in avoiding moral ignorance and wrongdoing due to her upbringing, cultural context, and so on, does not alter the content of the judgments that the reactive attitudes are sensitive to. Cleo's judgment that slavery is morally permissible reflects objectionable values, such as a commitment to severely restrict persons' autonomy and treat them as if they are a piece of property or a possession. Such bad values, and the social context that nurtures and reinforces them, make it hard for Cleo to discover the moral truth, or to derive the moral truth from what she knows and feels about slavery, but this difficulty does not excuse Cleo.

Thus, my response to the challenge that something is missing from the New Attributionist view is to add to New Attributionism what is missing but argue that it does nothing to change our evaluation of Cleo's blameworthiness. The challenge does not detract from New Attributionism as grounding blameworthiness and supporting a defense of the No Excuse thesis.

A person might reply to this that Cleo, and other morally ignorant agents, possess non-moral knowledge and affective mental states. Couldn't Cleo marshal a response to our challenge for her reasons for slaveholding in a way that exculpates? Couldn't Cleo answer the challenge by stating that her knowledge and sympathy concerning the suffering of the slaves mitigates blameworthiness; her epistemic states excuse her? Next, I will argue that this response from Cleo not only fails to exculpate, it makes things worse.

The question of relevance for the attributionist is: Does the interpersonal significance of Cleo's slaveholding change via Cleo's answer to our demands for her

reasons when she cites her knowledge that the slaves suffer or that she sympathizes with their suffering? At first glance, it may look like this sympathetic slaveholder is not blameworthy for slaveholding. After all, she's not completely oblivious to morally-relevant considerations regarding her actions. However, when looking at the standing of Cleo's victims (i.e., her slaves) they would rightly regard a response from knowledge and sympathy as blameworthiness exacerbating. When the slaves learn that their slaveholder knew that they suffered while under her care and that she sympathized with their suffering, a natural response would be an increase in resentment and indignation. They might rightly question, "given what you knew and experienced concerning our suffering, why did you keep us enslaved and continue to subject us to harsh treatment, poor work conditions, and no compensation for our labor?" Given the epistemic states Cleo possessed, her response to such a question will fail to diminish blameworthiness. If Cleo says, "most everyone in my culture was doing it," or, "the practice was good for economic and political reasons," this will fail to assuage the reasonable resentment and indignation her slaves would have toward her given her epistemic states of knowledge of their suffering and her experience as of their suffering. Thus, citing the fact that the practice and the values that underlie slavery are widely endorsed in her social context would not render unfair her slaves' resentment and indignation.

You might wonder whether there is anything special about an agent's epistemic states such that they block an exculpatory answer to the challenge for reasons? Does the request for the reasons Cleo took to settle the question of whether to keep slaves turn on the nature and content of Cleo's epistemic states that accompanied her keeping of slaves? I will argue that this is the case.

3.6 When Knowledge Grounds Blameworthiness

When an agent possesses knowledge of non-moral facts relevant to the act in question, such an epistemic state can ground blameworthiness. This is because of the normative significance of knowledge. As such, whether an agent is answerable for her action, such that her action is attributable to her, a response to that request for the reasons the agent took to justify her action that invokes moral ignorance will fail to exculpate. For instance, the false belief that her action is morally permissible will fail to excuse her from blameworthiness provided that the agent knew the non-moral (descriptive) facts relevant to the action. When moral ignorance is pure in that it is not also accompanied by factual ignorance, moral ignorance is not an excuse.¹⁶⁸

In cases testing whether moral ignorance ever exculpates, the agents possess knowledge of the non-moral facts. The moral ignorance is pure in that it doesn't stem from factual ignorance.¹⁶⁹ Were the moral ignorance to stem from factual ignorance this would fail to test whether moral ignorance ever exculpates. The focus would shift to factual ignorance and whether it is blameless. If it is blameless, it exculpates. If not,

¹⁶⁸ When moral ignorance is not pure in that it is accompanied by factual ignorance, then the priority is determining whether the factual ignorance is culpable or non-culpable. If it is non-culpable, then the agent is off the hook morally. In such a case, the moral ignorance does not do the excusing. This is why moral ignorance never excuses. When moral ignorance is pure, the subject possesses epistemic states, such as knowledge, that ground blameworthiness. So, moral ignorance is no excuse. And when moral ignorance is impure, the factual ignorance can excuse. Again, moral ignorance does not excuse.

¹⁶⁹ Wieland (2017b: 150) discusses this distinction. Moral ignorance is *pure* when it doesn't stem from ignorance of non-moral facts. *Impure* moral ignorance involves moral ignorance due to ignorance of non-moral facts. If someone doesn't know that slaves suffer because they falsely believe that slaves do not suffer, and this explains why they do not know that slavery is wrong, we need to ask whether their ignorance of the non-moral facts is culpable. If that person falsely believes that slaves don't suffer, and this is through no fault of their own (i.e., they tried to investigate the truth of the matter but were shielded by others from the slaves' suffering), the person's factual ignorance is blameless and excuses them from responsibility for wrongdoing. Proponents of the No Excuse thesis grant this possibility. They agree that blameless ignorance of non-moral facts excuses. They just don't think the same thing holds for moral ignorance that's not the result of a violation of procedural epistemic obligations. To test whether this is true it's important to stipulate that in cases involving moral ignorance the person is not ignorant of relevant non-moral facts. This zooms-in on whether moral ignorance *per se* ever exculpates.

it inculpates. Those who endorse the No Excuse thesis agree that blameless *factual* ignorance can exculpate. This would not show, though, that *moral* ignorance exculpates.¹⁷⁰

Knowledge is a normative epistemic state, and it arguably has more value as compared to other epistemic states.¹⁷¹ As John Greco (2010: 4) explains, “When we say that someone knows something we are making a value judgment. We imply, for example, that his or her judgment is preferable to someone else’s mere opinion...knowledge attributions and the like have a normative or evaluative dimension.” Additionally, knowledge is modally robust in a way that other epistemic states are not. Knowledge is often thought to satisfy modal conditions like safety and sensitivity.¹⁷² That knowledge requires satisfying modal conditions makes knowledge stable concerning truth in a way that other states, which need not satisfy such conditions, are not. Knowledge is also often thought to require doxastic justification.¹⁷³ When a person has a doxastically justified belief, they believe the

¹⁷⁰ See the previous footnote for more on this point.

¹⁷¹ A reason for thinking this is that epistemology is normative, and knowledge is a primary target of evaluation in epistemology. As Mark Schroeder (2015: 379) explains, “Epistemologists are concerned not simply with what people believe, how confident they are, or the conditions under which they believe it or are so confident. They are also—and more interestingly—concerned with what people should believe, how confident they should be, and the conditions under which they should believe it and be so confident. But epistemology is, paradigmatically, the study of knowledge.” Epistemology is normative and paradigmatically focuses on knowledge, so the normativity of knowledge is relevant in thinking about one’s total epistemic state.

¹⁷² Safety requires that if a person were to believe a proposition, then that proposition would not be false (where this is determined by reference to nearby possible worlds). The upshot is that the person could not have believed the proposition and easily fallen into error (i.e., believed a false proposition). Sensitivity is the contrapositive of safety: if the proposition were false, the person would not believe it. Defenders of safety include Duncan Pritchard (2007), Ernest Sosa (1999), and Timothy Williamson (2000). Robert Nozick (1981) is a prominent defender of the sensitivity condition on knowledge.

¹⁷³ I should clarify that I am not committed to knowledge always requiring justification and the satisfaction of an epistemic basing relation. As Greco rightly points out, holding these two theses together is highly controversial and leads to well-known skeptical problems and makes it hard to account for perceptual knowledge. I do not intend to take a controversial stand on the nature of knowledge. I only intend to point out its normative dimension and note that some ways of characterizing knowledge better capture this normative dimension. I am implicitly taking a stand that

proposition and base their belief that the proposition is true on the good reasons they have that support the proposition. That is, doxastically justified beliefs are properly based on good reasons.¹⁷⁴ Modal robustness and proper basing contribute to the idea that knowledge is more valuable than states falling short of knowledge.

Another way to capture the evaluative significance of knowledge is to embrace a virtue-theoretic conception of knowledge. Embracing a virtue-theoretic approach to knowledge is an effective way of capturing the normativity of knowledge or its special value as compared to epistemic states that fall short of it. This is because knowledge is an achievement, and achievements have value over and above mere lucky successes. An achievement is attributable to an agent in a way that a mere lucky success is not. Likewise, knowledge is attributable to an agent in a way that a mere lucky success (e.g., a true belief due to wishful thinking) is not. Greco explains that on this virtue-theoretic account of knowledge:

Knowledge is robustly normative....[K]nowledge is a kind of success from ability. Put another way, knowledge is a kind of achievement, or a kind of success for which the knower deserves credit. And in general, success from ability (i.e., achievement) has special value and deserves a special sort of credit. This is a ubiquitous and perfectly familiar sort of normativity. Thus we credit people for their athletic achievements, for their artistic achievements, and for their moral achievements. We also credit people for their intellectual achievements. Epistemic normativity is an instance of a more general, familiar kind. (Greco 2010: 7)

characterizing knowledge in a purely natural (i.e., descriptive) manner is not the way to go. Though, I take it that naturalizing knowledge and epistemology is the more controversial stand.

¹⁷⁴ For more on the epistemic basing relation see Ian Evans (2013).

Why are epistemic states that fall short of knowledge not successes due to the exercise of one's abilities? In the epistemic domain success occurs when one secures a true belief. States that fail in securing truth, such as justified false beliefs, are not successes, so they're not *successes* from one's abilities. Epistemic states that are successes are not necessarily successes *because of* one's abilities. True beliefs that fall short of knowledge are either not the result of reliable cognitive abilities or they are subject to luck. In the first case, epistemic virtues are conceived of as competences that tend to secure true beliefs, which include perception, memory, testimony, and induction. True beliefs that come from wishful thinking are not the product of such abilities. As such, they're not the result of a virtuous ability. Such a true belief doesn't manifest one's competence and so is not a success due to competence. In the second case, luck is incompatible with knowledge. A justified true belief falls short of knowledge when someone's belief is true due to luck. An agent in a standard Gettier case holds a true belief and exercises his intellectual abilities, but he does not hold a true belief *because of* his cognitive abilities. Instead his belief is ultimately true due to factors intervening to make his belief true; as such, his cognitive success is attributable to luck, and luck is incompatible with knowledge.¹⁷⁵

Granting the previous considerations, how does a creditworthy epistemic state, such as knowledge, underwrite, in part, a judgment that an agent is morally blameworthy? Phrased differently, why does knowledge ground blameworthiness?

¹⁷⁵ Following Pritchard (2005) the luck at issue in standard Gettier cases is called "intervening luck." This contrasts with what is called "environmental luck," which is at play in Barn Façade-type cases.

The normative and evaluative significance of knowledge produces strong reasons that favor the immorality of the agent's wrongful actions done from ignorance. For instance, Cleo's knowledge that slaves suffer produces strong reasons against keeping slaves. One might object, "Yes, but those reasons could never be reasons for Cleo to favor the impermissibility of slavery. She could not see those reasons as reasons counting against her actions of keeping slaves and forcing them to work without pay." However, again, this is to fall back into a pure accountability view of responsibility. It is true that, given this, Cleo could not have rectified her ignorance or believed differently, and later I will argue that someone can be morally responsible for their actions done from moral ignorance even if they could not have believed differently than they did. I will defend this attributionist stance. But now we are considering whether the actual profile of Cleo's total epistemic state captures her poor self-governance in relation to her keeping of slaves. We are asking what grounds the idea that her judgment is objectionable?

Cleo's all-things-considered judgment that slavery is permissible is a reflection on her and her poor weighing of the reasons she possessed. She had strong reasons to favor the immorality of slavery. That those reasons did not prove decisive, as other reasons outweighed those reasons, says something about Cleo. It reflects the absence of a good quality of moral judgment.

In addition to possessing non-moral knowledge in cases of moral ignorance, agent's total epistemic states often include additional mental states. For instance, Cleo had an experience as of the suffering of slaves. It matters that the content of the epistemic states is morally salient. For instance, going back to the Enlightened Mr.

Potter case, Mr. Potter had an experience as of his brother and others objecting to his business practices. Such objections provided evidence of moral disagreement. The cases involving the mob enforcer and the gang avenger involved *knowingly* causing people to suffer and die. And Cleo knowingly caused slaves to suffer by forcing them to work without pay and had experiences as of sympathy for their suffering.

Both the type of epistemic states and the content of the epistemic states produce strong moral reasons that favor acting in a way that shows due regard for the humanity of the victims of the person's wrongful actions. Such reasons show what is objectionable about the agent's all-things-considered judgment that their actions are morally permissible. They show that the persons affected by the agent's actions were not properly factored into an evaluation of what the agent had most reason to do. The suffering of the persons was outweighed by concerns such as what was best for the economy, for one's family or clan, or for one's political system. The persons' humanity was, in some sense, discounted by the agent's pursuit of ideals concerning what the agent took to be promoting economic, political, and prudential values. Such weighting of reasons reflects what the agent values in a way that is objectionable, as it reflects an unreliable mechanism of moral judgment.

Answering the challenge for the reasons the agent took to justify her actions and attitudes by citing such economic, political, and prudential values as justifying their all-things-considered judgment does not yield exculpation. Instead, it only aggravates and reinforces the agent's blameworthiness for her actions, as they were done in the absence of a good quality of moral judgment. It reinforces why the quality

of judgment is objectionable, as it involves lack of proper sensitivity to the humanity of the victims. Can I offer more in justifying this attributionist view?

You may wonder why Cleo did not come to the truth on the slavery issue? Doesn't history matter? After all, some people became abolitionists. They saw the moral truth. So, what does Cleo's moral ignorance indicate about her objectionable quality of moral judgment? Here it is helpful to look at why actual slaveholders and people in power did not recognize the truth of the immorality of those subject to their enslavement.¹⁷⁶ The work of Elizabeth Anderson is helpful in this regard. Those who persisted in moral ignorance regarding slavery did so in a way that reflected an unreliable mechanism of moral judgement. As Anderson indicates:

People in powerful positions tend to insulate themselves from the claims of those over whom they exercise power, to censor, discount, or misunderstand the claims of those beneath them, and to construct systems of law and moral accountability filled with loopholes through which they but not others can escape. So they rarely have the characteristic experiences through which they would learn that what they are doing to social inferiors is wrong. People holding powerful positions are also liable to confuse their own power with moral authority, and thereby confuse the self-serving orders they give to others with what others are morally obligated to do. Hence they are liable to misread challenges to their orders from below as signs of vice — of insubordination and insolence, irresponsibility, laziness, and so forth. The relatively powerless enjoy

¹⁷⁶ I thank Aaron Zimmermann for suggesting this turn to actual historical details.

no such luxuries. Hence people are prone to confuse their own desires with the right in rough proportion to their power. (2014: 8)

Cleo being in a position of power over her slaves and persisting in moral ignorance through weighting self-interest above the well-being of her slaves meant that she likely had in place some of these patterns of cognition. Such patterns of justification of morally wrong practices are a reflection on Cleo's moral cognition, namely that it is unreliable and distorted toward self-serving ends. But, one might protest that it is not Cleo's fault that her cognition is like this concerning the ignorance of her ways. Isn't it society's fault that the abolitionist movement did not yet grow strong enough to awaken her to the moral truth? Such a response is diversionary blame-shifting that is unhelpful. If a person's moral mechanism of judgment is broken and that broken mechanism reflects what the person judges true and good, then that broken mechanism and the outputs of it are attributable to that person. The person is rightly held accountable for such outputs even if, at the time, they could not see the needs and interests of those harmed by their actions as sufficiently weighty. Though a social movement of abolitionists could have awoken Cleo to her moral ignorance, it is problematic to claim that Cleo's moral ignorance was not attributable to her in a way sufficient for deserving blame because a social movement had not yet strengthened to the point of jarring Cleo into moral truth. It is problematic because it shifts blame away from the unreliable mechanism producing and perpetuating the moral ignorance onto external sources that are purported saviors aiming to deliver minds from moral darkness. With the broken mechanism before us, and with that mechanism owed by Cleo, as it reflects what she values and thinks true, such attributability secures accountability for the products of that unreliable mechanism of moral judgment.

3.7 Interlude on the Nature of Moral Responsibility

To this point, I have provided reasons to think attributability sufficient for accountability in diagnosing cases of moral ignorance. Such reasons might sway a reflective agnostic regarding the debate between attributability and accountability concerning the nature of moral responsibility. Yet a committed accountability theorist may want more. He may think I have begged the question against accountability by assuming attributability true and using its features in arguing that attributability best captures cases involving moral ignorance. A full exposition on the nature of moral responsibility goes beyond the scope of this current project, but I want to look at a small slice of the literature to show why one might think accountability theorists misguided when it comes to wanting more out of the desert-based notion of moral responsibility than what is provided by New Attributionism.

As mentioned in section 3.3, Gary Watson (1996, 2004) thinks of responsibility as coming in two distinct versions. Attributability focuses on aretaic appraisals of a person's moral character. Such views are sometimes called self-disclosure views or real self views. Accountability focuses on whether it would be appropriate to hold someone accountable for her behavior. This happens when it is fair to impose on the person certain burdens such as social sanctions licensed by the reactive attitudes. Attributability is not sufficient for accountability on this proposal. It is not enough that an agent's attitudes and actions reflect who she is as a practical agent. Being subject to adverse treatment in the form of the reactive attitudes and forms of punishment also requires that the agent had a fair opportunity to avoid being exposed to such

treatment.¹⁷⁷ If the agent had a morally deprived, insulated, or abusive upbringing that (in some sense) made them into who they are as the author of their attitudes and actions, then the agent is not accountable for those attitudes and actions even if they are attributable to her. Watson's approach is nice in that it captures why someone like, to use Watson's hallmark example, a "moral monster" like Robert Alton Harris is not subject to the full range of reactive attitudes. He suffered a horrific and abusive upbringing. Though he grew into an adult that inflicted harm, torture, and death on human and animal victims, he has suffered enough and is not a fair target of further adverse treatment. This is because he did not have a fair opportunity to avoid becoming the person he became. Yet, Watson's two-prong view of moral responsibility is subject to an important criticism. This criticism or, rather, confusion underlies why I think many accountability theorists think attributability insufficient for accountability.¹⁷⁸

Angela Smith (2008) argues that Watson's two faces of moral responsibility rest on an ambiguity in the notion of accountability.¹⁷⁹ Watson argues that conditions required for accountability go beyond conditions required for attributability. Attributability has "different and less stringent" conditions than those required for accountability. Further, as Smith summarizes regarding Watson's stance, "self-disclosure [attributability] views do not justify many of the activities and responses associated with our current practices of moral responsibility" (2008: 377). Smith

¹⁷⁷ Watson (1996: 237).

¹⁷⁸ Instead of embracing one form of moral responsibility with different varieties or species, the accountability theorist thinks there are two different genus's or types of moral responsibility. By contrast, New Attributionism collapses the genus's into one genus with different species.

¹⁷⁹ For an additional exposition of the distinction between being responsible versus holding responsible see Smith (2007).

responds to Watson that there is an ambiguity in “accountability”, namely judging blameworthy versus active blaming. Smith grants that thinking of accountability in terms of activity blaming is thinking of it as a notion that is governed by conditions beyond attributability, such as the condition Watson (1996: 237) cites as the agent having, “had a fair opportunity to avoid being subject to that adverse treatment.” Smith explains that after clearing up the ambiguity this shows that,

there is a distinction between judging a person to be responsible and culpable for her behavior, on the one hand, and engaging in various forms of blaming activity (including punishment), on the other. But the fact that there may be further conditions governing our blaming activities does not, I believe, support the claim that attributability views are working with a different and weaker conception of the relevant forms of moral appraisal. (2008: 377)

Smith grants that accountability disambiguated as involving blaming practices may have additional conditions on it. This is because actually blaming people for actions for which they are blameworthy imposes more of a burden of justification on those practices. In the Robert Alton Harris case, it may be inappropriate to subject him to various types of ill-treatment, such as actually blaming and punishing him, given his upbringing.¹⁸⁰ Though he is blameworthy for his actions, as he is the author of them in the attributability sense, imposing further burdens on him beyond that required to keep him from harming others in the future may seem unnecessarily harsh or cruel. Judgements of culpability for attitudes and actions are governed by attributability conditions. These conditions concern a different disambiguation of “accountability,”

¹⁸⁰ I discuss this case in more detail in Chapter 4. For details of the case see Watson (1993).

namely it being appropriate for you to give an account of the reasons you took to justify doing what you did. This concerns what you actually value, what you think true and good, and whether those things reflect who you are as a moral agent. Robert Alton Harris thought his horrific actions appropriate. He embraced them upon reflection and, obviously, they severely impaired the relationships people could enter into with him. They are attributable to him, he is answerable for them, and he is blameworthy and accountable for his actions.

Lastly, to close this interlude, let me mention that thinking of accountability in terms of actual blaming responses is governed by additional factors. That is, as Smith clarifies, active blame includes, “a variety of moral and non-moral norms...which have nothing to do with the basic conditions of moral responsibility.” These norms involving determining whether one is justified in blaming the person. Such actual blaming responses are governed by, “our relation to the person, our stake in the matter, the significance of the fault, and the person’s own response to her failure.”¹⁸¹ They are also governed, as already mentioned, by what Watson mistakenly took as definitive of accountability, such as whether the person had a fair chance to avoid adverse treatment in terms of being subject to the reactive attitudes. These concerns with history, what the agent could have done differently, and standing to blame determine whether public or social censure is appropriate and whether and to what extent the person should be punished.¹⁸²

¹⁸¹ For example, it may be inappropriate to actually blame someone when they express sincere remorse for their actions. However, cavalier dismissals of the moral magnitude of the faults and harms at issue may amplify the degree to which one is justified in actually blaming and punishing the person.

¹⁸² For someone who resists the attributionist line and posits three distinct forms of responsibility, which are underwritten by three different agential capacities, see the work of David Shoemaker (2011). Shoemaker targets Smith’s view, but the conceptions at play get hard to track, as Shoemaker is not

3.8 An Epistemic Frankfurt-style Case

In this section, I will introduce an objection to the idea that epistemic states can inculcate an agent concerning blameworthiness. I will analyze commitments underlying the objection, and then I will respond to the objection by introducing an epistemic Frankfurt-type case showing how, for the attributionist, an agent can be morally responsible for his acts done from ignorance even though it was not possible for him to have rectified his ignorance and believed otherwise.

William FitzPatrick (2017) objects to the sort of move I made in embracing New Attributionism as a way to make sense of why moral ignorance never exculpates.¹⁸³ In contrast to moral responsibility understood as accountability, New Attributionism does not focus on the history of what led to the moral ignorance or whether it would have been reasonable to expect the agent to have rectified her ignorance in the past. What matters for attributionism is whether, at the time of action or belief-formation, the action or attitude is attributable to the agent such that it reflected the agent's evaluative commitments, quality of will, or moral character. Recall, as Talbert (2016: 139) explains, "Attributionism is concerned with why a person *actually* did what she did, not with whether it would have been rational for her to do something else." FitzPatrick thinks this attributionist move results in a shallow form of moral

using the terms "attributability" and "accountability" as understood by Watson. This is why, in part, a full exposition of the nature of moral responsibility is its own project. Disentangling all the notions in play in the literature is a task unto itself. However, the basic notions stemming from Watson and the attributionist expositions and modifications are somewhat manageable. This is why, for simplicity, I followed Talbert in lumping them under the label New Attributionism. This just signals that New Attributionists do not restrict attributionism to aretaic appraisals. New Attributionism allows for character-based appraisals but is not limited to them because a person can do something out of character but it can still be attributable to them.

¹⁸³ FitzPatrick is targeting Harman's claim, which is a claim I embrace, that, "we are morally obligated to believe the moral truths relevant to our actions, and thus not to believe false moral claims relevant to our actions" (Harman 2011: 459).

criticism—one that's not able to capture desert-based, full-blooded blameworthiness. FitzPatrick holds that, "*deservingness* of blame...goes beyond mere *fittingness* of negative evaluation of character, attitude, judgment, or behavior, beyond expression of values, and beyond the sorts of reactive modifications of attitudes, intentions, and relationships that Scanlon (2008) takes to be the essence of blame" (2017: 33). FitzPatrick argues for this point by way of the following case:

Insular and Devout. Suppose Daniel was born into a small, insular community and raised to be deferential to the authoritarian leaders, whose teachings are based on religious texts they take to express God's will. They have deprived him of broad education and limited his exposure to alternative ideas or role models, and he accepts that this is to protect him from "corruption." As a young adult, he encountered outsiders on occasional trips to town, but he unsurprisingly views them with deep suspicion, seeing them as threatening players in the conspiracies that populate his imagination. When one of them tries to engage him in critical discussion of his group's virulently anti-homosexual beliefs (which he shares, reinforced by his own revulsion to homosexuality, which he interprets as a perception of its depravity), he responds with disdain and with contemptuous denunciation of a nearby homosexual couple, viewing this as an expression of righteous indignation that is morally justified by the magnitude of their sin and the danger that they pose to the whole community by inviting God's wrath. (FitzPatrick 2017: 34)

FitzPatrick thinks Daniel is a moral agent. He is sane and has the capacity to comprehend moral reasons, but given his insular upbringing the question becomes

whether Daniel had a reasonable opportunity to rectify his moral ignorance. FitzPatrick (2017: 35) claims that, “Since it is not his fault that he was raised in these conditions, under which he had little opportunity to have known better than he does or to have developed a better set of values and attitudes, it is, first of all, *not his fault that he came to be this way*. And if it is not his fault that he is this way, then it is hard to see how he could genuinely deserve *blame* for this...even granting that other moral criticisms are appropriate.” Daniel didn’t have sufficient control over the formation of his moral values and character. He also didn’t have sufficient opportunities to shape his moral character differently given his impoverished upbringing. So how can Daniel deserve blame for behavior that naturally emanated from a moral character he is not blameworthy for possessing? As FitzPatrick explains:

[I]f we accept that [Daniel] doesn’t deserve blame for his condition, and so is not blameworthy for his moral ignorance..., it is equally hard to see how he can *deserve blame* and thus be blameworthy for the behavior that straightforwardly issues from it. It may be tempting to posit such blameworthiness insofar as he chose his action based on attitudes and values he endorsed, without coercion, knowing the offence and hurt it would cause. But while that is significant in determining our response, it is not enough to ground deservingness of blame. The same consideration that undermines the thought that Daniel deserves blame for his ignorance—i.e., the fact that there was nothing he could reasonably have done to avoid it—equally undermines the thought that he deserves blame for the behavior that issues from those beliefs and attitudes. How else could he be expected to behave at this point? (FitzPatrick 2017: 35)

FitzPatrick reinforces his commitment to the claim that one is blameworthy for one's act done from moral ignorance *only if* one is blameworthy for one's ignorance. Daniel is not blameworthy for his ignorance, so he is not blameworthy for his acts done from ignorance. This is because a counterfactual is true: there is nothing Daniel could reasonably have been expected to have done to have corrected his ignorance. Given this counterfactual, Daniel is not blameworthy for his ignorance. Now FitzPatrick embeds the counterfactual claim about reasonable expectations within a biconditional. Concerning necessary and sufficient conditions for culpable (i.e., blameworthy) ignorance:

Reasonable Expectations. Ignorance is culpable just in case (and because) the agent *could reasonably have been expected* (in the normative, not merely predictive sense) to have remedied this ignorance.¹⁸⁴

In the Insular and Devout case, the Reasonable Expectations biconditional underlies the argument that Daniel is not blameworthy for his actions. I will argue against one direction of the biconditional. I will argue that even if, given an agent's cultural context and capabilities, the agent could not reasonably have been expected to have corrected his ignorance it can still be the case that his ignorance is culpable.

¹⁸⁴ FitzPatrick (2017: 29). This biconditional version of Reasonable Expectations contrasts with the conditional version found in the Culpable Ignorance (CI) principle, "Ignorance, whether circumstantial or normative, is culpable *if* the agent could reasonably have been expected to take measures that would have corrected or avoided it, given his or her capabilities and the opportunities provided by the social context, but failed to do so either due to akrasia or due to the culpable, nonakratic exercise of...vices." (FitzPatrick 2008: 609). In chapter 2 of this dissertation, I embraced (CI). What I say here doesn't conflict with that earlier commitment. I still agree that reasonable expectations to correct ignorance plus ignorance implies the ignorance is culpable. What I do not embrace is one direction of the biconditional version of the Reasonable Expectations claim. I do not embrace the claim that if the agent could not have been reasonably expected to have rectified her ignorance, then her ignorance is exculpatory.

To begin to question FitzPatrick's view, it is helpful to see that there is an attributionist reading of cases like the one involving Daniel available. On such a reading, the truth of the counterfactual in the Reasonable Expectations claim does not show that the actual etiology or reasons for which an action occurred does not disclose the agent's own sincerely held beliefs and values. Even if the agent could not reasonably have been expected to believe otherwise (i.e., to not be ignorant), the agent's attitudes and actions can still express an objectionable quality of moral judgment. The actual etiology of behavior can, contrary to FitzPatrick, ground blameworthiness precisely for the reason FitzPatrick mentioned. For the attributionist, Daniel is blameworthy because, "he chose his action based on attitudes and values he endorsed, without coercion, knowing the offence and hurt it would cause" (FitzPatrick 2017: 35). The actual etiology of Daniel's false belief is of the right kind.¹⁸⁵ It is inculpatory. His false moral belief that his act is permissible is not exculpatory.

FitzPatrick might call into question the attributionist reading of the Insular and Devout case by drawing a distinction between the moral reactions involved with full-blooded blameworthiness and lesser moral reactions.¹⁸⁶ The etiology of Daniel's moral ignorance makes him deserve lesser moral reactions but not full-blooded reactive attitudes. Regarding these lesser reactions to Daniel in the Insular and Devout case FitzPatrick says, "No doubt it is appropriate for those subject to his vitriol to feel offended and hurt, to demand that he cease and apologize, and perhaps to view him

¹⁸⁵ This differs from Rosen's view according to which agents are responsible for their false moral beliefs only when there is an akratic episode in the etiology of belief. FitzPatrick's earlier work adds a vice-condition to the grounding of blameworthiness.

¹⁸⁶ This line is suggested by FitzPatrick's comments.

as currently ineligible for trust and friendship” (2017: 35). Why stop short of Daniel deserving blame, indignation, and resentment? According to FitzPatrick, Daniel’s history is exculpating, as:

Blameworthiness for unwitting wrongdoing requires a greater agential role, in a context of decent epistemic opportunities, in the flawed formation of belief and character through reasonably avoidable poor choices....Notice (to take an extreme, fanciful case) that we would not think someone to be *deserving of blame* for nasty behavior expressive of false and offensive moral beliefs and vices that were simply induced by an evil neuroscientist shortly before, even though it’s true that the behavior is itself uncoerced and responsive to judgements he currently endorses, which are therefore all currently attributable to him. Many negative responses might be fitting and many critical things might be said about him; he may be dangerous and revolting and someone to be avoided. But he is not deserving of blame. The agent’s role in the epistemic and character-forming history matters. (2017: 36)

In response, the attributionist might claim that what FitzPatrick has done is aptly describe why Daniel is answerable for his attitudes and actions. They are attributable to Daniel, as they reflect his evaluative judgments about the value of persons and practices. He is answerable for them and blameworthy as a result. The “lesser” moral reactions are morally significant modifications of his victims and others relationship with him. Daniel’s failure to show reasonable regard for others’ moral interests warrants certain interpersonally significant modifications of relationships, such as him no longer being a candidate for trust and friendship. That such interpersonal

modifications are licensed by Daniel's actions is a sign that the full-blooded reactive attitudes are warranted. Daniel is blameworthy. He is deserving of blame.

Additionally, the attributionist might leverage the argument of Smith against Watson, as discussed in the previous section. Someone is only judging that Daniel is not blameworthy for his public slandering of homosexuals because they are thinking of accountability as involving actual blaming practices. Once being blameworthy is distinguished from conditions under which it is appropriate to actually blame an agent, which include norms involved in standing to blame, how the agent responded to their sins, and historical factors concerning correctability of ignorance, then it is possible to see that Daniel is blameworthy in the relevant sense. His actions are attributable to him and this suffices to make him accountable for his actions. Whether it is appropriate to actually hold him accountable is sensitive to the factors FitzPatrick mentions. Such factors apply most readily to whether it is fair to subject the agent to adverse treatment in actual applications of the reactive attitudes.

With the attributionist reading of the Insular and Devout case on the table, I will take the attributionist response to FitzPatrick a step further. I will argue that even if the counterfactual in the Reasonable Expectations claim is true, such that it's not reasonable to expect Daniel to have corrected his moral ignorance and believed otherwise, his moral ignorance is still culpable. I will do this via an epistemic Frankfurt-style case. But first I will provide more of the dialectic behind the case.

Returning to a point made by FitzPatrick in the previous quote, I disagree with FitzPatrick that belief and vice implantation by an evil neuroscientist right before action can undermine a person being deserving of blame, even though "it's true that

the behavior is itself uncoerced and responsive to judgements he currently endorses.” This is because there is a control condition on the mechanism that produces one’s actions and attitudes. If the evil neuroscientist implanted the beliefs and vices, as FitzPatrick says, shortly before the behavior in question, then there is a failure of control of the sort required for moral responsibility and blameworthiness. We are no longer talking about the epistemic condition *per se*. We are now talking about epistemic requirements on the control condition for moral responsibility.¹⁸⁷ In such a case, it is not moral ignorance that is exculpatory, it is lack of control that excuses.

One may object: Why doesn’t something similar apply to cases involving moral ignorance? Isn’t FitzPatrick’s analogy that Daniel’s insular religious upbringing was relevantly like an evil neuroscientist implanting beliefs and vices prior to the time of action?

In reply, I think that mechanism ownership creates a relevant difference between the cases taken to be analogous. Daniel had time to take ownership of his belief-forming mechanism, whereas a belief-forming mechanism implanted right before the time of action is not owned by the agent. As a result, beliefs and behaviors output by such an implanted mechanism would not be attributable to the agent.¹⁸⁸ Even if the agent’s nasty behavior aligned with his values, the thing that produced the ignorance and behavior would not be attributable to the agent. He would not be

¹⁸⁷ For a view that embraces such epistemic components on the control condition see Fischer and Ravizza (1998). On this account, control requires a reasons-responsive mechanism that one owns. And ownership of one’s mechanism of action involves having certain beliefs about it.

¹⁸⁸ If we add to the evil neuroscientist case that the agent had time to take ownership of the beliefs and vices producing his attitudes and actions, then I think the intuition FitzPatrick elicits falls away. Such an agent would be deserving of blame for bad behavior that is the result of moral beliefs and vices induced by an evil neuroscientist but were reflectively owned and endorsed by the agent. Such a belief-forming mechanism would be properly attributable to the agent and the agent would be answerable for the outputs of that mechanism.

answerable for his actions in the requisite sense. This differs from Daniel who took ownership of his mechanism of action over time. He owns his mechanism and is answerable for its outputs.

At this point, it is worth noting that FitzPatrick acknowledges the pull toward the view I defend. FitzPatrick (2017: 35) confesses that, “It may be tempting to posit such blameworthiness insofar as he chose his action based on attitudes and values [Daniel] endorsed, without coercion, *knowing* the offence and hurt it would cause.” My view embraces this temptation. Deciding to act knowing the offence and harm it causes is blameworthy-grounding. Such knowledge is an achievement that is creditworthy, and it expresses a connection between the action chosen and the agent’s evaluative commitments underlying that choice.¹⁸⁹ Thus, actions such as Daniel’s contemptuous denunciation of the homosexual couple generate direct responsibility and blameworthiness. We need not trace Daniel’s responsibility back to whether he is culpable for his moral ignorance. We need not determine whether he had a fair opportunity to rectify his ignorance or whether his insular religious upbringing prevented him from having such an opportunity. Instead, Daniel’s total epistemic state inculcates him. This makes Daniel a legitimate target of the reactive attitudes. It makes him morally responsible and deserving of blame. Although, I do grant we may need to do such tracing when justifying actually blaming and punishing him. How does FitzPatrick respond to this tempting line on the Daniel case that I embraced, and why is his response problematic?

¹⁸⁹ Knowingly offending and hurting others, though not knowing that doing so is wrong, implicates a poor quality of judgment in this way. Such knowledge is attributable to the agent, and it represents a link between the action and the agent’s values underlying the action.

In response to the idea that knowledge of non-moral facts can ground blameworthiness in the case involving Daniel, FitzPatrick argues that:

while it's true that [Daniel] has this knowledge, the problem is precisely that he fails to understand its significance, since he falsely believes that such offence and hurt are justified here (as offence and hurt sometimes in fact are). So in light of his moral ignorance, his circumstantial [i.e., factual] knowledge that he is offending and hurting people is not something that can be expected to register with him as a reason to refrain from what he is doing, and so it hardly seems like a sufficient ground of blameworthiness for his behavior. (FitzPatrick 2017: 35 n.3)

FitzPatrick's comments suggest the following principle:

You are not blameworthy for an action done from moral ignorance if you could not reasonably have been expected, given your whole history, to see a decisive moral reason as a reason to believe your action impermissible and to refrain from performing it.

FitzPatrick is thinking that if an agent could not have been expected to see a reason, which could have made a difference, as a reason against performing an immoral action, then the agent's moral ignorance is blameless, and the agent is not blameworthy for the action done from that ignorance. A principle underlying this claim is the principle that our reasonable expectations track whether the agent could have believed other than she did. If a decisive reason could not have mentally registered as a sufficient reason to believe an action morally wrong or refrain from the action, then it is not reasonable to expect the agent to have acted differently or to have rectified her ignorance. This suggests the following principle:

You must have had the opportunity to believe other than you did if it reasonable for us to expect that you could have rectified your ignorance.

The absence of opportunities to believe otherwise, given your history, cultural context, or capabilities, undermines it being reasonable to expect you to have corrected your moral ignorance. We can put these principles together to derive FitzPatrick's conclusion that Daniel is not blameworthy for his moral ignorance. The argument runs as follows:

1. If Daniel could have been reasonably expected to have remedied his moral ignorance, then he could have believed otherwise.
2. Daniel could not have believed otherwise.¹⁹⁰
3. So, Daniel could not have been reasonably expected to have remedied his moral ignorance. (1,2 MT)
4. If Daniel could not have been reasonably expected to have remedied his moral ignorance, then Daniel's moral ignorance is not culpable.¹⁹¹
5. Thus, Daniel's moral ignorance is not culpable. (3,4 MP)

The argument above lays bare how FitzPatrick is thinking that Daniel's moral ignorance is exculpatory. I will embrace the first part of the argument. I agree that Daniel could not have been reasonably expected to have remedied his moral ignorance, given his opportunities and upbringing. However, I do not think it follows from this that his moral ignorance is not culpable. Leveraging what I already argued against the accountability theorist in this section and the previous section, I will cast doubt on premise 4 in the argument, which is the right-to-left conditional in the Reasonable Expectations biconditional endorsed by FitzPatrick.

¹⁹⁰ Given that decisive reasons couldn't have registered as reasons to believe and act differently.

¹⁹¹ This is the right-to-left conditional from the Reasonable Expectations biconditional.

Using an epistemic Frankfurt-style case, I will argue that it is not reasonable for us to expect the agent in the case to have remedied his ignorance, as the agent could not have believed otherwise, but it is also true that the agent's ignorance is culpable on the basis of attributionist considerations. Premise 4 is problematic. Though Daniel could not have been reasonably expected to have remedied his moral ignorance, his moral ignorance is still culpable. I adapt the Insular and Devout case to make it a Frankfurt-style case. First a brief word about Frankfurt cases.

Harry G Frankfurt (1969) argued against the Principle of Alternative Possibilities (PAP): A person is morally responsible for what he has done only if he could have done otherwise. Frankfurt-style counterexamples involve an actual sequence of events whereby the agent decides to perform a reprehensible action of his own accord. Yet there is an intervener monitoring the agent's decision process via an electronic chip implanted in the person's brain. Were the agent to decide to perform an action other than the one the intervener wanted the agent to perform, the intervener would intervene and ensure that the agent performed the desired action. Yet, the agent makes the "right" choice, and the intervener does not interfere with the actual sequence of events. Thus, the agent is morally responsible for his action, and it is true that the agent could not have done other than he did. PAP is false.

Epistemic Frankfurt-style cases focus on whether an agent could believe other than they did, which is relevant to cases involving moral ignorance. Given that moral responsibility covers both praiseworthy and blameworthy actions, and we are only focusing on wrongful actions for which the agent may be blameworthy, we can formulate an epistemic version of PAP in such terms:

E-PAP: A person is blameworthy for what he has believed only if he could have believed otherwise.

With E-PAP in mind, let's return to the case involving Daniel. Recall that I grant premise 1 in FitzPatrick's argument. I grant that Daniel being able to have believed other than he did is required for it to be the case that Daniel could have been reasonably expected to have remedied his moral ignorance. Our reasonable expectations track whether Daniel could have believed other than he did. However, I do not think Daniel's blameworthiness for his moral ignorance is determined by our reasonable expectations. That is, I think E-PAP is false. Daniel is blameworthy for what he has believed even if he could not have believed otherwise. Premise 4 in FitzPatrick's argument is problematic. It is true that Daniel could not have been reasonably expected to have remedied his moral ignorance, given that decisive reasons couldn't register as reasons to believe differently, yet it is also true that Daniel is blameworthy for what he has believed (i.e., it is false that Daniel's moral ignorance is not culpable). Daniel's moral ignorance fails to exculpate. Thinking otherwise is to fall back into the confusion of thinking that deserving blame is the same as actually blaming.

To demonstrate the falsity of E-PAP and why premise 4 is troubling, here is an epistemic Frankfurt-style case involving a person named Jones:

Jones is a member of a small, insular community. He has been raised to defer to authoritarian leaders, whose teaching are based on religious texts they take to be expressions of God's will. Black is an authoritarian leader in the community and is also a skilled neuroscientist. Unbeknownst to Jones, Black has inserted a chip in Jones' brain. This chip allows Black to monitor and control Jones' behavior by

controlling his belief-forming processes, which output beliefs that provide reasons for Jones to perform actions. On a trip outside the community Jones encounters Smith—a person who challenges Jones group’s virulently anti-homosexual beliefs (which Jones shares, reinforced by his own revulsion to homosexuality, which he interprets as a perception of its depravity). Black is monitoring Jones’ belief-forming processes and is ready to intervene if Jones shows any sign of not strongly rebuffing Smith’s challenge to his moral beliefs about homosexuality. Though Jones knows that he is going to say and do something that could be offensive and hurtful, Jones falsely believes that such actions are justified, given what’s at stake in terms of God’s hatred of the sin of homosexual relations. Given his moral ignorance, his factual knowledge that he is about to say and do something that is offensive and hurtful does not register with him as a reason to refrain from what he is about to do. Black is aware of this via monitoring Jones’ mental processes. Jones has not formed the belief that he should refrain from a forceful rebuff of Smith’s challenge, and Jones shows no signs of doubting his moral convictions. So, Black need not intervene. Jones decides to respond forcefully to Smith’s challenge by making an example out of a nearby homosexual couple. He contemptuously denounces the couple, viewing this as a rebuff to Smith’s challenge and an expression of righteous indignation that is morally justified given the magnitude of the couple’s sin and the danger that they pose to the whole community by inviting God’s wrath.¹⁹²

¹⁹² This example is a hybrid of a Frankfurt case attributed to Fischer (2010) and the Insular and Devout case from FitzPatrick (2017).

Given this example, E-PAP is false. Jones could not have believed otherwise, yet Jones is blameworthy for believing that it is morally permissible to offend and hurt a homosexual couple to rebuff a challenge to his moral beliefs. This is the case precisely for the reasons the attributionist detailed. Even though, given Jones' moral values, that his comments and actions will hurt and offend the couple could not register as a reason to refrain from hurting and offending them, Jones' judgment reflected a lack of due regard for the humanity of the homosexual couple, such that Jones judged it permissible to publicly humiliate, hurt, and offend them to make a point in responding to a challenge to his moral beliefs. Jones' attitudes and actions are attributable to him.¹⁹³ Thus he is answerable and blameworthy for them for a similar reason—they were accompanied with the absence of good judgment.

Given the considerations above, blameworthiness for false moral beliefs does not require the ability to believe otherwise. Premise 4 in FitzPatrick's argument is false.¹⁹⁴ In the Insular and Devout case, Daniel could not have been reasonably

¹⁹³ In the actual sequence of events Black did not intervene. Unaided by Black's intervention, Jones' beliefs and actions reflected his values. Even though Jones' moral architecture was culturally determined, as he came to possess those values due to his insular upbringing, he still reflectively endorsed those values. Given that his beliefs and actions were not tampered with by the intervener Black, those beliefs and actions were responsive to his judgment about what he had most reason to do.

¹⁹⁴ Why couldn't FitzPatrick just resist the intuition that Jones is blameworthy? Thanks to Hieronymi for raising this worry. The standard way of proceeding in Frankfurt-style cases is that the person to whom they are addressed endorses the intuition that Jones is blameworthy. If I simply presented the Frankfurt-style case on its own or in isolation, FitzPatrick would simply deny the intuition. But, I have argued against FitzPatrick being able to simply deny the intuition. This is because I argued against his position that Daniel is not deserving of blame in the Insular and Devout case, which is a case of a deprived history due to inculcation. Now, when we get to the Frankfurt-case FitzPatrick could deny that he shares the intuition that Jones is blameworthy in this case. But now he needs to say why. I cast doubt on the reasoning why he is likely to hold such an intuition. He could try to cite a morally relevant difference between the cases, but inculcation plus brain monitoring due to chip implantation (monitoring that never alters the action) does not generate a morally relevant difference from the perspective of blameworthiness. Both cases involve a morally deprived history that leads to the bad values and the wrongful action done from ignorance. The fact that the Frankfurt-style case involved technology, and the other case did not, does not generate a morally relevant difference between the cases—at least simply regarding the intuition that Daniel and Jones are blameworthy for their harmful actions. What is common between the cases is that, at the time of action, Daniel and Jones knew what

expected to have remedied his moral ignorance, as he could not have believed otherwise, but Daniel's moral ignorance is culpable. He is answerable for his ignorance and acts done from that ignorance. Thus, the Insular and Devout case does not necessarily undermine the No Excuse thesis. Daniel's moral ignorance fails to excuse him from moral responsibility and blameworthiness.

3.9 Conclusion

In this chapter, I argued in favor of the No Excuse thesis. I argued for the No Excuse thesis by showing how a defender of the thesis, namely Harman (2011, 2014), embraced a conception of moral responsibility that subjected her view to a counterexample that undermined her defense of No Excuse. I identified the problem with Harman's view as the version of attributability about responsibility underlying the view. To overcome this deficiency, I embraced a more all-encompassing conception of attributability. Then I analyzed the case that created problems for Harman's view and demonstrated how that case supported the No Excuse thesis given the broader conception of responsibility.

I also argued that moral ignorance *per se* never excuses one from blameworthiness because, in the cases that test whether moral ignorance excuses, the agent possesses knowledge of the relevant non-moral facts. Given that unwitting wrongful acts in such cases are accompanied by factual knowledge and experiential epistemic states, such epistemic states point to the fact that the agent acted in the

they were doing, and their action reflected their deeply held values. So, I do not think going the morally relevant difference route will work for FitzPatrick in denying the Frankfurt-style case intuition. The other option is to counter my previous arguments of his analysis of the case in question. That remains an open option, and I would welcome being able to reply to his response were he to give such a response.

absence of good moral judgment. Such epistemic states point to something morally objectionable in the person's evaluative judgments regarding others. In fact, I argued that such epistemic states aggravate blameworthiness. As such, the No Excuse thesis is true in such cases—moral ignorance never excuses.

Lastly, I considered an objection to my view as put forward by FitzPatrick. After breaking down his argument and exposing the commitments underlying it, I proposed a Frankfurt-style case focused on the agent's epistemic states. This case undermined his argument against my defense of No Excuse. Our reasonable expectations about what agents could have done to correct their ignorance may track whether they could have believed differently than they did given their history and upbringing, but our moral evaluation of their blameworthiness is not hostage to such expectations.

In closing, as stated at the start of this chapter, the truth of the No Excuse thesis has practical, everyday implications. Given the arguments in this chapter, a posture of humility is warranted. We must admit that it is likely that our cultural context is limited concerning certain moral truths. We must admit that it is likely that we are in the grip of false moral views. We must accept that we are morally responsible and blameworthy for more than we realize. Our moral ignorance will not let us off the hook in this regard.

CHAPTER 4: EVIL DICTATORS, EPISTEMIC DIFFICULTY, AND MORAL LUCK

Look out behind you

Jojo's got his gun

He wouldn't mean to

But you know he likes his fun

–Boz Scaggs “JoJo”

4.0 The Skeptical Challenge Redux

In this final chapter, it helps to bring back online the issue animating this dissertation—the threat of skepticism concerning moral responsibility. In chapter 1, I introduced the skeptical argument as put forward by Rosen (2004) and formalized by FitzPatrick (2008). The conclusion of the skeptical argument is that it is not possible to know, concerning any case, whether a subject is responsible for an action. As Rosen claims, “You should hold that confident positive judgments of responsibility are never justified” (Rosen 2004: 295). To be justified in believing that a subject is responsible requires evidence that the subject satisfies conditions sufficient for responsibility. On Rosen’s internalist picture, culpable ignorance must trace back to acts of clear-eyed akrasia. However, there is always uncertainty from the agent’s perspective and the responsibility ascriber’s perspective regarding whether a benighting act was knowingly performed in full view of the relevant facts. Given such uncertainty, suspension of judgment is the only justified attitude to take toward whether an agent

is responsible for an act done from moral ignorance. In this chapter, I return to directly countering this skeptical challenge.

One might worry that I have not neutralized the skeptical challenge to moral responsibility. The arguments in chapters 2 and 3 were more metaphysically focused. In those chapters I defended the idea that the conditions sufficient for blameworthiness are broader than the narrow conditions proposed by the responsibility skeptic. I argued that the non-akratic exercise of vices can make a person's ignorance culpable, and I argued that moral ignorance *per se* is never exculpatory, as acts done from moral ignorance reflect objectionable judgments regarding how it is permissible to treat persons and they are accompanied by inculpatory knowledge of the relevant non-moral facts. However, the skeptical challenge is an *epistemic* challenge. It concerns the possibility of warranted ascriptions of responsibility. One might object that even if the metaphysical conditions I endorsed as sufficient for grounding responsibility are satisfiable in principle we cannot form justified beliefs in practice concerning whether those conditions are satisfied. We do not have good evidence to justify ascriptions of responsibility. To counter this worry, and to directly respond to the skeptical challenge, I will look at how people actually ascribe responsibility in specific cases. I will look at experimental results involving cases of moral ignorance. I will argue that these ascriptions of responsibility are justified and align with a version of New Attributionism. I will thus directly counter the skeptical challenge and bolster the earlier arguments of this dissertation.

4.1 Real Self Views

In this section, I analyze key components of the view of moral responsibility at issue. First let me connect New Attributionism to what, following Susan Wolf (1987), are called Real Self Views of moral responsibility.

In chapter 3, I introduced New Attributionism as a label for a species of views concerning moral responsibility.¹⁹⁵ These views combine aspects of attributionism and accountability. Attributability of bad behavior to a person licenses reactive attitudes toward that person. Such attributability makes it fair to judge the person blameworthy. Accountability-reactions are licensed by attributability concerns. Attributability is a synchronic (i.e., at-a-time) concern. It asks whether the action, at the time of action, reflected the agent's evaluative judgments and commitments. Attributionist views are less concerned with diachronic (i.e., across-time) considerations. Such views are less concerned with surveying the time-slices that preceded the wrongful action, such as whether the person's values were the result of a deprived childhood or whether the subject had opportunities in the past to rectify their moral ignorance.¹⁹⁶

A subcategory of New Attributionist views are Real Self Views. According to these views, acts that are attributable to an agent are acts that stem from the agent's "real self." Identifying an agent's real self involves finding, as Faraci and Shoemaker (2010) indicate, "a subset of an agent's motivating psychological elements as privileged for self-determination and responsibility, such that as long as one's actions

¹⁹⁵ I followed Talbert (2016) in making this useful classification.

¹⁹⁶ As such, New Attributionist views embrace moral luck. They deny the principle stating that we are responsible for something only if that thing was in our voluntary control. I discuss moral luck at the end of this chapter.

are ultimately governed by this subset, they count as one's own and thus render one eligible for responsibility-responses to them." Real Self Views diverge based on what they take to be required for a feature of an agent's psychology to count as a member of that set. Harry G. Frankfurt (1971) thinks that higher-order desires and first-order desires meshing or aligning make such psychological elements part of the privileged set comprising one's real self. Gary Watson (1975: 215) holds that an agent's real self is determined by what an agent values, and, "an agent's values consist in those principles and ends which he—in a cool and non-self-deceptive moment—articulates as definitive of the good, fulfilling, and defensible life." Such goals and normative principles that an agent endorses comprise an agent's lens through which they view the world. Such values comprise the core of an agent's moral identity. Contemporary Real Self Views include those endorsed by Pamela Hieronymi, T.M. Scanlon, and Angela Smith. I discussed these views in chapter 3. So, I will just mention that for Hieronymi commitment-constitutive attitudes are the focus, for Scanlon judgment-sensitive attitudes are the focus, and for Smith bearing a rational connection to an agent's evaluative judgments is the focus. These views share the idea that a person is morally responsible for an action or attitude if and only if that action or attitude reveals something fundamental about who that person is and what they value.

4.2 Wolf's JoJo Case

Susan Wolf (1987) argues against the Real Self View of responsibility. She thinks the Real Self View licenses a shallow form of blameworthiness. Such a view cannot account for the diachronic impact to responsibility of poor formative circumstances. It cannot account for the way a poor upbringing that shapes one's values can impair

one's capacity for telling right from wrong. The Real Self View cannot account for what Wolf calls "moral sanity." An agent may perform actions that match her values, but her values may be the result of formative circumstances that impair her normative capacity—her capacity to tell right from wrong. To illustrate this deficiency, Wolf introduces her famous case involving an evil dictator named JoJo:

JoJo is the favorite son of Jo the First, an evil and sadistic dictator of a small, undeveloped country. Because of his father's special feelings for the boy, JoJo is given a special education and is allowed to accompany his father and observe his daily routine. In light of this treatment, it is not surprising that little JoJo takes his father as a role model and develops values very much like dad's. As an adult, he does many of the same sorts of things his father did, including sending people to prison or to death or to torture chambers on the basis of whim. He is not *coerced* to do these things, he acts according to his own desires. Moreover, these are desires he wholly *wants* to have. When he steps back and asks, "Do I really want to be this sort of person?" his answer is resoundingly "Yes," for this way of life expresses a crazy sort of power that forms part of his deepest ideal. (Wolf 1987: 53-4)

Wolf's example targets a version of the Real Self View, which she calls the Deep Self View. This view focuses on desires. An action stems from one's deep self when it is the result of first-order desires (e.g., JoJo's desire to murder and torture citizens) and when one reflectively endorses those first-order desires via second-order desires (e.g., JoJo wants to be the sort of person that wants to murder and torture citizens). Wolf thinks a pre-theoretic intuition concerning the JoJo case is that JoJo is not responsible for his actions. As Wolf explains about the case:

In light of JoJo's heritage and upbringing—both of which he was powerless to control—it is dubious at best that he should be regarded as responsible for what he does. It is unclear whether anyone with a childhood such as his could have developed into anything but the twisted and perverse sort of person that he has become. However, note that JoJo is someone whose actions are controlled by his desires and whose desires are the desires he wants to have. (Wolf 1987: 54)

For Wolf, the Deep Self View is insufficient. A subject can act in a way governed by desires representing the subject's deep self yet still be unfree and not responsible for his actions. What is missing from such views? For Wolf a necessary condition—sanity—is missing. JoJo is morally insane. This explains why he is not responsible. To overcome this deficiency Wolf adds to the Deep Self View. Not only must there be a mesh between an agent's first-order and second-order desires, but the agent must also “know the difference between right and wrong,” and JoJo fails on this score because “a person who, even on reflection, cannot see that having someone tortured because he failed to salute you is wrong plainly lacks the requisite ability” (Wolf 1987: 56). This conclusion about JoJo sheds light on cases that tested, in chapter 3 of this dissertation, whether moral ignorance ever excuses. In the cases at issue, the people are not morally responsible because they are not morally sane. As Wolf articulates:

The slaveowners of the 1850s, the Nazis of the 1930s, and many male chauvinists of our fathers' generation...their false beliefs in the moral permissibility of their actions and the false values from which these beliefs derived may have been inevitable, given the social circumstances in which they developed. If we think that the agents could not help but be mistaken about their values, we do not blame them

for the actions those values inspired. It would unduly distort ordinary linguistic practice to call the slaveowner, the Nazi, or the male chauvinist even partially or locally insane. Nonetheless, the reasons for withholding blame from them is at bottom the same as the reason for withholding it from JoJo. Like JoJo, they are, at the deepest level, unable cognitively and normatively to recognize and appreciate the world for what it is...their deepest selves are not fully *sane*. (Wolf 1987: 56-7)

Wolf adding a sanity condition to the Deep Self View is a challenge to the thesis I defended in chapter 3. Wolf thinks moral ignorance is often an excuse condition because it is often the result of normative incapacity. It is often the result of an inability to tell right from wrong, which stems from forces external to the agent shaping the agent into a person benighted toward the moral truth. So, I need to respond to Wolf's challenge.¹⁹⁷

The first way I will respond to Wolf's challenge is to look at a recent empirical study focusing on ascriptions of responsibility concerning the JoJo case. Contrary to Wolf's pre-theoretic intuition that JoJo is not morally responsible for his actions, participants judged that JoJo is responsible. Next, I will propose a way to explain ascriptions of moral responsibility, and this view will capture the main data in the study. It will also provide a theory from which I will argue that ascriptions of responsibility are justified in particular cases, which will counter the skeptical challenge. Let's start by looking at the study.¹⁹⁸

¹⁹⁷ At this point you may wonder how this relates back to the stated focus of the chapter, which is on the epistemic concerns that originally animated the skeptical challenge. After articulating and defending the relevant conditions, they will be related back to the epistemic challenge.

¹⁹⁸ Thanks to Greco for bringing to my attention the need to address the use of experimental philosophy (X-Phi) at this point in the project. The use of X-Phi seems to come out of nowhere and does not include discussion of the pros and cons of using X-Phi methodology. While I cannot go into the large literature

4.3 Testing the JoJo Intuition

David Faraci and David Shoemaker (2010) conducted an empirical study regarding Wolf's JoJo case.¹⁹⁹ They tested Wolf's pre-theoretic intuition that JoJo is not morally responsible. Participants in the study were presented with three scenarios. The first scenario is the control scenario. It features JoJo senior (Jo the First) without the details of a deprived upbringing.

JoJo1 Scenario: JoJo is an evil and sadistic dictator of a small, undeveloped country. JoJo does many things as a dictator, including sending people to prison or to death or to torture chambers on the basis of whim. He is not coerced to do these things. When he steps back and asks, "Do I really want to be this sort of person?" his answer is resoundingly "Yes," for this way of life reflects his deepest values and ideals. (Faraci & Shoemaker 2010: 325)

The second case is Wolf's JoJo case slightly modified to avoid phrases suggesting an interpretation of the judgment of responsibility regarding JoJo.

JoJo2 Scenario: JoJo is the favorite son of Jo the First, an evil and sadistic dictator of a small, undeveloped country, entirely cut off from the outside world. Because of his father's special feelings for the boy, JoJo is given a special education

on the pros and cons of using X-Phi in philosophical inquiry, I can point to a helpful resource for wading into those waters. For good discussion of the pros and cons of X-Phi, including specific topics X-Phi addresses, see Machery and O'Neill (2014). I use X-Phi in this project because of the form of responsibility at issue that seeks to do justice to our ordinary social practices of ascribing responsibility. That form of responsibility should be consistent with the way people actually ascribe responsibility. In chapter 4 we focus on the epistemic status of judgments of responsibility. So, given the nature responsibility at issue (i.e., a desert-based conception that is broadly Strawsonian) and the epistemic status of responsibility judgments at issue, it is relevant to look at how people actually form judgments of responsibility in relation to how attributionist views handle such cases.

¹⁹⁹ It is worth mentioning what sort of evidential value I take these X-Phi results to have. I take it that the experimental results provide *prima facie* evidence that Wolf's intuition about the JoJo case is not supported. Recalling the dialectic, Wolf makes a move against the Real Self View by assuming an intuition holds. Showing that the intuition does not hold provides *prima facie* evidence against the need for a Sanity Condition. Could another study defeat the first study and restore the need for a Sanity Condition (i.e., a modified Deep Self View)? I think, yes. But there have not been many empirical studies done regarding the epistemic conditions of moral responsibility. I can only go on what is available, and it provides *prima facie* evidence that the intuition is not reflected in our ordinary practices of ascribing responsibility to agents in such cases.

and is allowed to accompany his father and observe his daily routine. In light of this treatment, little JoJo takes his father as a role model and develops values very much like Dad's. As an adult, JoJo does many of the same sorts of things his father did, including sending people to prison or to death or to torture chambers on the basis of whim. He is not coerced to do these things. When he steps back and asks, "Do I really want to be this sort of person?" his answer is resoundingly "Yes," for this way of life reflects his deepest values and ideals. (Faraci & Shoemaker 2010: 325)

The third case is Wolf's JoJo case plus exposure to alternative ways of dealing with people as a leader. JoJo lives abroad for a year.

JoJo3 Scenario: JoJo is the favorite son of Jo the First, an evil and sadistic dictator of a small, undeveloped country, entirely cut off from the outside world. Because of his father's special feelings for the boy, JoJo is given a special education and is allowed to accompany his father and observe his daily routine. In light of this treatment, little JoJo takes his father as a role model and develops values very much like Dad's. When he turns 21, JoJo is sent to live in a developed country for a year, and there he becomes aware that other leaders treat their subjects with respect and goodwill because they value the lives and well-being of their subjects. Nevertheless, when he returns to lead his country, JoJo does the same sorts of things his father did, including sending people to prison or to death or to torture chambers on the basis of whim. He is not coerced to do these things. When he steps back and asks, "Do I really want to be this sort of person?" his answer is resoundingly "Yes," for this way of life reflects his deepest values and ideals. (Faraci & Shoemaker 2010: 325-6)

For each scenario, participants were asked to, "circle the number that best represents the degree to which you think JoJo is blameworthy for his actions (sending people to prison/death/torture chambers),' with '1' being not at all blameworthy, '4' being somewhat blameworthy, and '7' being completely blameworthy." The results were as follows:

The Results: On average, subjects did judge JoJo1 (the control) quite sternly: the mean was 5.8, with exactly half of the respondents assigning a solid 7 to him. In the JoJo2 case, however—essentially Wolf’s original JoJo case—the mean was 4.77. This means that, while subjects did deem JoJo2 to be less blameworthy than JoJo1, they still judged him to be more than somewhat blameworthy. (The mode score was actually a 6, and the median reply was a 5). What of the responses to JoJo3, the one seemingly aware of moral alternatives? The responses here were unexpected: the mean was 4.93, but the mode and median scores were both 5, trends suggesting that subjects found that someone with JoJo’s background is less blameworthy than Jo the First, even if he has been exposed to moral alternatives. The differences in mean between JoJo2 and JoJo3, however, were not statistically significant. (Faraci & Shoemaker 2010: 326)

There are two main upshots of these results. The first upshot is that Wolf’s pre-theoretic intuition concerning JoJo not being morally responsible is cast into doubt. Participants did not excuse JoJo2 from moral responsibility given the way his values were inculcated through his upbringing. These results confirmed the Deep Self View, namely that what matters most for pre-theoretic judgments of responsibility is whether the agent reflectively endorses the values issuing in the actions in question.²⁰⁰

²⁰⁰ A possible worry is that pre-theoretic intuitions are unreliable. Such intuitions about responsibility might be unreliable because they are influenced by false views of the self or soul that stem from religious modes of thinking. And, as history shows, what is accepted as intuitively true at one time is sometimes shown to be radically false by future generations. However, this worry relies on contentious assumptions. It assumes that religious views about the self and soul are false, and it assumes that religious modes of thinking are not a natural phenomenon. First, as McCauley (2011) explains, religion is natural but science is not. Religion is natural in that it is intuitive, where intuition for McCauley is, “the principle manifestation of natural cognition in our mental lives” (14). So even if religious thinking impacts intuitions about responsibility this does not show that they are misguided or contrary to what is natural and intuitive. Second, there is a debate over the nature of the self and soul. It is not the case that materialism and naturalism concerning the self and soul are true and dualism and supernaturalism concerning the self and soul are false. There is much debate to be had on this point. For instance, consider Alvin Plantinga’s evolutionary argument against naturalism. As Plantinga argues, “naturalism, insofar as it implies materialism about human beings, has no room for the essential features of our mental life, including in particular *belief*” (Plantinga 2009: 1). Further, even if intuitions are influenced by religious modes of thinking this may only reflect that religious thinking (about the soul and responsibility) is a natural phenomenon. Theoretical reflection in the hands of ethicists and philosophers is not necessarily more accurate. The claim that philosophers’ intuitions are more reliable

They fail to affirm Wolf's claim that the Deep Self View needs to be augmented with a sanity condition.

The second upshot of the results is confirmation of the No Excuse thesis that I defended in chapter 3 of this dissertation. The results suggest that, according to the intuitions of those surveyed, moral ignorance is not exculpatory. In fact, Faraci and Shoemaker indicate that *no one* they surveyed considered JoJo2 not blameworthy to some degree. No one judged JoJo2 excused from responsibility.

An additional upshot of the results is the lack of a statistically significant difference between case 2 and case 3. This surprised Faraci and Shoemaker. They hypothesized that isolation from moral alternatives and the accompanying moral ignorance would create a significant difference. It did not. Exposure to moral alternatives in the JoJo3 case failed to elicit a major difference. Participants failed to judge JoJo3 much more blameworthy than JoJo2 because it was already the case that JoJo should have known better. In interpreting these results, Faraci and Shoemaker still attempt to make moral ignorance *the* key factor.

Faraci and Shoemaker claim that participants judged JoJo2 and JoJo3 equally ignorant yet both JoJos *should have known better*. On their view, it is reasonable to expect the JoJos to have rectified their ignorance. Faraci and Shoemaker think that JoJo2 and JoJo3 should have comprehended the wrongness of their actions.²⁰¹ According to Faraci and Shoemaker, participants judge that JoJo2 should have known

because philosophers are experts has been challenged. See Buckwalter (2016). I thank Aaron Zimmerman for urging me to think about this worry.

²⁰¹ Faraci and Shoemaker might think this follows from the JoJos knowledge of the suffering of the people they are torturing plus the JoJos introspective/counterfactual knowledge that they would not want to be treated in the way they are treating the people they are torturing or causing to suffer.

better and the year abroad that exposed JoJo3 to other moral viewpoints was redundant. This is why it only made JoJo3 slightly more blameworthy than JoJo2. But why think that participants are reading into the cases culpable moral ignorance such that the JoJos should have known better? I will cast doubt on this interpretation of the data. But, first it is good to set aside another way a Deep Self theorist might go wrong or take on problematic theoretical resources given the results of the study.

4.4 A Scalar Deep Self View

Faraci and Shoemaker suggest ways of explaining the exact pattern of the data in terms friendly to a Deep Self view. Yet, following their suggestions in this regard is unnecessary and theoretically problematic. They also unnecessarily push an interpretation of the data.²⁰² In contrast, I think the Deep Self theorist should stick to her guns. She should insist that the main details of the study provide *prima facie* support for the Deep Self view. No one let any of the JoJos off the hook in light of historical formative factors or exposure to moral alternatives. The Deep Self theorist should remain agnostic about what explains the decrease in blameworthiness between JoJo1 and JoJo2. The point that bears emphasizing is that, “while subjects did deem JoJo2 to be less blameworthy than JoJo1, they still judged him to be more than somewhat blameworthy. (The mode score was actually a 6, and the median reply was a 5)” (Faraci & Shoemaker 2010: 326). Despite the importance of this central point, one might wonder whether the Deep Self theory can exactly fit the data?

One way to press the Deep Self view to fit the exact contours of the data is to posit a Deep Self theory that is scalar or comes in degrees. On a scalar characterization

²⁰² I address this interpretation in the next section.

of the Deep Self view the degree to which an agent is blameworthy for an action corresponds to the degree to which the action is attributable to the agent's deep self. In interpreting the results of their study, Faraci and Shoemaker make a recommendation along these lines.

[W]e cannot forget that JoJo2 was deemed to be quite blameworthy (indeed, no one thought he was not blameworthy at all). So it is not at all clear that his being (slightly) less blameworthy than JoJo1 warrants judging the DSV [Deep Self View] as it stands to be insufficient. Indeed, one might think that he is as blameworthy as he is precisely because he meets the conditions of the DSV, just not as fully as does JoJo1. Perhaps subjects think that JoJo2's actions are not as attributable to his deep self as JoJo1's, given the former's upbringing, or perhaps subjects think that the depth of JoJo2's deep self is somewhat limited. But in any event, there may be scalar resources within the DSV itself that could explain both the significant degrees of blameworthiness attached to both JoJos as well as the disparity between them. (Faraci & Shoemaker 2010: 327)

This scalar Deep Self view must account for the fact that none of the JoJos are excused from blameworthiness for their actions, yet there is a difference in the degrees of blameworthiness of the JoJos. Let me start with the former.

All the JoJos are judged blameworthy because their actions are all attributable to their deep selves in a way that passes a threshold for outright attributability in the binary sense. What common thread is there between all the JoJo cases? Here is the crucial bit of verbiage that all the cases have in common, "When [JoJo] steps back and asks, 'Do I really want to be this sort of person?' his answer is resoundingly 'Yes,' for this way of life reflects his deepest values and ideals" (Faraci & Shoemaker 2010: 325-6). JoJo desires to be the sort of person that desires to rule his people through intimidation and harming innocent people. This suggests that if an action is

attributable to an agent's deep self because the person acts on desires they reflectively endorse as expressing values they hold—normative ideals they endorse as representative of who they want to be and how they want to treat people—then that action passes an attributability threshold. In an outright sense, the action is attributable to the agent, as it expresses the agent's deep self. The agent is blameworthy for the action.²⁰³

There is variation among the degrees to which the JoJos are blameworthy. Factors mitigate the degree to which they are blameworthy, but those factors do not excuse the JoJos from being blameworthy. In the JoJo2 case, poor formative circumstances dampen blameworthiness. Compared to the JoJo1 case, in the JoJo2 case attributability might be diverted away from JoJo's deep self. Some attributability for the wrongful actions is directed toward the conditions that shaped JoJo's values, namely his upbringing lacking in a moral education.²⁰⁴ The JoJo3 case adds that JoJo

²⁰³ Given that participants in the study are given that information concerning JoJos reflective endorsement of his actions as aligning with his sincerest values, the participants are given evidence that justifies their ascriptions of responsibility and blameworthiness. Contrary to the responsibility skeptic, this is an example of ascribers of responsibility having justified judgments of responsibility concerning a specific case. Therefore, it is not the case that there are not any justified ascriptions of responsibility concerning specific cases. I thank Hanser for raising the worry that this is question begging. Rosen's skeptical challenge works against the background of an accountability view of responsibility that endorses the Reasonable Expectations claim as a necessary condition. Isn't it begging the question against the skeptic to assume a different view of responsibility is true and then use that account, which would be denied as true by the skeptic, to counter the skeptical challenge? In response, I need not argue in a question begging way. Rather, I can argue that skeptical worries are less pressing on a Deep Self view because the factors that bear on responsibility seem in principle ones that we have epistemic access to in comparison to the worries that arise out of claiming culpability must trace to acts of clear-eyed akrasia. The attributability theorist focuses on whether we can ascribe mental states to others (i.e. whether we can ascribe judgments and values to them in light of their self-reports or details about their mental life). Continuing to press the skeptical challenge against the attributability view yields a very deep form of skepticism. That the skeptical worry shifts to such a general worry about the possibility of ascribing mental states to people counts in favor of attributability as compared to accountability. It is a mark in favor of the Deep Self view that it pushes the skeptic into these deep waters and away from a skepticism specific to responsibility and the need to locate acts of knowing wrongdoing in the history of the origin of the mental states. Thanks to Hanser for suggesting a response somewhat along these lines.

was exposed to alternative moral conceptions regarding permissible behavior toward persons under one's authority. This is a slight blameworthiness amplifying factor, but it does not generate a statistically significant difference between the JoJo3 and JoJo2 cases. As a result, the factor that reduces blameworthiness, namely JoJo's formative circumstances, does not reduce blameworthiness beyond the attributability threshold, as that threshold is surpassed by the action being attributable to the agent's deep self.

How can a deep self theorist claim that differences in upbringing affect differences in blameworthiness when it does not make one subject identify with her vicious motivations or values more than the other?²⁰⁵ One way of explaining why this holds is in terms of intellectual difficulty. Moral ignorance from a morally deficient childhood decreases blameworthiness, as it generates intellectual difficulty in overcoming that ignorance. This is the line Faraci and Shoemaker take. In the next section I argue against this line. Instead, I think the worry posed in the form of a question that started this paragraph is onto something. Allowing JoJo2's upbringing to mitigate his responsibility cuts against the core idea of the Deep Self view, namely that how JoJo2 became the person he is does not mitigate blameworthiness. JoJo2's morally deficient upbringing helps explain why he has the deep self that he has, but it does not decrease his blameworthiness to a significant degree. JoJo2 is outright responsible for his actions. That there is a slight decrease in blameworthiness between JoJo1 and JoJo2 should not worry the Deep Self theorist or cause her to take on

²⁰⁵ Thanks to Aaron Zimmerman for raising this question and worry.

unnecessary theoretical baggage that is inconsistent with the main idea behind the Deep Self view.²⁰⁶

4.5 The Difficulty with Difficulty

In this section, I will argue against Faraci and Shoemaker's proposed explanation for the variation in degrees of responsibility and blameworthiness in the JoJo cases. They postulate that epistemic difficulty in ascertaining the truth is an inversely mitigating factor: The harder it is for a subject to attain moral truth, the less blameworthy the subject is; the easier it is to secure truth, the more blameworthy the subject is.

In presenting Faraci and Shoemaker's suggestion concerning difficulty it is important to put in place their commitment to moral ignorance playing a factor in the assessment of responsibility, especially between the JoJo2 and JoJo3 cases. They postulate that, "what may mitigate the full-blown excusal of JoJo2 is the belief that his ignorance itself is rather culpable, that even though he did not in fact know better (and his ignorance is deep-seated), he *should have*, where this means there were plenty of opportunities for him to infer that expressions of ill will were wrong, if only he had paid closer attention or been sufficiently sensitive" (Faraci & Shoemaker 2010: 328). Earlier, I suggested that this perspective is implausible. Now we can see why. As Faraci and Shoemaker continue, "Indeed, subjects may well have thought that JoJo3's exposure to the moral alternative was *redundant*: he had already had sufficient exposure to render his ignorance culpable. Nevertheless, because of the deep-seated nature of the ignorance, subjects could be cutting both JoJo2 and JoJo3 roughly equal slack, given the degree of difficulty attached to their actually succeeding in identifying

²⁰⁶ Thank you to Matthew Hanser for helpful suggestions along these lines.

and eliminating their zones of ignorance” (Faraci & Shoemaker 2010: 329). Against this line, exposure to moral alternatives is not redundant from the perspective of grounding blameworthiness. It is neither the case that, “JoJo3’s exposure to the moral alternative was redundant,” nor that, “he had already had sufficient exposure to render his ignorance culpable.” Let me explain.

JoJo3’s exposure to an alternative way to treat people and give them the due regard owed to them as persons was not redundant for a couple of reasons. First, that exposure was not a redundant form of evidence concerning the immorality of his evaluative stance toward the treatment of persons. Prior to his year abroad, JoJo3 lived in a moral echo chamber. JoJo lacked robust exposure to moral alternatives because of the nature of his father’s position and power. As the vignette explains, JoJo’s father was, “an evil sadistic dictator of a small, undeveloped country, entirely cut off from the outside world.” Evil and sadistic dictators rule through fear and intimidation. In such environments citizens do not question their leader without significant risk to life and property.²⁰⁷ Public disapproval of Jo the First would likely be minimal. Little moral dissent would likely reach JoJo3’s ears, and what dissent did reach his ears he learned to handle as his father did. As the case explains, JoJo grew up to do, “the same sorts of things his father did, including sending people to prison or to death or to torture chambers on the basis of whim.” When JoJo3 spent a year abroad he was confronted with a new perspective. Such exposure was not redundant. He was exposed to new counter-evidence to his attitudes and actions. However, exposure to such moral alternatives did not translate into changed values and

²⁰⁷ It is also likely the case that out of fear of retribution, imprisonment, or death, Jo the First’s inner circle did not publicly question or dissent from how he governed his people.

behavior. JoJo returned home from his time abroad and proceeded to run his country in an evil and sadistic manner like his father did. JoJo also endorsed this manner of governing his country as reflecting the sort of person he wants to be and reflecting his “deepest values and ideals.” These considerations go some way to explaining why participants reading the JoJo vignettes did not drastically increase the degree to which JoJo3 was blameworthy as compared to JoJo2.

Second, JoJo3’s exposure to moral alternatives was not redundant because JoJo2’s moral ignorance was not inculpatory in the way suggested by Faraci and Shoemaker. It is unreasonable to claim that JoJo2 should have known better, “where this means there were plenty of opportunities for him to infer that expressions of ill will were wrong, if only he had paid closer attention or been sufficiently sensitive” (Faraci & Shoemaker 2010: 329). Given how difficult it would have been for JoJo2 to rectify his ignorance it is unclear why it is still true, as Faraci and Shoemaker claim, that JoJo2 should have known better. Wouldn’t it be unfair to expect him to have taken advantage of the opportunities for moral improvement by just being more attentive to the moral factors in play in his actions and their consequences? To see that it would be unfair we need look no further than Faraci and Shoemaker’s adept articulation of why JoJo2 and JoJo3’s moral ignorance was deep-seated:

JoJo2 and JoJo3 may be thought to be closely aligned in this respect precisely because of the particularly insidious type of ignorance they have in common, a type we believe is different in *kind* from the moral ignorance people usually experience....Now when (either) JoJo inflicts his injuries on the peasants he actually intends to do so: this is the only way such peons will learn, he thinks, or perhaps it expresses a power suitable to his station. We might thus say that JoJo intentionally expresses *ill* will to them. What is he ignorant of, then? He does not

know that expressions of *ill will* are wrong. That is, his is a more fundamental ignorance than ours: he is unaware of which moral properties supervene on which act-types.²⁰⁸ Indeed, he thinks that what he is doing is morally right, that he is following squarely in the footsteps of his “admirable” and “morally good” father. This is a kind of ignorance, though, that mere exposure to the relevant moral alternative—the basic demand that expressions of ill will are wrong—may not be sufficient to displace. (Faraci & Shoemaker 2010: 327-8)²⁰⁹

Given Faraci and Shoemaker’s analysis of JoJo2 and JoJo3’s deep-seated ignorance I fail to see how they can embrace a Reasonable Expectation claim regarding the JoJos, such that the JoJos *should have known better*. If JoJo could not know that moral properties such as wrongness supervene on act-types involving expressions of ill will, then how could Faraci and Shoemaker claim, as they do, that for JoJo, “there were plenty of opportunities for him to infer that expressions of ill will were wrong, if only he had paid closer attention or been sufficiently sensitive” (Faraci & Shoemaker 2010: 329)? This is the difficulty that focusing on epistemic difficulty poses for their interpretation of the results. It forces Faraci and Shoemaker down the path toward endorsing Wolf’s stance that JoJo lacked normative competence and was morally

²⁰⁸ By contrast, our ignorance involves, “our lacking knowledge that some particular action is generally construed as an expression of ill will, where such expressions are what we typically believe make the action wrong. In other words, we are aware of which moral properties supervene on which act-types, but we may lack knowledge of which act-tokens fall under those act-types” (Faraci & Shoemaker 2010: 327-8).

²⁰⁹ Recent real-life examples of such tyrants include Kim Jong-il and Kim Jong-un of North Korea and Saddam Hussain and Uday Hussein of Iraq. Saddam regularly tortured and killed his people. As John F. Burns (2003) of the *New York Times* explains, “[Saddam] used murder on a mass scale as a political tool, to create the fear that discouraged challenges to his power.” As Faraci and Shoemaker explain, such a leader often murders and tortures people because, “it expresses a power suitable to his station.” Such a megalomaniac may not consciously think about morality, but he does believe that what he is doing is morally fine given his power and station. He knowingly and intentionally acts with ill will toward his people, but he does not recognize such acts as morally wrong. For instance, Saddam Hussein in his public trial remained defiant and pled not guilty to crimes against humanity despite the mountains of evidence of his systematic brutality toward his people throughout his reign. Thanks to Zimmerman for encouraging me to consider real life cases.

insane.²¹⁰ Faraci and Shoemaker's attempt to give an ignorance-focused interpretation of the data is part of avoiding Wolf's claim that JoJo is insane. The problem with Wolf's claim is that such incapacity should prove exculpatory. But it did not in the minds of the respondents surveyed in this study. No participants in the study judged JoJo not blameworthy to any degree. Yet if JoJo2 and JoJo3 suffer from deep ignorance of the sort Faraci and Shoemaker postulate, and this explains why exposure to moral alternatives fails to make a difference between the cases, this lends credence to the idea that the JoJos capacity to tell right from wrong is positively impaired. Unawareness of which moral properties supervene on which act-types results in thinking that what one is doing is morally permissible when it is not. And exposure to questions concerning the morality of such behavior will not dislodge this ignorance. If so, such normative incapacity should prove exculpatory, especially given the difficulty that rectifying it poses the agent. But, it does not.

Notice that New Attributionism does not fall victim to Faraci and Shoemaker's interpretation, which lands Faraci and Shoemaker back at the front door of Wolf's interpretation of the JoJo case. The Deep Self View need not embrace a Reasonable Expectation claim as a requirement for establishing the culpability of ignorance, though satisfying such a claim may prove sufficient. Requiring culpable ignorance to trace to what it is reasonable to expect regarding rectifying that ignorance, where this can be impacted by how intellectually difficult it is to overcome that ignorance, brings into focus questions of normative competence. By contrast, for a deep self theorist, even if it is unreasonable to expect JoJo to have rectified his ignorance, as it is of the

²¹⁰ This holds even though they do not actually embrace Wolf's stance. They are on the path toward embracing it whether they recognize it or not.

deep-seated sort, JoJo is still morally responsible and blameworthy for his actions that reflect his values and a morally objectionable stance toward how it is permissible to treat people—a stance that impairs relationships as it violates the basic demand for due regard owed to people.

4.6 Justified Responsibility Ascriptions

Here I return to directly responding to the skeptical challenge concerning moral responsibility. I will argue that ascriptions of responsibility are justified in some specific cases. Judgments of responsibility and blameworthiness are justified on the basis of evidence. Typical sources of evidence and justification include perception, memory, testimony, and induction. In specific cases, justified ascriptions of responsible are based on testimonial reports concerning what an agent has done, how they think about what they have done, or by conversation with the agent, where possible.

To show that responsibility ascriptions are justified concerning specific cases I will start by breaking down the JoJo2 case discussed earlier. The details of the case can be broken down as follows:

Setting: JoJo is the favorite son of Jo the First, an evil and sadistic dictator of a small, undeveloped country, entirely cut off from the outside world.

Value Inculcation: Because of his father's special feelings for the boy, JoJo is given a special education and is allowed to accompany his father and observe his daily routine. In light of this treatment, little JoJo takes his father as a role model and develops values very much like Dad's.

Actions: As an adult, JoJo does many of the same sorts of things his father did, including sending people to prison or to death or to torture chambers on the basis of whim.

Free Action: He is not coerced to do these things.

Deep Self: When he steps back and asks, “Do I really want to be this sort of person?” his answer is resoundingly “Yes,” for this way of life reflects his deepest values and ideals. (Faraci & Shoemaker 2010: 325)

In this specific case, the ascriber of responsibility is given evidence that the freely performed actions by JoJo, which express morally objectionable judgments about how it is permissible to treat people, are attributable to JoJo’s deep self. The ascriber gets a glimpse into how JoJo thinks about the actions in relation to what he values. JoJo affirms that how he treats people under his authority represents who he is and wants to be. This “deep self” dialogue provides the ascriber of responsibility with evidence that justifies the verdict of JoJo being blameworthy for his wrongful actions done from moral ignorance.

But, what if the “deep self” details of the case were left out? How could an ascriber reach a justified ascription of responsibility? In answering this question, it helps to remember the nature of the skeptical challenge. The challenge claims that there are *not any* justified judgments of responsibility about particular cases involving wrongful acts done from ignorance. I just showed there is such a case. This is enough to counter the skeptical challenge. However, I admit that, in many cases, the details of the cases are under-described. When this occurs suspension of judgment pending further investigation and discovery of fact concerning what the action reveals about what the agent values and reflectively endorses is in order. However, in the cases in discussion in the literature, and in many everyday cases, ascriptions of responsibility

are justified. Here are details from cases discussed in chapter 3. These cases provide evidence justifying the judgment that the agents are blameworthy:

Enlightened Mr. Potter. He comes to the view that it is better for the economy overall if companies operate to maximize profits, and that economic progress always has casualties. He believes that by pursuing what is in the best interests of his shareholders, he is playing a valuable role in the economy.

Gang Avenger. A gang member kills a member of an opposing gang who killed his friend (and fellow gang member). He believes he is doing the right thing.

Mob Enforcer. A mafia “family” member kills a shop owner who refuses to pay a weekly extortion fee. (The mafia demands a weekly payment for “protection”; otherwise they threaten violence.) He believes he is doing the right thing.

Sympathetic Slaveholder. 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.).

Given the examples above, the ascriber of responsibility can determine that the actions of the agents reveal a certain stretch of the agent’s mind. The actions and attitudes toward persons are attributable to the agent. They reveal that the agent has settled the question of whether to perform the acts in question. They reveal that the agent’s attitude toward others damages the relationship that other people can have with them. From the New Attributionist perspective, and the Deep Self View in particular, there is evidence that the actions reflect objectionable judgments held by the agents concerning how it is permissible to treat people, attitudes that are reflectively endorsed or normatively-relevant and inculcating (i.e., such as knowledge of the suffering of others).

How does this connect to the arguments of chapter 2? In chapter 2, I modified the culpable ignorance principle a couple of times. The main modification of the principle involved allowing for the non-akratic exercise of vices to count as sufficient to make ignorance culpable. How are ascribers of responsibility provided evidence of the exercise of vices? Let's consider the original Potter case articulated by FitzPatrick (2008). Unlike JoJo, Potter was not isolated from opportunities to consider alternative moral views. Potter had the capacity to critically consider his own values and had the opportunities to do so. But Potter did not critically engage with the opportunities in his environment due to the exercise of vices. FitzPatrick (2008) mentions that Potter did not capitalize on the opportunities to overcome his ignorance because of, "a cavalier dismissal of the relevant considerations, viewpoints, and arguments—an easy labeling and dismissing of them as 'socialist' or 'liberal,' without honest reflection open to the possibility that they may point to blind spots in his own views" (603-4). Ascribers of responsibility in this case can assign culpability for his ignorance and blameworthiness for his bad actions because his "cavalier dismissal" of alternative views. Potter's perseverance in ignorance strongly suggests that his ignorance is not the result of one-off behaviors that are out of character for Potter. Rather, to maintain ignorance in the face of such opportunities, given that one has the capacity to utilize that readily available evidence, provides the ascriber with evidence that Potter's ignorance is the result of stable dispositions to ignore counter-evidence. Potter's lack of reflection and gathering of evidence relevant to the moral permissibility of his actions is the result of "voluntary exercises of vices such as overconfidence, arrogance, dismissiveness, laziness, dogmaticism, incuriosity, self-indulgence, contempt, and so on" (2008: 605). As this illustrates, responsibility

ascriptions can be justified based on the evidence provided by the details of a particular case.

I conclude this section with forestalling a worry. In addressing this worry I will connect what I have just argued back to the argument in chapter 3 that moral ignorance is never exculpatory. The worry is that if moral ignorance is never exculpatory, then all judgments of blameworthiness concerning wrongful actions done from moral ignorance will be sufficiently justified given that they are conjoined with knowledge that the person (intentionally) did the thing for which she was blamed and had knowledge of all the relevant non-moral (or descriptive) facts of the case. Evidence of moral ignorance will be sufficient to ground such justified judgments. There is no need to verify that the actions and attitudes are connected to the agent's deep self in order to form justified judgments of responsibility.

In responding to this worry it is important to remember that the argument in chapter 3 that moral ignorance is never exculpatory is a metaphysical argument. It would be unreasonable to require ascribers of responsibility to comprehend and endorse metaphysical arguments concerning the nature of culpability regarding moral ignorance to form justified judgments of blameworthiness. For evidence of moral ignorance, when conjoined with evidence that the agent intentionally committed the action in full knowledge of all relevant non-moral facts, to function alone as sufficient evidence for justifying ascriptions of responsibility it would need to do so against the backdrop of an argument to the effect that moral ignorance is never exculpatory. Such an argument (or its conclusion) would function as an enabling condition allowing recognition of moral ignorance to provide justification for ascriptions of

responsibility. Otherwise, evidence of moral ignorance may be an unreliable guide to ascriptions of blameworthiness. If moral ignorance is sometimes an excuse, the presence of moral ignorance alone will not reveal what conditions must be satisfied for that ignorance to inculcate as opposed to exculpate. It is also reasonable to think that agents do not possess a default entitlement to an argument to the effect that moral ignorance is never exculpatory, such that, against the backdrop of this entitlement, claims of moral ignorance provide sufficient evidence justifying ascriptions of blameworthiness.²¹¹ Thus, even if from a metaphysical perspective, given what I have argued in chapter 3, moral ignorance is never exculpatory, it *is* the case that, from an epistemic perspective, justified ascriptions of responsibility are based on sufficient evidence when they connect actions and attitudes with what those actions and attitudes reveal concerning what the agent thinks and values. This often will involve key details in the case drawing the connection to the agent's deep self or making available a reasonable inference from the details of the case to the conclusion that there is such a connection. Simply reading that the agent performed the action in moral ignorance, perhaps because the agent possessed a false belief that the action was morally permissible, is insufficient to provide the ascriber of responsibility with a justified ascription of responsibility.²¹²

²¹¹ Plausibly, this is because of the nature of moral responsibility as embodied in social practices. Unlike warranted ascriptions of perceptual knowledge, warranted ascriptions of responsibility identify agents as fair targets of the reactive attitudes. It is less plausible that the reactive attitudes and ascriptions of responsibility generally are warranted based on entitlement to certain hinge propositions, such as the falsity of skeptical hypotheses about responsibility. This contrasts with what some epistemologists have argued regarding ascriptions of perceptual knowledge, as they are licensed against the backdrop of certain hinge propositions, such as the falsity of skeptical hypotheses about the external world. For such a move regarding rational claims to warrant regarding perceptual knowledge, see Crispin Wright (2014).

²¹² In addition, if the details of the case are under-described in this regard or a reasonable inference from the details of the case to the connection to the agent's evaluative judgments cannot be drawn, then withholding judgment is the justified attitude for the ascriber to take. This can coincide with it being

Lastly, it is worth mentioning that the argument of chapter 3 in defense of the No Excuse thesis is not unconnected to what I have argued in this chapter. The sort of evidence ascribers of responsibility must secure in forming justified judgments of responsibility track the metaphysical argument of chapter 3. Even though the argument of chapter 3 was strict liability in spirit it was not a strict liability argument in the sense that liability (i.e., culpability) is incurred in cases of moral ignorance irrespective of the *mens rea* (i.e., the guilty mind) of the individuals. Rather, I argued for the strict liability in spirit claim based on its connection to the *mens rea* of the individual. I argued that actions done from moral ignorance can directly ground responsibility because they embody a morally objectionable stance toward the treatment of persons. That stance toward persons reflects a mind not appropriately sensitive to what is owed to others. It reflects objectionable evaluative judgments. So, even if discovery of moral ignorance alone, when conjoined with intentional performance of the act in full knowledge of its descriptive properties, is not sufficient to *justify* ascriptions of responsibility, the same considerations that justify ascriptions of responsibility also connect to what metaphysically grounds the agent as being responsible. Thus, the main argument of chapter 3 is connected to what I have argued so far in this chapter.

4.7 Summary of Key Results

As part of countering the skeptical challenge to there ever being justified responsibility judgments I have articulated metaphysical conditions on blameworthiness. It is

the case that, metaphysically-speaking, the agent's moral ignorance will not excuse her from responsibility.

helpful to take stock of these conditions prior to the last topic of this dissertation, namely moral luck. The Blameworthiness principle is:

Blameworthiness: S is morally blameworthy for performing wrongful act X if and only if (1) S freely and intentionally performs X, (2) S possesses full knowledge of the factual (descriptive) properties of X, or S is culpably ignorant of such non-moral facts relevant to X-ing, and (3) S expresses a morally objectionable quality of will or quality of judgment in X-ing,²¹³ or S performing X aligns with S's evaluative moral commitments.²¹⁴

The first disjunct in (3) captures the arguments of chapter 3 of this dissertation, and the second disjunct in (3) captures earlier considerations in this chapter. However, clause (2) is incomplete. I need to further qualify necessary condition (2) in order to account for the arguments in chapters 2 and 3. Instead of making a monstrosity of the Blameworthiness principle I will add a separate principle specifying conditions individually sufficient for making non-moral ignorance culpable. The arguments of chapter 3 simplify things regarding moral ignorance. We need not consider whether moral ignorance is the result of the culpable exercise of vices or an act of clear-eyed akrasia. When S performs an act from moral ignorance that person satisfies requirement (3). Performing an act from *pure* moral ignorance, where the moral ignorance is accompanied by knowledge of the non-moral facts, expresses a morally objectionable quality of judgment concerning how it is permissible to treat persons. Acts done from such moral ignorance can directly ground blameworthiness. We need

²¹³ Whether such a morally objectionable quality of judgment expressed by the action is in or out of character for the agent. Attributability to the deep self (i.e., S's evaluative commitments) is sufficient but not necessary for blameworthiness. Agents can be morally responsible for out of character acts.

²¹⁴ That is, judgments she endorses or would endorse on reflection concerning the good and the true.

not keep tracing backward to discover whether the moral ignorance is culpable. However, cases of *impure* moral ignorance involve ignorance of non-moral facts. As such, we still need to ask whether the non-moral ignorance is blameless or culpable. Here is the principle that details requirement (2) for blameworthiness.

Ignorance: S is morally blameworthy for performing wrongful act X only if (2) S possesses factual knowledge relevant to X-ing, or S is culpably ignorant of non-moral facts relevant to X-ing, where (2a) S's ignorance of non-moral facts is culpable *if* S's ignorance stems from either: an act of clear-eyed akrasia via beliefs held consciously or unconsciously that directly rationalize conformance with one's epistemic obligations, or a decision to not act in accordance with one's belief that one has sufficient yet non-decisive reason to perform acts that would rectify one's ignorance, or the non-akratic exercise of vices, or not acting in accordance with a justified true belief that is consciously or unconsciously held that indirectly rationalizes conformance with one's epistemic obligations.

The Ignorance principle is an elaboration of the Culpable Ignorance** principle.²¹⁵ As such, it captures the considerations of chapter 2. The Blameworthiness and Ignorance principles secure important considerations regarding the epistemic condition for moral responsibility.

As stated in the Introduction to this project, the epistemic condition is often left under-developed or totally ignored in the moral responsibility literature. For instance, consider this gloss on Blameworthiness by Ishtiyaque Haji (2013: 267):

²¹⁵ Culpable Ignorance** principle is found in chapter 2.

Blameworthiness: S is morally blameworthy for doing x if and only if (1B) S does x freely, (2B) S knows that x is morally wrong, (3B) S expresses a deficient (or a morally objectionable) quality of will in x-ing...

Notice that necessary condition (2B) specifies that S is not morally ignorant. But, what happens when S *is* morally ignorant? Is it possible that moral ignorance is an excuse condition rendering S not blameworthy for X-ing? What about factual ignorance? When is S's factual ignorance culpable such that S is blameworthy for performing their act because they are culpable for the ignorance from which they acted? Relatedly, when is factual ignorance blameless? These epistemic considerations are relevant to determining whether S is blameworthy for doing X. It is true that most theorists focus on free will in discussing moral responsibility, but the neglect of the epistemic condition is problematic. A moral responsibility theorist focusing on free will might respond, "that's just not what I'm up to here." But given that ignorance is an excuse condition that can exculpate under certain conditions, the moral responsibility theorist cannot ignore addressing the epistemic condition. Otherwise, a complete account of moral responsibility is never developed, as an almost universally acknowledged necessary condition for responsibility is not given the attention it deserves. Hopefully, this project has contributed to rectifying this imbalance.

4.8 The Problem of Moral Luck

In the final sections of this dissertation, I will show how the key results of this project afford a strategy for addressing a problem.²¹⁶ To further situate my project and show its utility in addressing the literature on moral responsibility I turn to a challenge to

²¹⁶ To clarify, I do not claim that these remarks establish this strategy. Effectively arguing for the strategy would require its own larger project.

whether moral responsibility judgments are true or false. Is what a person is morally responsible for sensitive to factors over which the person did not exercise control?²¹⁷ States of affairs that obtain outside of a person's control are considered lucky occurrences.²¹⁸ It is intuitive that chance events should not impact whether a person is morally responsible. Yet, we tend to blame person X more than person Y where the only difference between why X did A and Y did not to A is due to a factor beyond the agents' control. Our intuitive commitment that luck should not make a moral difference and our practice of allowing luck to make a moral difference shows that something is amiss.

The deep problem posed by moral luck is that given how many factors agents lack control over in the causal sequence of action, and our commitment to agents only being responsible to the extent that what they're morally responsible for was within their control, means that we should stop our blaming practices all together. Moral responsibility vanishes in some sense. Responsibility claims are without truth value.²¹⁹

²¹⁷ It is possible to view this concern as lurking in the background of much of what I have discussed in contrasting the accountability approach to responsibility with the attributability approach.

²¹⁸ Here I set aside the question of whether this is a coherent conception of luck. There is reason to think this conception of luck is incomplete at best. As discussed in chapter a previous chapter, luck includes a value or significance dimension. A low probability because outside the person's control event that is not significant for the person is not lucky. Moral luck may be sensitive to considerations of value and not just a purely statistical matter in terms of likelihood of occurrence given control over its occurrence.

²¹⁹ Nagel (1979) considers the deep problem of moral luck. This is the worry about whether we can (in principle) conform our ordinary judgments to the control principle. If this is possible in principle, then we would not blame people at all. Closer inspection reveals that human agency is impacted by all sorts of things over which people exercise no control. If such judgments are not warranted, as the Control Principle suggests, then responsibility ascriptions begin to vanish as Nagel (1979) explains:

[t]he area of genuine agency, and therefore of legitimate moral judgment, seems to shrink under this scrutiny to an extensionless point...as the external determinants of what someone has done are gradually exposed, in their effect on consequences, character, and choice itself, it becomes gradually clear that actions are events and people things. Eventually nothing remains which can be ascribed to the responsible self, and we are left with nothing but a portion of the larger sequence of events, which can be deplored or celebrated, but not blamed or praised.

While I will not pretend to settle the deep problem of moral luck in the remaining sections of this project I will explain how the New Attributionist approach that has emerged in this project provides theoretical support and motivation for pursuing a strategy to handle moral luck. Overall this strategy accepts and explain moral luck. It does so by countering a common way of denying the existence of two forms of moral luck, and it argues for acceptance of a different form of moral luck. Showing that the results of this project commend such a strategy does not prove that the strategy is the correct strategy for handling moral luck. But, it does show that this project offers insight into important challenges to moral responsibility beyond the epistemic challenge to the justification of responsibility judgments. This highlights an additional source of utility of the project.

With a high-level description of what I am up to in considering moral luck in mind, let me more systematically lay things out. First, I will get on the table key principles and different types of moral luck. Then, I will show how New Attributionism lends itself to handling the problem of moral luck.

In previous sections I argued that circumstances, history, culture, and so on, are not sufficient to excuse an agent from blameworthiness for actions done from moral ignorance.²²⁰ Recall the real-life case of Robert Alton Harris. His horrific childhood influenced his moral commitments and values. Who he was raised by, and how he was raised by them, was not under his control. In contrast to what I argued, a natural reaction is to think it unfair to consider him blameworthy for killing innocent people as an adult. Given his upbringing Robert Harris should be excused from

²²⁰ Provided that the other necessary conditions are met for blameworthiness.

blameworthiness for his wrongful actions. A principle that captures this idea that factors beyond our control should not be held against us when we are assessed morally is called the Control Principle (CP):

(CP) We are morally assessable only to the extent that what we are assessed for depends on factors under our control.²²¹

(CP) lends itself to a corollary principle. Imagine a person named Tom Barnaby Smith. Tom Smith only differs from Robert Harris regarding the circumstances of his upbringing. Tom was raised in a loving home. Yet were Tom to have had the bad fortune of being born into Robert's developmental circumstances Tom would have done the same wrongful things as Robert. Given that, by stipulation, the difference between Tom and Robert is a factor outside of their control (i.e., what environment they were raised in) is Tom less culpable than Robert? It is (bad) luck that Robert was born into an abusive upbringing, and it is (good) luck that Tom was born into a loving upbringing? A corollary of (CP) captures the idea that luck should not make a difference in how we morally assess Tom and Robert:

(CP-Corollary) Two people ought not to be morally assessed differently if the only other differences between them are due to factors beyond their control.²²²

If (CP) and its corollary are true, then we would do something we ought not do in differentially assessing the moral blameworthiness of Tom and Robert.²²³ Perhaps

²²¹ This is the principle as articulated by Dana Nelkin (2013).

²²² Dana Nelkin (2013).

²²³ CP appears supported by our ordinary judgments. If someone runs into me in a hallway, and I learn that someone pushed them into me without their consent to being so pushed, then I am inclined to consider the person not deserving of blame. They did not control whether they ran into me or not. Likewise, if two people are identical from the inside, and one person is driving and a dog runs in front of their car and is killed, and the other person in an otherwise identical circumstance does not have a dog run in front of their car, then the first driver hardly seems more blameworthy than the other driver. Whether the dog ran out in front of his car was not in his control.

Robert Harris should be excused from blameworthiness for his actions. Had Robert Harris been born into a supportive environment for healthy moral development he would have grown up to be like Tom Smith. He would not have committed murders.

Despite the intuitive pull of (CP) and its corollary many of our ordinary judgments about responsibility and punishment pull in the opposite direction. We are inclined to judge more deserving of blame the successful murderer as opposed to the mere attempted murderer though the only difference between the two is a factor out of their control, such as whether the intended victim tripped just prior to the gunshot being fired. Though legal responsibility is not equivalent to moral responsibility Robert Alton Harris was given the death penalty and executed for his crimes.²²⁴ Our ordinary judgments of responsibility align with Moral Luck (ML):

(ML) Moral luck occurs when an agent can be correctly treated as an object of moral judgment, despite the fact that a significant aspect of what he is assessed for depends on factors beyond his control.²²⁵

How do we reconcile the intuition that luck should not impact what a person is morally responsible for with the fact that luck does seem to make a moral difference in judging

²²⁴ After killing two young innocent boys in cold blood Robert Harris proceeded to eat the boys' food, he called his accomplice a sissy for getting sick after seeing him blow one of the boys brains out, and, as his accomplice recounted, "Harris pulled out the Luger, noticed blood stains and remnants of flesh on the barrel as a result of the point-blank shot, and said, 'I really blew that guy's brains out,' And then, again, he started laughing" (Watson 1993: 132). Likewise, I have argued that Robert is morally responsible and blameworthy precisely because he intentionally committed the murders of those innocent people, he did so while possessing full knowledge of the non-moral descriptive facts concerning what he was doing, his actions reflected a morally objectionable stance regarding how it is permissible to treat persons, and such actions reflected his evaluative moral commitments. On the other hand, Tom did not kill innocent people. Though he would have done so were he raised like Robert was raised, Tom did not actually intentionally do so. On my view, differential judgments of blameworthiness between Robert and Tom are correct.

²²⁵ This is a distillation by Nelkin (2013) of the definition of moral luck given by Nagel (1979), "Where a significant aspect of what someone does depends on factors beyond his control, yet we continue to treat him in that respect as an object of moral judgment, it can be called moral luck."

people responsible and blameworthy for wrongful actions? There are two general strategies. The first denies the control principle. It holds that there is nothing wrong with moral luck. The second approach retains the control principle and argues that, despite appearances otherwise, luck never actually impacts how a person is morally assessed. I will show how New Attributionism motivates pursuing the first strategy. Before showing how this works let's get on the table the three main types of moral luck: resultant, circumstantial, and constitutive.²²⁶

Resultant luck is luck concerning the outcome of actions. Earlier I mentioned such a case. For instance, if there are two people that are identical from an internal perspective regarding their intentions, knowledge, and so on, and both people want to kill a person, but in one case the potential victim trips and the shot misses and in the other case the person doesn't trip and is shot, we tend to judge the successful murderer more blameworthy than the attempted murderer, even though the only difference between the agents was a factor beyond their control, namely whether the intended victim tripped right before the shot.

Circumstantial luck is luck concerning a person's situation.²²⁷ It includes the culture in which one finds oneself and how that culture impacts one's opportunities for moral improvement and opportunities for praiseworthy and blameworthy choices. Nelkin (2013) discusses a case from Nagel illustrating circumstantial luck:

[C]onsider Nazi collaborators in 1930's Germany who are condemned for committing morally atrocious acts, even though their very presence in Nazi

²²⁶ These classifications stem from Thomas Nagel (1979). I leave out a fourth form of moral luck that significantly overlaps with the free will debate, which is causal luck. Many theorists view this type of luck as redundant with a combo of constitutive and circumstantial luck. See Nelkin (2013) for more on this point. I also set aside causal luck because the focus of this project is not free will.

²²⁷ Circumstantial moral luck was involved in the cases discussed in chapter 3 of this dissertation.

Germany was due to factors beyond their control (Nagel 1979). Had those very people been transferred by the companies for which they worked to Argentina in 1929, perhaps they would have led exemplary lives.²²⁸ If we correctly morally assess the Nazi collaborators differently from their imaginary counterparts in Argentina, then we have a case of circumstantial moral luck.

In the example above, we see that a change in circumstances is the difference between that person being a Nazi supporter versus that person living a morally decent life in Argentina. Had the person living in Germany during Hitler's reign been transferred by his company to Argentina he would not have performed the morally wrong actions for which he is blameworthy. Given that whether his company transferred him to Argentina or he remained in Germany was not within his control, yet he remained in Germany and is blameworthy for doing bad things, circumstantial moral luck factors into our moral assessment of him.

An additional form of moral luck is constitutive moral luck. This form of luck is often intertwined with circumstantial moral luck.²²⁹ But it is possible to pry apart these forms of luck. Constitutive luck concerns who you are, where this is determined by your genetics, environment, and the dispositions you develop through the interaction of these factors. It also concerns how you were raised. Though you do not exercise control over these factors, these factors can have a profound impact on who

²²⁸ Aaron Zimmerman suggests a helpful qualification of this retelling of Nagel by Nelkin. To say that such people would have, as Nelkin explains, "led exemplary lives" is an exaggeration. However, at least in Nagel (1979), I do not see Nagel casting the Argentinian counterpart as morally problematic. He explicitly claims that the counterpart involves a person, "who led a quiet and harmless life in Argentina." I think it best to claim that the counterpart in Argentina had anti-Semitic dispositions but did not act out of those dispositions. This includes not demonstrating microaggressions toward Jewish people living in Argentina.

²²⁹ For instance, the circumstances you happen to find yourself in can be the result of the interplay of nature and nurture making you the sort of person that was in a position to find yourself in circumstances that exposed you to an instance of moral luck.

you are and what you value. For example, JoJo observed his father doing his dictatorial work. JoJo admired his father as a great leader and a moral role model. Had JoJo not been born into a dictatorial family he would not have been raised to have such dispositions, such as the disposition to, on a whim, send his subjects to the torture chamber or sentence them to death. Yet, JoJo is rightly assessed based on these dispositional factors, even if he did not exercise control over coming to possess them. As Nelkin (2013) explains about this form of luck:

[I]f we correctly blame someone for being cowardly or self-righteous or selfish, when his being so depends on factors beyond his control, then we have a case of constitutive moral luck. Further, if a person acts on one of these very character traits over which he lacks control by, say, running away instead of helping to save his child, and we correctly blame him for so acting, then we also have a case of constitutive moral luck. Thus, since both actions and agents are objects of moral assessment, constitutive moral luck undermines the Control Principle when it comes to the assessment of both actions and agents.

Constitutive moral luck comes into play regarding the cases in chapter 2. The main case we discussed was the Potter case. In that case, Potter possesses a cadre of vices. Potter was not a hopeless victim to his vices, as if they were compulsions. But, it is plausible that he lacked control over coming to possess those vices. Perhaps this is because of how he was raised combined with his genetic predispositions regarding his personality. He “naturally” came to be overly self-assured, factor his own interests more highly than the interests of others, and to see capitalistic values as the ultimate values to guide one’s decisions. If so, then Potter is being assessed for decisions that stem from traits that he ultimately did not exercise control in coming to possess. Potter is subject to constitutive luck.

This raises a question of how active children are in the development of their own moral character?²³⁰ Children are not entirely passive in this process, but they do not typically exercise control over their formation of their moral character. Moral psychology posits different stages of moral development.²³¹ As children go through early stages of development, morality is largely externally controlled and based on imitation of role models and authority figures. In the later stage of moral development (i.e. postconventional), abstract principles come to govern the child's morality (e.g. rules about justice), but those principles are largely inculcated. It is not typically until reaching adulthood that one reevaluates such principles and either takes ownership of those principles (i.e., comes to see them as reflecting what one truly values) or rejects those moral principles and exercises control over shaping their moral character by considering a different set of values.²³² What about morally decent children that emerge from the worst conditions regarding their moral upbringing? Such children do not directly choose to reject the distorted moral values they are raised with. Rather, as is most often the case, it is a combination of genetic and personality dispositions coupled with good moral role models in the form of other adults or kids at school that help such children emerge from horrific upbringings to become morally decent people. Such interplay of genetics, personality, and social influences (i.e. factors outside their control) shape their development.²³³

²³⁰ I thank Aaron Zimmerman for suggesting this consideration.

²³¹ This is especially true of Kohlberg's Stage Theory of moral development. For more on moral psychology, including a critical evaluation of Kohlberg's theory, see Lapsley (2018).

²³² This most obviously occurs with children raised in a religious context. As adults they either make those values their own or reject them. Or, sadly, some people continue blindly following the religious precepts of their youth without examining the evidential and moral merits of those precepts.

²³³ Another possibility that I do not discount is a religious possibility. If the God of traditional theism exists, then it is possible that such an all-good, all-powerful, and all-knowing God empowers such children to become morally decent in the face of countervailing forces heading in the other direction.

With the types of moral luck on the table, in the next two sections I will outline a strategy for handling moral luck. This strategy is commended by the picture of moral responsibility that has emerged in this project. Again, I do not claim this solves the problem of moral luck. Rather, the New Attributionist perspective suggests such a strategy and this demonstrates a way that the view has utility in extending it to additional problems in moral responsibility.

4.9 Denying the Epistemic Argument

A common approach to explaining away moral luck is to endorse what is called the “epistemic argument.”²³⁴ According to this argument it only appears like moral luck exists because when we differentially judge people based on factors over which they did not exercise control we possess differential evidence about those agents. We are in a different epistemic position regarding the evidence we have relevant to evaluating those agents. This explains why we judge the agents differentially. This argument is applied to resultant luck and circumstantial luck. Let’s see how this works by looking at a case we discussed earlier that involves resultant luck.

Consider two prospective assassins. One assassin is successful, and the other is not because the intended target trips prior to the bullet hitting them. When the person succeeds in killing their target, ascribers of responsibility gain evidence concerning the plans, intentions, and commitments of the murderer. This same evidence is not typically available with mere attempts. When a plan is not brought to full completion

This may occur independently of genetics, personality, and social influences or God may orchestrate such influences to providentially produce the intended result, perhaps in the service of a higher good.

²³⁴ Latus (2000) calls this the epistemic argument. Proponents of this argument include Richards (1986), Rescher (1993), Rosebury (1995), and Thomson (1993).

we often do not know what the person intended to do or to what extent they were committed to a plan of action. Successful attempts provide fallible evidence of what a person intended to do or how committed they were to a plan. This does not show that moral luck affects the truth value of what a person is morally responsible for. Instead, it shows that there is an epistemic difference concerning moral responsibility judgments. By contrast, as Nelkin (2013) explains citing Thomson (1993),

If we were in the unrealistic situation of knowing that both agents had exactly the same intentions, the same strength of commitment to their plans, and so on, then we would no longer be inclined to treat them differently. Thomson represents a number of those who employ this strategy when she asks, “Well do we regard Bert [a negligent driver who causes a death] with an indignation that would be out of place in respect to Carol [an equally negligent driver who does not]? Even after we have been told about how bad luck figured in his history and good luck in hers?” And Thomson answers: “I do not find it in myself to do so” (1993, 205).

Going back to the assassin case, it is only in the unrealistic situation that we are told that the two assassins had the same intentions, plans, and commitments to their plans that we possess the same evidence concerning the agents. In real life, we rarely possess such information. One might claim that were we to possess that same evidence concerning the agents, we would equally ascribe responsibility to the agents. What appears as moral luck differentially impacting ascriptions of responsibility across agents is just the fact that successful actions provide us evidence of the intentions of

the agents. Unsuccessful actions do not provide this same evidence. This is the evidential argument concerning resultant luck.

Now let's look at the epistemic argument as applied to circumstantial luck. Recall that circumstantial luck is the most relevant form of luck given the arguments of chapter 3 whereby moral ignorance is not an excuse, even in cases whereby the agent has been shaped by finding himself in a benighted culture. Nelkin (2013) aptly explains how the epistemic argument works regarding circumstantial luck:

Consider again the Nazi sympathizer, and a counterpart who moved in 1929 to Argentina on business. The counterpart has exactly the same dispositions as the Nazi sympathizer, but lives a quiet and harmless life in Argentina. According to this line of reasoning, while it is true that the counterpart is not responsible for the same deeds as the Nazi sympathizer, he should be judged precisely for what he would have done. Richards argues that we do judge people for what they would have done, but that what they do is often our strongest evidence for what they would have done. As a result, given our limited knowledge, we might not be entitled to treat the counterpart in the same way as the Nazi sympathizer, even though they are equally morally deserving of such treatment (Richards 1986, 174 ff.). Thus, circumstantial luck, like resultant luck, affects the basis available to us when we judge agents, but does not affect what those agents deserve.

According to the epistemic argument explaining away the existence of moral luck, we judge people responsible considering what they would have done given their dispositions. Yet we often lack evidence of what people would have done were they in different circumstances. The Argentinian counterpart to the Nazi sympathizer cannot be magically transported to Nazi Germany for us to see what that person would have done. We observe him living a harmless life in Argentina while a person with the same dispositions living in Nazi Germany commits moral injustices against Jewish people.

Had we possessed the evidence of what the Argentinian counterpart would have done, even though they did not commit the same acts, the Argentinian would have deserved equal blame as the Nazi sympathizer.

The key results of this dissertation offer a nuanced analysis of the epistemic argument. The argument gets something right, and it gets something wrong. New Attributionism about the epistemic condition on moral responsibility affords insight into why what a person actual does is evidence for what they would have done. Recall that the Blameworthiness principle requires intentional action with knowledge of the non-moral facts in full view. Successful actions provide us with fallible evidence of what a person intended to do. It provides us with evidence of their commitment to a course of action. This same evidence is not available with mere attempts. This much the epistemic argument gets right.

The epistemic argument also gets right that, “If we were in the unrealistic situation of knowing that both agents had exactly the same intentions, the same strength of commitment to their plans, and so on, then we would no longer be inclined to treat them differently” (Nelkin 2013). Not being inclined to differentially judge the agents is the case because in that unrealistic situation of possessing the same evidence of the intentions and dispositions of the agents we learn that performing the action in question aligns with the agents’ evaluative moral commitments. It reflects their deep selves.

Despite what the epistemic argument gets right, according to New Attributionism it is a mistake to leap from the fact about how we are inclined to treat the agents given possession of similar evidence about their intentions to the fact that

it is true that the agents are equally blameworthy. Outright blameworthiness tracks what an agent actually did, not what they would have done had luck not intervened. This is to hold that we correctly judge the agents morally responsible for what they actually did; we do not merely judge them for what they would have done, even in the unrealistic situation where we possess knowledge of the truth of such a counterfactual.

Such a New Attributionist take on the epistemic argument highlights the fact that life is not always fair. Not everyone is dealt an equal hand in the outcomes of their actions or in the circumstances of their actions. Sometime morality aligns with the unfairness of life. Sometimes just deserts result from factors outside one's control.²³⁵ New Attributionism captures this reality.

That New Attributionism offers a nuanced analysis of the epistemic argument shows the utility of the approach. It is not flat-footed in how it can be used to approach arguments in the literature on moral luck. This also shows that New Attributionism is inclined to accept moral luck or deny the control principle. In the next section, I will show how this approach lends itself to such a strategy.

4.10 Accepting Constitutive Moral Luck

The brand of New Attributionism developed in this project aligns with an approach to rejecting the control principle. This approach accepts constitutive moral luck, which is moral luck concerning our traits and dispositions. Many people come to possess traits and dispositions due to genetics and environmental and developmental influences. Moral luck comes into play when a person deserves blame for performing

²³⁵ As Nagel (1979) indicates, it is a Kantian impulse to try to make morality immune from luck.

an action that stems from a trait or disposition that they formed due to influences they did not have control over. For instance, if someone is raised by parents that instill in that child a selfish disposition and they grow into an adult that performs wrongful acts (from ignorance) due to their selfish disposition, yet we still correctly judge the person blameworthy for such acts, then the person is subject to constitutive moral luck. Establishing that such constitutive moral luck exists undermines the control principle.²³⁶ So, one way of arguing against the control principle is arguing that we should accept constitutive moral luck. In what follows, I sketch why the key results of this project support or motivate such an approach.

New Attributionism is an approach to moral responsibility that takes its inspiration from a paper by Robert Adams (1985). Adams focuses on our ordinary practices and judgments regarding traits and dispositions largely outside of a person's control. For instance, it is common to judge a person blameworthy for having a racist attitude toward a minority group even if they are not in control of that attitude and they came to possess that attitude due to factors outside their control. Being raised in a racist household does not ordinarily excuse a person for having racist attitudes and performing racist actions that reflect a morally objectionable stance toward persons. According to Adams, intentional states of mind (i.e., states of mind directed toward or about something) "for which we are directly responsible are those in which we are responding, consciously or unconsciously, to data that are rich enough to permit a fairly adequate ethical appreciation of the state's intentional object and of the object's

²³⁶ If moral luck exists, people are morally responsible for actions and attitudes that depend on factors beyond their control. The control principle claims that such judgments reflecting moral luck are illegitimate as such judgments do have as their propositional content a genuine luck-based phenomenon. Such judgments are explained in terms of a different non-luck-based phenomenon.

place in the fabric of personal relationships” (Adams 1985: 27).²³⁷ For instance, a white woman may harbor racist attitudes toward black men. She may associate black men with danger. Even if a black man approaching her poses no threat she may clench her purse for fear of the black man trying to snatch her purse. Imagining that she does not know that such an attitude or reaction toward the black man is morally wrong, as she is ignorant of its wrongness, she is still morally responsible and blameworthy for her racist attitude and for her actions stemming from that attitude.²³⁸ This holds even if she does not control the attitude and did not control coming to possess it. Applying the Blameworthiness principle, such an attitude reflects a morally objectionable stance toward persons of color, and it reflects her evaluative commitments. It reflects the fact that she judges it true that black men are dangerous, and it reflects the fact that it is good to protect your belongings when in the presence of black men. Thus, the blameworthiness principle is applicable to an approach to accepting constitutive moral luck.²³⁹ It is useful in this way. It further illuminates why, on such an

²³⁷ Responsibility for such an attitude does not require that a person actually appreciates the ethical dimension of the intentional object of the attitude, even though such a “fairly adequate ethical appreciation” is permitted by the richness of the dataset. Instead the data make such a basic understanding possible. Though, in cases of moral ignorance, there will be varying degrees to which it is difficult for persons to overcome that ignorance and appreciate the moral dimension of the intentional object of the attitude afforded by the data. If such ethical appreciation were required, then the question of direct moral responsibility for acts done from moral ignorance would not enter the picture. Attitudes absent such an appreciation of the moral dimension of the object of the attitude (such as the moral wrongness of racism) would not be proper targets of moral judgment. Further, in such a case, moral ignorance would be a way to escape moral responsibility.

²³⁸ Here I am not arguing for this. I am assuming it both given what I have already argued and the goal of this section. The goal is to show that my approach can explain this intuition or judgment. This is not simply: attributability is correct therefore attributability is correct. Rather, an older form of attributability is bolstered and extended by the key developments of this dissertation.

²³⁹ Again, I do not claim to have conclusively established the falsity of the control principle. Arguing against it was not the goal of this section. Rather, I showed that the Blameworthiness principle is applicable to an approach that denies the control principle.

attributionist approach, we are morally responsible and blameworthy for things beyond our control.

Relating this back to the Control Principle, one way of looking at what Adams (1985) is up to is showing that the Control Principle is mistaken for a more general principle.²⁴⁰ The more general principle is that we are responsible for things that spring from us in the right way. We are responsible for things that are correctly attributable to us. According to the Blameworthiness principle, such actions must be freely and intentionally committed, they must be committed while possessing knowledge of the relevant non-moral (descriptive) facts or be non-culpably ignorant of them, and they must reflect a morally objectionable stance toward persons or reflect a person's evaluative moral commitments. When these conditions hold the attitude or action springs from the person in the right way for it to be a proper basis of moral evaluation in terms of moral responsibility. By contrast, the Control Principle is overly restrictive. According to the New Attributionist, control over the basis of what is morally assessed may or may not obtain. If it does not obtain, then what matters is that the agent bear the right relation to what they are assessed for. The attitude and action must have a proper link to the agent's cognition. The Control Principle is often confused for a broader vantagepoint of what is needed for moral responsibility. Thus, again, the key results of this project motivate such an approach to countering the Control Principle. They extend an older attributionist model.

²⁴⁰ This is how Nelkin (2013) reads part of Adams' project in his paper.

4.11 Chapter and Dissertation Conclusion

In this chapter, I brought together the developments of chapters 2 and 3. I tied them together into a Blameworthiness principle and an Ignorance principle. I picked up on the Cleo case by introducing Wolf's famous JoJo case. That case involved an evil dictator named JoJo who was unlucky to have been born into a situation devoid of alternative moral perspectives that could have led him to rectify his ignorance. As in chapter 3, I argued that moral ignorance is not an excuse. JoJo is blameworthy for harming his citizens. Supporting this argument were empirical results that suggested that this is how people actually respond to the JoJo case. They do not let JoJo off the hook because he was raised by a dictator father who allowed him to watch how he ruled his people by the whims of moral depravity, even though exposure to moral alternatives, as in a different case, slightly increased the degree of blameworthiness. This showed that moral ignorance due to factors outside one's control does not excuse one from responsibility. Instead, JoJo reflectively endorsed his actions as reflecting what he most valued, who he wanted to be, and what kind of life he wanted to live. This supported the attributionist perspective that reflectively endorsing one's evaluative judgments or acting in a way that evidenced an objectionable stance toward persons—as viewing it permissible to torture and kill them on a lark—is indictment enough that the person is blameworthy. The Deep Self View is a version of New Attributionism that captures this idea. Such a view does not emphasize that it is only fair to judge people responsible for things they control. We rightly judge people blameworthy based on factors they do not control, such as the outcomes of their actions, the circumstances surrounding them, and the environmental pressures that help shape them into who they are.

I applied this New Attributionist position by considering the problem of moral luck. I introduced the key principles in the moral luck debate, and I outlined the types of moral luck. I proceeded to show how the results of this project could be used to motivate undermining an argument (i.e. the epistemic argument) to the effect that moral luck does not exist. According to the key results of this project the argument gets some things right and some things wrong. This facilitated a nuanced response to the epistemic argument, which provided utility for better understanding an influential argument in the literature. Lastly I discussed the old attributionism (or original attributionism) as put forward by Adams (1985). I showed how the New Attributionist approach to moral responsibility in this project could extend and further illuminate how attributionism accepts the existence of constitutive moral luck. Again, this showed the utility of the results of this dissertation for important challenges to moral responsibility.

In closing, I have expanded the scope of the epistemic condition on moral responsibility. This has produced responsibility beyond belief—responsibility beyond whether an agent acts contrary to what they believe they should do overall in a given situation. Instead of the restrictive internalist perspective, a more permissive and nuanced picture has emerged—one that aligns with and vindicates our ordinary practices of judging people responsible and blameworthy for their actions. It is my hope that this, coupled with other recent efforts, gives the epistemic dimension of moral responsibility the adequate attention it deserves along with the notion of free will that is likewise a necessary condition for moral responsibility.

REFERENCES

- Adams, R. M. (1985). 'Involuntary Sins'. *The Philosophical Review*, 94, 3-31.
- Anderson, E. (2014). 'Social Movements, Experiments in Living, and Moral Progress: Case Studies from Britain's Abolition of Slavery'.
- Anscombe, G. E. M. (1957). *Intention*: Harvard University Press.
- Anscombe, G. E. M., & Morgenbesser, S. (1963). 'The two kinds of error in action'. *Journal of Philosophy*, 60, 393-401.
- Aristotle. (1985). *Nicomachean Ethics translated by Terence Irwin*: Hackett Publishing Company.
- . (2014). *Nicomachean Ethics* (C. D. C. Reeve, Trans.): Hackett Publishing Company.
- Arpaly, N. (2003). *Unprincipled Virtue: An Inquiry Into Moral Agency*: Oxford University Press.
- Battaly, H. (2015). *Virtue: Polity*.
- Bennett, J. (1974). 'The Conscience of Huckleberry Finn'. *Philosophy*, 49, 123-34.
- Benson, P. (1987). 'Freedom and Value'. *The Journal of Philosophy*, 84, 465-86.
- . (2001). 'Culture and Responsibility: A Reply to Moody-Adams'. *Journal of Social Philosophy*, 32, 610-20.
- Bradford, G. (2017). Hard to Know. In P. Robichaud & J. W. Wieland (Eds.), *Responsibility - The Epistemic Condition* (pp. 180-98): Oxford University Press.
- Buckwalter, W. (2016). 'Intuition Fail: Philosophical Activity and the Limits of Expertise'. *Philosophy and Phenomenological Research*, 92, 378-410.
- Burns, J. F. (2003). Aftereffects: Reign of Terror; Soccer Players Describe Torture by Hussein's Son. *The New York Times*. Retrieved from <https://www.nytimes.com/2003/05/06/world/aftereffects-reign-of-terror-soccer-players-describe-torture-by-hussein-s-son.html>
- Clarke, R. (2005). 'On an Argument for the Impossibility of Moral Responsibility'. *Midwest Studies In Philosophy*, 29, 13-24.
- Coates, D. J., & Tognazzini, N. A. (2013). *Blame: Its nature and norms*: Oxford University Press.
- Coffman, E. J. (2007). 'Thinking about luck'. *Synthese*, 158, 385-98.
- Driver, J. (1989). 'The Virtues of Ignorance'. *Journal of Philosophy*, 86, 373-84.
- Echeñique, J. (2012). *Aristotle's Ethics and Moral Responsibility*: Cambridge University Press.
- Evans, I. (2013). 'The problem of the basing relation'. *Synthese*, 190, 2943-57.
- Faraci, D., & Shoemaker, D. (2010). 'Insanity, Deep Selves, and Moral Responsibility: The Case of JoJo'. *Review of Philosophy and Psychology*, 1, 319-32.
- Fields, L. (1994). 'Moral Beliefs and Blameworthiness'. *Philosophy*, 69, 397-415.
- Fischer, J. M. (2010). 'The Frankfurt cases: The moral of the stories'. *Philosophical Review*, 119, 315-36.
- Fischer, J. M., & Ravizza, M. (1998). *Responsibility and Control: A Theory of Moral Responsibility*: Cambridge University Press.
- FitzPatrick, W. J. (2008). 'Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge'. *Ethics*, 118, 589-613.
- . (2017). Unwitting Wrongdoing, Reasonable Expectations, and Blameworthiness. In P. Robichaud & J. W. Wieland (Eds.), *Responsibility: The Epistemic Condition* (pp. 29-46): Oxford University Press.

- Frankfurt, H. G. (1969). 'Alternate possibilities and moral responsibility'. *The Journal of Philosophy*, 829-39.
- Frankfurt, H. G. (1971). 'Freedom of the Will and the Concept of a Person'. *The Journal of Philosophy*, 68, 5-20.
- . (1988). *The Importance of What We Care About: Philosophical Essays*: Cambridge University Press.
- Ginet, C. (2000). 'The Epistemic Requirements for Moral Responsibility'. *Noûs*, 34, 267-77.
- Goldman, A. I. (1986). *Epistemology and Cognition* (Vol. 98): Harvard University Press.
- Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*: Cambridge University Press.
- Guerrero, A. A. (2007). 'Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution'. *Philosophical Studies*, 136, 59-97.
- . (2017). Intellectual Difficulty and Moral Responsibility. In P. Robichaud & J. W. Wieland (Eds.), *Responsibility - The Epistemic Condition* (pp. 199-218): Oxford University Press.
- Haack, S. (2001). The Ethics of Belief 'Reconsidered'. In M. Steup (Ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue* (pp. 21-33): Oxford University Press.
- Haji, I. (2013). 'Conversation and responsibility'. *Canadian Journal of Philosophy*, 43, 267-86.
- Hanna, N. (2014). 'Moral Luck Defended'. *Noûs*, 48, 683-98.
- Harman, E. (2011). 'Does Moral Ignorance Exculpate?'. *Ratio*, 24, 443-68.
- . (2014). 'Ethics is Hard! What Follows?'. (ms).
- . (2017). When is Failure to Realize Something Exculpatory? In P. Robichaud & J. W. Wieland (Eds.), *Responsibility: The Epistemic Condition* (pp. 117-26): Oxford University Press.
- Hieronymi, P. (2006). 'Controlling Attitudes'. *Pacific Philosophical Quarterly*, 87, 45-74.
- . (2008). 'Responsibility for believing'. *Synthese*, 161, 357-73.
- . (2009). 'Believing at Will'. *Canadian Journal of Philosophy, Supplementary Volume*, 35, 149-87.
- . (2014). 'Reflection and Responsibility'. *Philosophy & Public Affairs*, 42, 3-41.
- Houlgate, L. D. (1968). 'Knowledge and Responsibility'. *American Philosophical Quarterly*, 5, 109-16.
- Lapsley, D. K. (2018). *Moral Psychology*: Routledge.
- Latus, A. (2000). 'Moral and Epistemic Luck'. *Journal of Philosophical Research*, 25, 149-72.
- Le Morvan, P. (2011). 'Knowledge, Ignorance and True Belief'. *Theoria*, 77, 32-41.
- Levy, N. (2009). 'Culpable Ignorance and Moral Responsibility: A Reply to FitzPatrick'. *Ethics*, 119, 729-41.
- . (2011). *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*: Oxford University Press.
- . (2016). 'Culpable Ignorance: A Reply to Robichaud'. *Journal of Philosophical Research*.
- Machery, E., & O'Neill, E. (2014). *Current Controversies in Experimental Philosophy*: Routledge.
- Markovits, J. (2010). 'Acting for the Right Reasons'. *Philosophical Review*, 119, 201-42.

- McCauley, R. N. (2011). *Why Religion is Natural and Science is Not*: Oxford University Press.
- McHugh, C. (2011). 'Epistemic Deontology and Voluntariness'. *Erkenntnis*, 77, 65-94.
- Montmarquet, J. (1999). 'Zimmerman on Culpable Ignorance'. *Ethics*, 109, 842-5.
- Moody-Adams, M. M. (1994). 'Culture, Responsibility, and Affected Ignorance'. *Ethics*, 104, 291-309.
- Nagel, T. (1979). *Mortal Questions*: Cambridge University Press.
- Nelkin, D. K. (2013). 'Moral Luck'. *Stanford Encyclopedia of Philosophy*.
- Nozick, R. (1981). *Philosophical Explanations* (Vol. 92): Harvard University Press.
- Peels, R. (2010). 'What Is Ignorance?'. *Philosophia*, 38, 57-67.
- . (2011). 'Tracing Culpable Ignorance'. *Logos & Episteme: An International Journal of Epistemology*, 2, 575-82.
- . (2014). 'What Kind of Ignorance Excuses? Two Neglected Issues'. *Philosophical Quarterly*, 64, 478-96.
- Pereboom, D. (2013). Free Will. In R. Crisp (Ed.), *The Oxford Handbook of the History of Ethics* (pp. 609-37): Oxford University Press.
- . (2014). *Free Will, Agency, and Meaning in Life*: Oxford University Press.
- Plantinga, A. (2009). *Knowledge of God*: John Wiley & Sons.
- Pritchard, D. (2005). *Epistemic Luck* (Vol. 29): Clarendon Press.
- . (2007). 'Anti-Luck Epistemology'. *Synthese*, 158, 277-97.
- Rescher, N. (1993). Moral luck. In D. Statman (Ed.), *Moral Luck* (pp. 141--66): Suny Press.
- . (1995). *Luck: The Brilliant Randomness of Everyday Life*: Farrar, Straus and Giroux.
- Richards, N. (1986). 'Luck and desert'. *Mind*, 95, 198-209.
- Rivera-López, E. (2006). 'Can There Be Full Excuses for Morally Wrong Actions?'. *Philosophy and Phenomenological Research*, 73, 124-42.
- Robichaud, P. (2014). 'On Culpable Ignorance and Akrasia'. *Ethics*, 125, 137-51.
- Robichaud, P., & Wieland, J. W. (2017). *Responsibility: The Epistemic Condition*: Oxford University Press.
- Rosebury, B. (1995). 'Moral responsibility and "moral luck"'. *Philosophical Review*, 104, 499-524.
- Rosen, G. (2003). *IV—Culpability and Ignorance*. Paper presented at the Proceedings of the Aristotelian Society.
- . (2004). 'Skepticism about Moral Responsibility'. *Noûs*, 38, 295-313.
- Sauvé Meyer, S. (2011). *Aristotle on Moral Responsibility: Character and Cause*: Oxford University Press.
- Scanlon, T. M. (1998). *What We Owe to Each Other*: Belknap Press of Harvard University Press.
- . (2008). *Moral Dimensions*: Harvard University Press.
- Schroeder, M. (2015). 'Is Knowledge Normative?'. *Philosophical Issues*, 25, 379-95.
- Sher, G. (2006). *In Praise of Blame*: Oxford University Press.
- . (2009). *Who Knew?: Responsibility Without Awareness*: Oxford University Press.
- Shoemaker, D. (2011). 'Attributability, Answerability, and Accountability: Toward a Wider Theory of Moral Responsibility'. *Ethics*, 121, 602-32.
- Smith, A. M. (2005). 'Responsibility for Attitudes: Activity and Passivity in Mental Life'. *Ethics*, 115, 236-71.
- . (2007). 'On Being Responsible and Holding Responsible'. *The Journal of Ethics*, 11, 465-84.

- . (2008). 'Control, responsibility, and moral assessment'. *Philosophical Studies*, 138, 367-92.
- . (2012). 'Attributability, Answerability, and Accountability: In Defense of a Unified Account'. *Ethics*, 122, 575-89.
- . (2015). 'Responsibility as Answerability'. *Inquiry*, 58, 99-126.
- Smith, H. (1983). 'Culpable Ignorance'. *Philosophical Review*, 92, 543-72.
- Smith, H. M. (2014). 'The Subjective Moral Duty to Inform Oneself before Acting'. *Ethics*, 125, 11-38.
- Sosa, E. (1999). 'How to defeat opposition to Moore'. *Philosophical Perspectives*, 13, 137-49.
- Steup, M. (2008). 'Doxastic freedom'. *Synthese*, 161, 375-92.
- Strabbing, J. T. (2016). 'Responsibility and Judgment'. *Philosophy and Phenomenological Research*, 92, 736-60.
- Strawson, G. (1994). 'The Impossibility of Moral Responsibility'. *Philosophical Studies*, 75, 5-24.
- Strawson, P. F. (1962). 'Freedom and resentment'. *Proceedings of the British Academy*, 48, 1-25.
- Strawson, P. F. (1993). Freedom and Resentment. In J. M. Fischer & M. Ravizza (Eds.), *Perspectives on Moral Responsibility* (pp. 45-66): Cornell University Press.
- Talbert, M. (2012). 'Accountability, Aliens, and Psychopaths: A Reply to Shoemaker'. *Ethics*, 122, 562-74.
- . (2013). Unwitting Wrongdoers and the Role of Moral Disagreement in Blame. In D. Shoemaker (Ed.), *Oxford Studies in Agency and Responsibility Volume 1*: Oxford University Press.
- . (2016). *Moral Responsibility: An Introduction*: Polity.
- Thomson, J. J. (1993). Morality and Bad Luck. In D. Statman (Ed.), *Moral Luck*: State University of New York Press.
- Unger, P. K. (1975). *Ignorance: A Case for Scepticism*: Oxford University Press.
- Van Woudenberg, R. (2009). 'Ignorance and Force: Two Excusing Conditions for False Beliefs'. *American Philosophical Quarterly*, 46, 373-86.
- Vargas, M. (2012). *Building Better Beings: A Theory of Moral Responsibility*: Oxford University Press.
- Wallace, R. J. (1996). *Responsibility and the Moral Sentiments*: Harvard University Press.
- Wallace, R. J., Kumar, R., & Freeman, S. (2011). *Reasons and Recognition: Essays on the Philosophy of T. M. Scanlon*: Oxford University Press.
- Waller, B. N. (2017). *The Injustice of Punishment*: Routledge.
- Watson, G. (1975). 'Free Agency'. *The Journal of Philosophy*, 72, 205-20.
- . (1993). Responsibility and the Limits of Evil. In J. M. Fischer & M. Ravizza (Eds.), *Perspectives on Moral Responsibility*: Cornell University Press.
- . (1996). 'Two Faces of Responsibility'. *Philosophical Topics*, 24, 227-48.
- . (2004). *Agency and Answerability: Selected Essays*: Oxford University Press.
- Weatherson, B. (2008). 'Deontology and Descartes's Demon'. *Journal of Philosophy*, 105, 540-69.
- Wieland, J. W. (2017a). Introduction: The Epistemic Condition. In P. Robichaud & J. W. Wieland (Eds.), *Responsibility - The Epistemic Condition*: Oxford University Press.
- . (2017b). 'What's Special about Moral Ignorance?'. *Ratio*, 30, 149-64.
- . (2017c). 'Willful Ignorance'. *Ethical Theory and Moral Practice*, 20, 105-19.

- Wiland, E. (2012). *Reasons*: Continuum.
- Williamson, T. (2000). *Knowledge and its Limits*: Oxford University Press.
- Wolf, S. (1987). Sanity and the Metaphysics of Responsibility. In F. D. Schoeman (Ed.), *Responsibility, Character, and the Emotions: New Essays in Moral Psychology* (pp. 46-62): Cambridge University Press.
- . (1990). *Freedom Within Reason*: Oxford University Press.
- Wright, C. (2014). On Epistemic Entitlement (II): Welfare State Epistemology. In D. Dodd & E. Zardini (Eds.), *Scepticism and Perceptual Justification* (pp. 213-47): Oxford University Press.
- Zimmerman, M. J. (1988). *An Essay on Moral Responsibility*: Rowman & Littlefield.
- . (1997). 'Moral Responsibility and Ignorance'. *Ethics: An International Journal of Social, Political, and Legal Philosophy*, 107, 410-26.
- . (2008). *Living with Uncertainty: The Moral Significance of Ignorance*: Cambridge University Press.