On Culpable Ignorance: Answering the Debate between Epistemic Vice and Internalism

by

Cathy Zhu

A thesis

presented to the University of Waterloo

in fulfillment of the requirement for the degree of

Master of Arts

in

Philosophy

Waterloo, Ontario Canada, 2025

© Cathy Zhu 2025

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Imagine I invite you over for tea. I ask you if you want any sugar in your tea and you say "yes, please!" So I take the sugar bowl out of my cupboard and spoon some sugar into your tea. Little do I know, someone had snuck into my house in the middle of the night and switched out the sugar in my sugar bowl with arsenic. You drink the arsenic tea and fall deathly ill. By putting what I think is sugar into your tea, I cause you to fall ill and potentially cause your death. But all along, I had no idea that the contents of my sugar bowl was arsenic, not sugar. Am I blameworthy for poisoning you?

The topic of when an agent can be held responsible, or culpable, in cases where they did not know what they were doing has gained attention in recent years. The field of culpable ignorance aims to look at when agents are culpable for unwitting wrongful acts performed from ignorance, and what criteria must be fulfilled to judge an agent as morally responsible for her wrongdoing.

I am to defend a version of what is called the Epistemic Vice View held by William FitzPatrick, which states that normative ignorance can be culpable if it is a result of epistemic vices such as arrogance, laziness, or dismissiveness, given that the agents *could have* reasonably remedied their ignorance had they not exercised such epistemic vice. However, it has been met with critique. I focus on Neil Levy's Internalist critique, which states that what an agent can *reasonably* know and have done is a state *internal* to the agent herself.

I argue that Levy is incorrect in his internalism because his theory is founded on two false assumptions. In doing so, I hope to defend FitzPatrick's Epistemic Vice View. I argue that the Epistemic Vice View captures our intuitions about blameworthiness, while also avoiding the extremes of both excusing all ignorance and condemning all unwitting wrongdoing. My defense

contributes to the broader debate by offering a nuanced standard for culpable ignorance, one that preserves the integrity of moral responsibility without requiring radical revision of our practices of blame.

In the first chapter, I introduce the topic of culpable ignorance, and delineate the main ways theorists aim to solve the issue. In the second chapter, I outline the debate between William FitzPatrick and Neil Levy. Finally, in the third chapter, I pose critiques to Levy's internalist position.

Acknowledgements

I consider myself lucky to be a graduate student at the University of Waterloo. And for it, I have many people to thank.

First, I would like to thank my supervisor, Dr. Mathieu Doucet. Without his comments, guidance, and suggestions, this thesis would not be what it is today (and may not even exist). I am grateful and humbled for having worked with such a brilliant and calming person on the biggest project of my life thus far. On days where I felt like this endeavour was hopeless, I relied a lot on your assurance. Thank you for your work on this project.

I need to give an endless "thank you" to my parents. My academic journey began at home with books and storytimes and perfectly thought out answers to my constant, relentless questions. It continued with financial support through my undergraduate degree. And while I do not know how it will conclude (I doubt it ever will), I know that every text, call, and visit means the world to me. It is the best support I could ask for. I also need to thank my first ever student, my little brother. I hope to never stop listening, teaching, and guiding you through life. To my family, thank you for your faith in my abilities and your unconditional love. I would not have made it to this point in my life without you.

I am also infinitely thankful of Ewa Rygielski, Darryl Firsten, Pawel Rygielski, and the entire team at IN8 who kept me employed throughout the completion of my Masters. Ewa, thank you for taking a chance on me when I was a recent graduate, and for taking an even bigger chance by keeping me on the team when I told you I was starting school again. Although you didn't give lectures, you were one of my greatest teachers during this time. Your passion, love, work ethic, and business sense (among many other things) have been an inspiration to me and the lessons I

have learned from you will stick with me for a lifetime. You have shown me what dedication truly is, and I wish you all the best.

Thank you to my husband, for your unfettered support. I am so grateful to have you in my life. You calm me, motivate me, and advise me in so many facets of my life. While I may not say it enough, I think you are brilliant. And thank you to my friends, without whom I would never leave the confines of my home office (you know who you are). I have loved taking trips with you all, sharing meals, and exasperatingly unloading the stress of my education. Your community is priceless and irreplaceable.

To everyone I have mentioned and those that I did not, I want to extend an endless thank you. I am lucky I have crossed paths with each and every one of you. I have learned so much from every interaction I have, no matter how insubstantial or short. And for that, I am grateful and humbled.

Table of Contents

Author's Declaration ii

Abstract iii

Acknowledgements v

Table of Contents vii

1.1. INTRODUCTION 1

1.2. WHAT IS CULPABLE IGNORANCE? 2

1.3. TRACING VIEWS 6

1.4. CONSERVATIVE VIEWS 11

1.5. EPISTEMIC VICE VIEWS 12

1.6. INTERNALIST VIEWS 16

1.7. QUALITY OF WILL VIEWS 18

1.8. CONCLUSION 22

CHAPTER 2 24

2.1 INTRODUCTION 24

2.2. BACKGROUND: ROSEN AND SKEPTICISM ABOUT MORAL RESPONSIBILITY 24

2.3. WHAT AGENTS CAN REASONABLY HAVE DONE: THE DEBATE BETWEEN FITZPATRICK AND LEVY 29

2.4. CONCLUSION 42

CHAPTER 3 43

3.1. INTRODUCTION 43

3.2. FACTUAL IGNORANCE AND NORMATIVE IGNORANCE 43

3.3. CULTURAL "BLINDNESS" 54

3.4. WHAT THIS MEANS FOR EPISTEMIC VICE VIEWS 63

3.5. CONCLUSION 64

Works Cited 66

CHAPTER 1

1.1. INTRODUCTION

Imagine I invite you over for tea. I ask you if you want any sugar in your tea and you say "yes, please!" So I take the sugar bowl out of my cupboard and spoon some sugar into your tea. Little do I know, someone had snuck into my house in the middle of the night and switched out the sugar in my sugar bowl with arsenic. You drink the arsenic tea and fall deathly ill. By putting what I think is sugar into your tea, I cause you to fall ill and potentially cause your death. But all along, I had no idea that the contents of my sugar bowl was arsenic, not sugar. Am I blameworthy for poisoning you?

The standard conditions required for moral responsibility are the control condition and the knowledge condition. The control condition claims that agents can only be blameworthy or praiseworthy for acts that they voluntarily perform. And the knowledge condition claims that agents can only be blameworthy for an action if they knew all the pertinent facts about their action prior to acting. But what are we to make of cases where wrongful agents fail to fulfill, specifically, the knowledge condition?

In this thesis, I aim to answer the larger question of "when should an agent have known better?" In answering this question, I hope to outline a criteria for which ignorance can be culpable. To do so, in Chapter 1, I explain exactly what the problem of culpable ignorance is and the main ways theorists have tried to solve it. In Chapter 2, I outline a particular debate between William FitzPatrick and Neil Levy. This debate aims to answer the question of "what can we reasonably expect agents to do to remedy or avoid their ignorance?" I defend the view held by William FitzPatrick, which argues that normative ignorance can be culpable if it is a result of epistemic vices such as arrogance, laziness, or dismissiveness, given that the agents *could have* reasonably

remedied their ignorance had they not exercised such epistemic vice, against Levy's view, which argues that what an agent can reasonably know and have done is a state *internal* to the agent herself. In Chapter 3, I defend FitzPatrick's view against Levy's by illustrating two false assumptions he makes in his argument. Particularly, the false assumptions are that responsibility for mere factual ignorance can be treated the same as responsibility normative ignorance, and that culture and upbringing can effectively "blind" agents from knowing how to act better.

1.2. WHAT IS CULPABLE IGNORANCE?

The standard conditions required for responsibility and blameworthiness are the control condition and the knowledge condition (Haji 530; Zimmerman 411-412; Rosen, "Skepticism About Moral Responsibility", 295). The control condition claims that agents can only be blameworthy or praiseworthy for an action if it was performed voluntarily. Consider the poisoned tea example again. It would be a different story if I was being held at gunpoint and forced to pour the "sugar" into your tea. Or if I was sleepwalking and unknowingly poured arsenic into my sugar bowl. Or if I had accidentally knocked over a vial of arsenic into my sugar bowl when I was in a hurry. Or if I was under a strong hypnosis and instructed to put arsenic into my sugar bowl. I would not be acting *voluntarily*, and it would seem unfair to judge me as blameworthy in that scenario. Both these situations violate the control condition, and hence changes judgements of blameworthiness. Now, if I put the arsenic in voluntarily, but only because I thought it was sugar, that would violate the knowledge condition, not the control condition. The knowledge — or the epistemic — condition holds that agents can only be blameworthy or praiseworthy if they know the pertinent facts about their action (Haji 530). Given that understanding, I would only be blameworthy for poisoning your tea with arsenic if I a) voluntarily mixed arsenic into your tea, and b) knew that I was putting arsenic into your tea and that arsenic is deadly.

However, it is not the case that agents always fulfill both conditions before acting. Oftentimes, an agent will not satisfy the knowledge condition. These cases are cases of agents *acting from ignorance* (Zimmerman 423-424; Rosen, "Skepticism About Moral Responsibility", 300). And while *ignorance* has various different meanings when used colloquially, I define *ignorance* as simply lacking knowledge, or justified true belief (Zimmerman 412; Peels 480).[1] From here, we can start noticing how ignorance puts pressure on our understanding of responsibility, and, more importantly, judgements of blameworthiness (and praiseworthiness[2]). Considering the basic understanding of responsibility, if an agent does not fulfill the knowledge condition, then she cannot be justifiably responsible for her action. However, there are often cases where an agent is ignorant, does something wrong, *and* seems blameworthy for having done it. But if ignorance, according to the tenets of responsibility, is inherently exculpating, then what are we to make of blameworthiness?

The problem with acting from ignorance is that it seems we can handle cases like the poisoned tea in two ways, neither of which are intuitively comfortable. On one end of the spectrum, it could be argued that since agents who unwittingly perform wrongful acts from ignorance do not know all the pertinent information before they act, and hence lack the knowledge condition, then they cannot be fairly held blameworthy for any acts committed out of ignorance (Smith 559; Rosen, "Culpability and Ignorance", 64; Zimmerman 414-415; Peels 577; FitzPatrick 591-593; Montmarquet 41-42). In other words, ignorance is always excusing. Surely, this conclusion appeals to our intuitions about moral responsibility because, after all, these actions do not fulfill

[1] Further investigation into what ignorance may be and the forms it may come in is irrelevant to my topic, but I acknowledge the complexity in my claim.

[2] While the tension exists for praiseworthiness as well, ascriptions of praiseworthiness in cases of agents performing morally wrong actions from ignorance is not my main concern. I will be focusing on morally wrong actions committed from ignorance.

all the requirements of moral responsibility. I might forget to tell you "happy birthday" or drive without knowing that the brake pads in my car are worn out or accidentally email confidential information to the wrong person. There are many wrongful acts performed from ignorance that would be unfair to judge as blameworthy. Perhaps I could have checked my calendar to see if it was your birthday and taken my car in for servicing more often and double checked the name of the email recipient, but it is harsh to say all these acts are blameworthy, especially since I did not fully know what I was doing (or neglecting to do) in each instance. However, suggesting that ignorance is *always* excusing would require us to do intense revisions to our intuitive understanding of responsibility. Sometimes an act is such that it is hard to find the agent exculpatory, even if she did not know what she was doing was wrong. There are times where it seems obvious that an agent is blameworthy for her unwitting wrongful actions despite the fact that she performed them from ignorance.

Thus, on the other end of the spectrum, it could be argued that even if the ignorance that is directly related to the wrongful act is culpable, it still does not provide any excuses for the ignorant agent; wrongful acts are always blameworthy regardless of fulfilling the knowledge condition (Smith 548; Haji 523). In other words, ignorance is never excusing. There is an intuitive aspect of this position, as well. There are times when it is clear that an agent *should have known better*, and her unwitting wrongful act performed from ignorance is blameworthy because she did not know what she *should have known*. Suppose a psychotherapist often slips up and uses his patients' stories as funny anecdotes with his friends. Or perhaps the controller of a profitable company unknowingly leaks confidential financial information to her competitors. Or perhaps a babysitter forgot that he left the child in his hot car while he was grocery shopping. Although these examples are all actions performed from ignorance, they no longer seem excusable. They seem to warrant

blameworthiness, at least to some extent, because in each one of these instances, it can be argued that they should have known not to do what they ended up doing. Ignorance becomes a poor excuse. However, with this view, we are never given the chance to deny culpability with ignorance. The defense of "I didn't know any better" no longer holds. Again, this view challenges our intuitive understanding of responsibility because there certainly are instances where ignorance can (and should) be excusing. Even if we think people should be blamed for unwitting wrongful acts performed from ignorance, at least sometimes, they should probably be blamed *less*.

There are some cases where we genuinely want to say ignorance excuses entirely. Suppose that the arsenic in my sugar bowl was there because someone snuck into my house and swapped it in. If this is the case, then no one should blame me: They should blame the person who swapped it. Here, ignorance seems to excuse entirely. A tempting response to this is that in many cases, blame is appropriate because the unwitting wrongful agent *should have known* not to do what she did. Clearly, if I unknowingly put arsenic into your tea, thinking it was sugar, is still a *bad thing* to do. But, it is not as bad as if I were to *intentionally* put arsenic into your tea, with the goal of poisoning you. It is inherently worse to do something with full knowledge of the action's wrongfulness than to do it with no knowledge of the action's wrongfulness.

Neither account that I have illustrated thus far satisfies this desire. We could either let ignorance be genuinely excusing in all cases, meaning no one can ever be responsible for not knowing. Or we could let every instance of unknowing wrongdoing be blameworthy, even though it often seems like it diminishes blame.

What we want is an approach that can show when it is justified to say that someone is blameworthy because they should have known, and when ignorance is genuinely excusing. The first approach is often called the Tracing View, which says that we can only be responsible for acts

done from ignorance when the ignorance can be traced to actions that were *not* done from ignorance. While not always (Smith 544-548), endorsers of Tracing Views often require us to hold significantly less people accountable for their actions (Zimmerman 412-418, Rosen, "Skepticism About Moral Responsibility", 300-304). The second course of action, which I will follow Holly Smith's lead and call the Conservative View (Smith 548), is also deeply concerning. We consider ignorance to be exculpating and excusing very frequently — just look at how uncomfortable it would be to blame me for mistakenly putting arsenic in your tea because I was unaware that someone snuck into my house and switched the contents of my sugar bowl.

With neither of the extremes seeming to reconcile the requirements of moral responsibility and our intuitive judgements about blameworthiness and responsibility, further investigation into the issue is prompted. The investigation into this matter is often based on whether ignorance can be culpable. That is, whether possessing ignorance *itself* ever be blameworthy, and if so, when. Let me breakdown possible answers to this question.

1.3. TRACING VIEWS

I begin with views that fall closer to the first end of the spectrum. These are often called Tracing Views (Peels 576, Fischer 531-533, Timpe 5-7). Tracing Views propose that agents cannot be blameworthy for wrongful acts performed from ignorance unless it can be traced back to a blameworthy act that was not done from ignorance. For many in the debate, this requires that it can be traced back to an instance of *clear-eyed akrasia*. *Clear-eyed akrasia* is acting in full knowledge that what you are doing is wrong, judging that you should not do it, but doing it anyway (FitzPatrick 590; Rosen, "Skepticism About Moral Responsibility", 306). More precisely, defenders of Tracing Views want to say that an action from ignorance can only be blameworthy if the ignorance can be traced back to a blameworthy action that was not done from ignorance (i.e.

an action that was done in full knowledge that it was wrong). For instance, suppose I have two identical sugar bowls. One is filled with sugar and the other is filled with arsenic. I knowingly put the arsenic bowl next to the sugar bowl, knowing that this is risky and could lead to inadvertent poisoning, and that I should store it safely elsewhere, but I store the two bowls beside each other anyways. Later, when you ask for sugar in your tea, I do not double check to see which bowl I am grabbing, and so give you arsenic instead of sugar in ignorance.

However, *clear-eyed akrasia* is not the only form that akrasia can take on. There are also instances of what is called *bleary-eyed akrasia*. While it is similar to the general idea of akrasia — that (tout court) I know something is wrong, I know I should not do it, but I do it anyway — *bleary-eyed akrasia* refers to instances where the knowledge that something is wrong is less obvious. It is the knowledge aspect that is bleared. For instance, suppose a CEO of a large company tells an employee to "get me the profits and don't tell me how you did it." Or perhaps the chief editor of a newspaper says to a journalist, "get me the story, I don't care if you have to hack phones." It seems that the CEO has an inkling that her profits will be ill gotten gains and that the chief editor knows that phone tapping is needed to get his story. Both agents are acting akratically, but their akrasia is only *bleary-eyed*.

The distinction I point out between clear-eyed and bleary-eyed akrasia is salient to the discussion of Tracing Views because endorsers of Tracing Views only contend that the former is culpable. A further distinction Tracing Views suggest is that genuine cases of clear-eyed akrasia are rare. They propose that it is, for the most part, unlikely for an agent to knowingly do something wrong while judging that they should not do it. This is not to say that agents cannot ever knowingly *and* non-akratically do something wrong. However, the basic idea is that it is actually quite rare for people to knowingly and intentionally do things that they believe are morally wrong. More

often, people think that their immoral actions are justified, even if they recognize that those actions are often wrong when done by someone else. The truth of their claim is beyond the scope of my discussion in this section, but it is pertinent to note here.

Given that these views look for instances in an agent's causal history where she acts in full knowledge, or clear-eyed akrasia — in other words, *not* ignorance — it would appear that they are suggesting that agents can *never* be blameworthy for wrongful acts performed out of ignorance. If you can only be blameworthy for acts performed in full knowledge, the view seems to be arguing that ignorance is always exculpating. While some proponents come to a conclusion that is somewhat adjacent to that, Tracing Views do not necessarily terminate in that way. The tracing that is central to the view looks for instances in an agent's past where they may have knowingly produced their own ignorance. The personal production of ignorance can itself be culpable, which would make any wrongful acts that stem from said ignorance culpable, as well. For instance, suppose I had two identical sugar bowls in my cupboard. It is true that one is filled with arsenic and one is filled with sugar, and I know this. It is also true that I do not have any identifiable information on the bowls to differentiate between the two. Now, suppose that when you tell me you want sugar and I grab a sugar bowl, I am unsure whether the sugar bowl I grabbed is the one that contains sugar. As the example goes, I put arsenic in your tea, which means I grabbed the sugar bowl containing arsenic. However, I am acting ignorantly because I am not sure of which sugar bowl I grabbed, and since my ignorance is a direct result of my act of knowingly not labelling the two bowls despite being aware of the risk, then my wrongful act of feeding you arsenic is culpable.

I acted akratically when I did not label the two sugar bowls; I judged that I should label the sugar bowls because one contains a commonly used spice and another is a deadly substance when

consumed, yet I did not. And since my akratic (and, thus, culpable) action of not labelling the two sugar bowls directly caused my unwitting wrongful act of feeding you arsenic, then my feeding you arsenic would also be culpable. My unwitting wrongful act of feeding you arsenic can be traced back to an instance where I acted akratically (and culpably), which also makes it culpable. By knowingly not labelling the two sugar bowls, I performed what some theorists call a *benighting act*. A *benighting act* is an "initial act, in which the agent fails to improve (or positively impairs) his cognitive position" (Smith 547). Benighting acts are culpable because an agent is knowingly creating her own ignorance by preventing herself from knowing the right thing to do in the future. In relation to the example, I am culpable for feeding you arsenic because that act is a direct result of my benighting act, which starts from a place of knowingly doing something wrong. So the thing that makes me culpable is my knowing wrongful act from which my ignorant act can be traced to. Thus, Tracing Theorists would conclude that ignorance is only culpable if the unwitting wrongful act itself comes from a place of clear-eyed akrasia or it is a result of a benighting act.

The Tracing View argument often resembles the following form:

1) If agent *A* is culpable for some action *S*, then *A* does *S* either from akrasia or from ignorance.
2) Cases of akrasia are highly exceptional.
3) So, if *A* is culpable for *S*, then in nearly all cases, *A* does *S* from ignorance.
4) *A* is culpable for an action *S* performed from ignorance only if *A* is culpable for that ignorance.
5) If *A* is culpable for an action *S*, then in nearly all cases, *A* performed *S* from culpable ignorance.

6) *A* is culpable for *A*'s ignorance only if *A* is culpable for some past action *T* from which that ignorance issued.
7) If *A* is culpable for an action *S*, then in nearly all cases, *A* is culpable for another culpable action *T* that preceded *S*.
8) If *A* is culpable for an action *S*, then in nearly all cases, *A* is culpable for an infinitely long chain of culpable actions, each of which precedes the other.
9) There are no infinitely long chains of culpable actions each of which precedes the other.
10) Therefore, for any person *A* and action *S*, *A* is hardly ever culpable for performing *S* (Peels 577).

Their main conclusion is that "[w]henever an agent acts from ignorance, whether factual or moral, he is culpable for the action if he is culpable for the ignorance for which he acts"[3] (Rosen, "Culpability and Ignorance", 64). In other words, an agent can be blamed for a wrongful act performed from ignorance only if the agent either a) acted akratically, or b) the ignorance is a causal upshot of a prior akratic act (Rosen, "Skepticism About Moral Responsibility", 307; Zimmerman 416).

At this point, it is also worth noting that Tracing Views can be co-opted to support skeptical arguments about responsibility. They are not skeptical by nature since there is a criteria for which ignorance can be culpable (Smith 559). But many theorists will use it to advance skeptical arguments. For instance, Gideon Rosen uses the logic of the Tracing View to support his radical conclusion that positive judgements of responsibility are never justified (Rosen, "Skepticism About Moral Responsibility", 295). In essence, his position argues that we can only blame agents for unwitting wrongful actions performed from ignorance if that ignorance is a case of or is the

[3] The difference between factual and moral ignorance will be explained in depth in the second section. It is pertinent to point out now, though, that all the examples in this section fall under the category of *factual ignorance*.

causal upshot of a case of akrasia. However, akrasia is both rare and difficult to identify. So, we are never in a justifiable position to blame agents who perform unwitting wrongful acts from ignorance. In tandem, Michael Zimmerman also contends that culpable ignorance occurs far less frequently than commonly supposed (Zimmerman 411). This skepticism runs into the problem I outlined earlier in my introduction. Namely, if it is true that we should always suspend judgement in cases of unwitting wrongful acts performed from ignorance, or that we should at least accept that far less people are culpably ignorant than we suppose, then we would have to conclude that ignorance is, for the most part, excusing. However, there are numerous cases where we seem to have justifiable reasons to blame agents for their unwitting wrongful acts performed from ignorance even when it does not trace back to an instance of akrasia. To accept the skeptical argument proposed by Tracing Views, then we would also have to engage in extreme revisions to our common understanding of responsibility. And when it comes to extreme revisions to concepts such as responsibility, it should be done cautiously.

1.4. CONSERVATIVE VIEWS

As mentioned, it seems that if we accept Tracing Views to be true, then we would not be able to hold agents blameworthy when we normally would. This is an uncomfortable outlook for various theorists. I will proceed to present three dimensions of the culpable ignorance debate that fall closer to the Conservative View and that find ways to reconcile our intuitions about blameworthiness and moral responsibility without arguing that it is unfair to blame the majority of ignorant wrongdoers. The answer to "can ignorance ever be culpable" to many of these theorists is "yes."

I am using Holly Smith's categorization, but I am using it loosely. Based on her definition, Conservative Views suppose that "the agent's culpable ignorance provides no excuse for his

performing the unwitting wrongful act; he is fully blameworthy" (Smith 548). Her definition concisely summarizes the opposing end of the argument where every wrongful act is culpable, regardless of an agent's ignorance. However, theorists who pose Conservative Views would not go as far as to say *all* wrongful acts done from ignorance are culpable. Despite the difference between the theories proposed and Smith's definition, these theories that I consider Conservative fall into this category because they all want to say that you can be responsible for ignorance that does *not* trace back to knowing wrongdoing. They argue against the revisionism and regressive structure of Tracing Views, but with various different methods. Thus, the debate is about the criteria – determining which cases and causes of ignorance are blameworthy. In general, the approach here is to say people can be blameworthy if they *should have known* even if they did not. These accounts owe us an explanation of the standards of "should."

1.5. EPISTEMIC VICE VIEWS

One dimension of responsibility contends that ignorance is culpable if the ignorance is rooted in one or many epistemic vices (FitzPatrick 609; Montmarquet 45). These theorists do not necessarily disagree with the conclusion presented by Tracing Views. It is perfectly acceptable that akratic acts are culpable (FitzPatrick 609). However, proponents of Epistemic Vice Views reject the idea that awareness is required for responsibility. While Tracing Views rely on akrasia, and akrasia requires an awareness that an action is wrongful and a judgement to not perform that action, Epistemic Vice Views hold that unwitting ignorant wrongdoers can still be culpable even if the benighting act that produced their ignorance was also performed from ignorance. In other words, an agent need not know that what she is doing is wrong in order for her to be justifiably blameworthy for her unwitting wrongful act.

Additional factors need to be considered when judging an agent's blameworthiness. Specifically, if an agent was blinded to the reasons why she should or should not perform a certain act because she exercised an epistemic vice such as intellectual irresponsibility, overconfidence, arrogance, dismissiveness, laziness, dogmatism, incuriousity, self-indulgence, or contempt, then that agent may be culpable (Montmarquet 45; FitzPatrick 609). These views hold that even if an agent is unaware of what she is doing, she may still be blameworthy because exercising epistemic vices that prevent her from knowing and performing her epistemic obligations is an intellectually irresponsible and culpable act. Thus, culpability does not necessarily terminate with akrasia. Epistemically vicious behaviour can also make an agent's unwitting wrongful act blameworthy. Blame does not trace back to a specific blameworthy wrongful action, but instead emerges to a morally blameworthy epistemic character or disposition.

For instance, suppose a doctor is assessing which medication will treat her patient's illness. Although recent studies have been published that say medicine X is better at treating the patient's specific illness, the doctor is unaware because she has not read the recent studies because she is over-confident in her own judgments. Due to this over-confidence, the doctor gives her patient medicine Y, and the patient's illness is not treated. Had the doctor not been over-confident in her own judgment, her patient would not have been harmed. Epistemic Vice Views would consider the doctor blameworthy because she is culpably ignorant due to an epistemic vice.

Or suppose a hiring manager refuses to hire women under the age of 40. His refusal to hire this demographic is because he believes that women of childbearing age make worse employees since they will get pregnant and take maternity leave. His belief is rooted in reading the one business magazine subscription he has that states time and time again that women make bad employees because they take maternity leave. And he only has one business magazine subscription

because he is too lazy to do the research required for finding other magazines he is interested in reading.

Or perhaps a fashion enthusiast has become complacent in following fashion trends. She unthinkingly buys the clothes that she sees her favourite celebrities wearing and the newest shoes on the feet of runway models and the newest drops from her favourite stores. Her complacency in following fashion trends leads to an ignorance surrounding the environmental and social effects of fast fashion. She contributes to the fast fashion market that harms people, economies, and environments, but she is so complacent in her enthusiasm for shopping that she continues to fund harmful practices.

Arguably, in each case the ignorance is blameworthy, not because it emerges from a knowing case of wrongdoing, but because it reflects a lack of epistemic responsibility more generally: Each agent characteristically fails to be curious, to take time to double check their informative sources, and to challenge their views.

Investigation into Epistemic Vice Views begins with asking, generally, *can possessing ignorance itself be culpable?* William FitzPatrick interprets this question to be asking "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?" His answer is that an agent is culpable for an unwitting wrongful act performed from ignorance if she violates an epistemic obligation that she could reasonably have been expected to follow because she exercised an epistemic vice. FitzPatrick establishes an agent's direct culpability for the culpable formation of epistemically irresponsible beliefs. He outlines the following three criteria for culpable ignorance:

1. There has to be no relevant limitations to the agent's social context that may have made broader reflection and information gathering unreasonably difficult for her.

Given this requirement, extraordinary cases such as those brainwashed by an isolated cult would not be able to reasonably reflect broadly and easily gather information (FitzPatrick 605).

2. The failure to adequately reflect and gather information must have been the result of a voluntary exercise of an epistemic vice (FitzPatrick 605).
3. And, the agent could have reasonably been expected to take steps that would have eliminated her ignorance by refraining from exercising the epistemic vice (FitzPatrick 605).

James Montmarquet interprets the original question as asking whether culpability of a belief-state can sometimes be fundamental and underived (Montmarquet 41)? He takes his answer a step further than FitzPatrick in that he believes that, beyond *forming* epistemically culpable beliefs, merely holding epistemically culpable attitudes can, itself, be culpable. He contends that "if an agent is culpably ignorant of some fact, there will always have been some possible act or omission, but for which this ignorance would not have existed" (Montmarquet 42). And since, he argues, agents exhibit direct (although incomplete) control over their actions, then there would have *always* been an instance for them to have known better and, hence, done better. The true source of culpability is in an agent's epistemically irresponsible attitude.

Epistemic Vice Views do not require akrasia for culpable ignorance. It is a weaker standard than Tracing Views because it does not demand as much, and more unwitting ignorant agents who act from ignorance would be considered culpable. These views have also not been spared criticism. A notable one is that, when looking at what agents can do to better their epistemic position, we must consider the perspective of the agent, not what is in general the *right* thing to do. Such criticism leads us into the second Conservative View I will discuss.

1.6. INTERNALIST VIEWS

Internalist Views do not agree with Epistemic Vice Views that epistemic vices are what make ignorance culpable. Internalist Views hold that proponents of Epistemic Vice Views are asking the right question, but come to the wrong answer. If we are to answer the question of what an agent can *fairly* have done in their past to avoid or remedy their ignorance, the fair answer would be *nothing*: "there is nothing that [an agent] could reasonably have been expected to do to avoid or remedy his normative ignorance" (Levy 732). Internalists contend that "blameworthiness is affiliated not with the objective wrongness of an action but with whether a person takes herself to be doing wrong in performing an action" (Haji 526). When we centre the question around culpability, we are actually concerned with an agent's capacity (Levy 735). Neil Levy claims that "agents are culpable for actions if it is reasonable to expect them to conform their behaviour to the appropriate normative standards, and this expectation is in place only if they have the capacity to behave accordingly" (Levy 735).

Internalist Views hinge on another intuitive aspect of moral responsibility: Internalist rationality. Here, internal is meant to refer to "the reasons which explain what it is rational for the agent to do, not with what she has reason to do" (Levy 735). What an agent has the most reason to do is, for the most part, the *right* thing. However, an agent may not have the proattitudes that would prompt her to do the reasonable, right thing. An agent can only act according to her internal beliefs. Thus, an agent is morally blameworthy for performing an action only if she has the belief that the action is wrong and if this belief plays an appropriate role in the agent's acting (Haji 529). It would be unreasonable to hold agents blameworthy for actions they could not avoid through some reasoning procedure (Levy 739).

To illustrate the significance of internal reason, and why Internalist Views deem it important, I will reproduce two cases of differing internal reason from Ishtiyaque Haji's (529-529) work. First, suppose Doctor A is a caring physician and wants the best for all her patients. When treating Patient X for an illness, she is under the belief that Medicine S will cure her patient. However, and unfortunately, Doctor A is unaware that administering Medicine S is the wrong thing to do and ends up killing Patient X. In the second case, suppose Doctor B is malicious, hateful, and does not want the best for her patients. When Doctor B is met with Patient Y, she is under the belief that Medicine T will kill her patient and so she administers Medicine T. However, Medicine T was the appropriate treatment for her patient and Patient Y is cured.

Although Doctor B did the right thing and Doctor A did the wrong thing, it seems inappropriate to say that Doctor B is praiseworthy while Doctor A is blameworthy. Rather, it should be the opposite. Doctor A did her best to ensure her patients lived, while Doctor B should probably have her doctor's license revoked. An agent's blameworthiness depends on the agent's internal state. If the agent is unaware that she is doing wrong, such as Doctor A, then it is unreasonable to blame her because based on her own rationale, she was doing the right thing.

While this argument sounds like Tracing Views — indeed, Levy concludes his argument by stating that Rosen was right all along (Levy 741) — I understand Internalist Views as falling closer to Conservative Views. Unlike Tracing theorists, Internalists are not thinking primarily about akrasia. Rather, Internalists look to what an agent can *reasonably* have done or known. It is only when an agent could not have done otherwise based on her own reasoning do Internalists conclude that agents are only blameworthy for wrongful acts committed from ignorance if the agent has the belief that the action was wrong. Further, unlike Tracing theories, Internalists decentre *objective* rights and wrongs. Rather, they focus on an agent's internal understanding.

Whether an agent knows the objective rightness or wrongness of an action is not of concern. Whether an agent can conform her behaviour to the requirements of morality is crucial to her blameworthiness. The beliefs an agent has about those actions is what factors into her blameworthiness.

Another aspect that makes Internalist Views different from Tracing Views is that it lacks the same regressive structure (Montmarquet 43, Peels 576). Since the regressive structure of Tracing Views is not foundational to my purposes in future chapters, I will only briefly discuss it here. Tracing Views suffer from a regressive structure in all cases except akratic ones because they aim to trace ignorance back to an akratic act. However, since Internalist Views are not working from a place that necessitates tracing akrasia, they do not suffer from the same. Although it appears that Internalist Views are concluding that, once again, only akratic actions can be blameworthy (which they are), they arrive at this conclusion by considering a weaker, internal awareness of the agent. If an agent could not be reasonably expected to have known that an action was wrongful, then the agent's ignorance is exculpating. We are not required to look back into the agent's history to see if her ignorance was the causal upshot of a prior akratic act or omission, as we would be in Tracing Views.

1.7. QUALITY OF WILL VIEWS

The final group of views I will be discussing also does not believe that akrasia is a requirement for culpable ignorance. These views focus on the idea that an agent's blameworthiness is grounded in the quality of will from which she acts. While each theorist defines quality of will slightly differently, it can generally be described as: "an agent is praiseworthy for an action or an attitude that accords with the demands of morality if the performance of the action or the holding

of the attitude arises from proper regard or concern for another person's morally significant interests. Conversely, an agent is blameworthy for an action or attitude that conflicts with the demands of morality if the performance of the action or the holding of the attitude arises from lack of proper regard or concern for those same interests" (Rudy-Hiller). Here is another group of theories that believe culpable ignorance does not require akrasia. Further, it is similar to Epistemic Vice Views in that it does not require akrasia or any direct awareness of the wrongfulness of an action for culpability. A weaker standard is sufficient to establish culpability. And the standard here is an agent's quality of will.

The most prominent theorist within Quality of Will Views is Nomy Arpaly. Her main assertion is that an agent's praiseworthiness is determined by how an agent responds to moral reasons, where moral reasons are those that make an action right (Arpaly 71-84). Acting upon these moral reasons, then, is what makes an action praiseworthy. The same can be said about the contrary. An action is blameworthy when an agent acts upon the wrong making features of an action. To act with good will is praiseworthy while acting with ill will is blameworthy. There is a further distinction that Arpaly draws between cases of blameworthy actions. An agent can act with ill will, or she can act with deficient good will. To illustrate the difference, let us, once again, return to the specific version of the arsenic sugar bowl example where I have two identical sugar bowls — one with sugar and one with arsenic. Once again, supposing that I know that there are two identical sugar bowls, but only one contains sugar, the right thing to do would be to double check if the sugar bowl I am grabbing is, in fact, the one that contains sugar. And this is the right thing to do because I care about your wellbeing. Feeding you arsenic would be antithetical to that. However, when I do not double check, which leads to my false belief that it is okay to spoon the contents into your tea, I am displaying a deficit of good will. Although I did not act with the

intention of harming you, I still am blameworthy because I lacked sufficient regard for causing you harm, and my ignorance is culpable because it revealed that I was not sufficiently sensitive to the relevant moral reasons at hand. If I had acted with the intent of harming you, then I would be displaying ill will.

A benefit of this view is that it shows how some cases of ignorance can be blameworthy because they involve a lack of sufficient moral concern (i.e. good will). But they also acknowledge that these cases are typically much less blameworthy than knowingly doing something wrong. For instance, intentionally poisoning someone shows ill will, which is worse than not displaying enough good will.

What is relevant in the cases of deficient good will and ill will is that, in both, I display substandard consideration for the moral reasons at hand. I am unresponsive to you and the moral reasons you give to my actions.

Matthew Talbert also presents an argument for the Quality of Will View, which begins by arguing that *moral competence requirements* do not need to be satisfied for an agent to be blameworthy (Talbert 90). These are the psychological, emotional, and motivational resources agents must have in order to respond to the moral considerations of their behaviour (Talbert 89). Rather, he contends, ignorant agents may still be blameworthy, despite their lack of moral competence because moral competence is not what separates the blameworthy and the excused (Talbert 90). The blameworthiness of an act truly comes from an ill quality of will because we have a "tendency to judge that our needs and interests should have normative significance for other people, and our corresponding concern with whether others treat our needs and interests as giving them reasons to constrain their behavior in various ways" (Talbert 94). When others fail to meet these concerns — when they fail to consider our needs and interests as having normative value —

then that is when we may justifiably blame them. And we are justified in our blame because the unwitting wrongful agent failed to take in that the welfare of other people is an action-generating reason to perform or refrain from a certain course of action.

For instance, suppose my birthday is a really important event for me. I enjoy celebrating my birthday with people I love, I enjoy making the day special, and I enjoy the comfort in knowing that there is one dependable day in the year where I will hear from all those I care about. Also suppose that I share this information with my partner. However, when it comes around to my birthday, I not only fail to hear "happy birthday" from him, but when I bring up to him that he might have forgotten what day it is, I am told "sorry I forgot, but it's not a big deal anyway" For Talbert, I would be justified in blaming my partner because his forgetfulness and nonchalant dismissal of my birthday is displaying an ill quality of will. He neglected to recognize that my interest in my birthday should generate a reason for him to (among other things) remember to tell me "happy birthday."

The definition of *blame* plays an important role here. Talbert appeals to P. F. Strawson's characterization of blame, which is defined as the reaction of the "quality of others' wills towards us, as manifested in their behaviour: to their good or ill will or indifference or lack of concern" (Strawson 83). Thus, to say that someone is *blameworthy* is to say that they can be fairly targeted with the negative reactive attitude of blame. Contrarily, to deny the judgment of blameworthiness to an agent's wrongdoing, then it would be unfair for victims to express the negative feelings involved in blame (Talbert 97). Considering the example just presented, what would normally be a reason to be upset and harbour feelings of blame turns into one that mischaracterizes my experience. The significance of my partner's actions would be lost. And the significance would be that my partner acted from an ill quality of will because, had he considered that my love for my

birthday was an interest that he should factor into his decision making, he would have remembered. But since he failed to take into account my interests and displayed an ill quality of will, he is an adequate target of blame.

1.8. CONCLUSION

When Rik Peels joined the conversation on culpable ignorance, he gave a warning against accepting Tracing Views: "we should be very suspicious of philosophical arguments that call for a substantial revision of one of our socially most important and deeply entrenched practices. We should take this route only if each of the premises of the argument is at least as convincing as the thesis that people are frequently blameworthy for their ignorance and for actions performed from ignorance" (Peels 578). I could not have expressed his warning better myself. Tracing Views ask us to reconsider *who* we blame and *how* we do it. They call for a more demanding and stricter criteria for responsibility than we typically resort to, and such revisions should not be taken without careful consideration.

I contend that the criteria Tracing Views provide is far too strong to accurately categorize culpable ignorance. Thus, we need a weaker standard for culpable ignorance. In particular, the problem with Tracing Views is that in some cases it seems reasonable to think that someone *should have known* something even if their ignorance cannot be traced directly to a clear-eyed akratic act, but rather, it can be traced to an epistemic vice or a lack of good will or moral concern.

The next chapter aims to distinguish between two competing approaches to this. In the next chapter I will defend Epistemic Vice Views. Specifically, I will defend FitzPatrick's version of the epistemic vice criteria for culpable ignorance. In doing so, I will delineate his argument and criticisms his view faces, particularly from Levy. The next chapter of my paper is dedicated to outlining the debate between these two theorists.

CHAPTER 2

2.1 INTRODUCTION

In this section, I will explain the debate between William FitzPatrick and Neil Levy. The debate begins with FitzPatrick's attempt to allay the worries that Tracing Views present to moral responsibility by introducing the Epistemic Vice View. In essence, he argues that agents are culpable for unwitting wrongful acts performed from ignorance when they act upon epistemic vices, which is an epistemically vicious and reproachable thing to do. If an agent *could have* remedied their ignorance, but chose not to because she exercised an epistemic vice, then she would be culpable for any wrongful acts she performs due to that ignorance. However, Levy believes that FitzPatrick is mistaken, and he proposes his Internalist View to illustrate why. In particular, Levy believes that we are only justified in blaming agents for what they *could have known*, which relies on an agent's internal reasoning. If an agent *could not* have known better, then the exercise of epistemic vices should not count them as culpable.

To do so, I will first introduce and explain Gideon Rosen's skeptical position. Then, I will introduce FitzPatrick's Epistemic Vice View, followed by criticisms Levy has. Levy's criticisms lead into his Internalist View, which I will also outline. I conclude by arguing that Levy's criticisms appear correct to me, but I am unwilling to lend him full acceptance. Instead, I will set forth in the next section by defending FitzPatrick's view against Levy's criticisms.

2.2. BACKGROUND: ROSEN AND SKEPTICISM ABOUT MORAL RESPONSIBILITY

While my main question does not directly address Rosen's position, both theorists I am drawing on make ample reference to Rosen's work. Thus, it would be remiss for me to begin

describing the debate without first mentioning Rosen, who is a starting point for both FitzPatrick and Levy.

Rosen's most radical conclusion, the one FitzPatrick is concerned by and Levy endorses, is centred around what he calls the *skepticism about moral responsibility*. The argument begins with a distinction between original and derivative responsibility, where original responsibility means that an agent is *directly* responsible for an event occurring and where derivative responsibility means that an agent is causally responsible for an event occurring by virtue of being directly responsible for a previous act. Think of a set of dominoes where the set up entails that by knocking over domino 1, domino 2 will fall and knock over domino 3, and so forth. If I were to knock over domino 1, I am *originally responsible* for its falling over, but I am only *derivatively responsible* for domino 2 falling over. Perhaps the set of dominoes has 100 dominoes lined up. In such a case, I am only *originally responsible* for domino 1 falling over, but due to the consequence of my original responsibility, I am derivatively responsible for dominoes 2 through 100 falling over. Due to the causal history of events that led to dominoes 2 through 100 falling over, it can always (let's assume) be traced back to my knocking over domino 1.

Rosen proposes that for any action, an agent can either be originally responsible or derivatively responsible. For there to be either sort of responsibility, there requires a locus of responsibility (for example, the act of knocking over domino 1). Now, an agent can act out of *clear-eyed akrasia*[4] — against her better judgement — and she would be culpable for it because this involves knowingly and intentionally doing something that the agent judges to be morally

[4]Again, *clear-eyed* akrasia is not simply knowing that something is wrong and doing it anyways. To put it into context, as applied to the sugar bowl example, it is not just knowing that the sugar was arsenic, but serving it anyway that would make it *clear-eyed*. It is knowing it is arsenic, judging and intending *not* to serve it, and then knowingly doing so anyway. In other words, akrasia is not just knowingly doing something morally wrong; it is knowingly doing something morally wrong while judging that you should not do it or while intending not to do it.

wrong. It therefore satisfies both the control and the knowledge condition for responsibility. As I mentioned in the first section of this paper, endorsers of Tracing Views, including Rosen, contend that only *clear-eyed akrasia* can qualify as originally responsible:agents must know that what they are doing is wrong, judge that they should not do it, but proceed to do it anyway. Returning to the sugar bowl example in the first section, suppose I had full knowledge that the contents of my sugar bowl was arsenic and not sugar, but I served it in your tea anyway. There is no question that my act is culpable. Anything short of full knowledge of all the relevant facts of an action and the judgement to avoid doing it would not be judged culpable in Rosen's eyes.

Up until this point, there is no direct mention of actions from ignorance. Ignorance poses further problems to conventional understandings of moral responsibility. There needs to be some sort of anchor for what makes unwitting wrongful actions performed from ignorance blameworthy. Rosen sets forth the following intuitive answer as a general principle for culpable ignorance: "*If X does A from ignorance, then X is culpable for the act only if he is culpable for the ignorance from which he acts*" (Rosen, "Skepticism About Moral Responsibility", 300). Ignorant actions can only be blamed if the ignorance itself is culpable, either originally or derivatively.

Due to his criteria for culpable ignorance, Rosen contends that culpable ignorance can never be a locus of original responsibility — culpably ignorant acts are always derivatively responsible. To be derivatively responsible for an ignorant act, the agent must have failed to fulfill her procedural epistemic obligations somewhere in its causal history, to wit: "[W]henever you are culpable for failing to know some pertinent truth, this is because you are culpable for failing to comply with one of these procedural epistemic obligations" (Rosen, "Skepticism About Moral Responsibility", 302). Rosen does not argue that an agent must know *every* fact about the situation prior to acting. Rather, he claims that agents are bound by the standing obligation that we do not

cause unjustifiable harm to others as we move through the world. Such an obligation does not entail full knowledge of all pertinent facts about our actions — to enforce that requirement would be impossible to satisfy and unreasonable to implement into a condition of moral responsibility. It is an obligation to take all the *reasonable* steps to ensure that we would know what the right course of conduct is if we find ourselves in such a situation. The relevant procedural epistemic obligations will vary between different cases, but it generally involves obligations to do certain things — "to ask certain questions, to take careful notes, to stop and think, to focus one's attention in a certain direction" (Rosen, "Skepticism About Moral Responsibility", 301). However, there is no guarantee that at the moment of acting I have a full understanding of what the right thing to do is. I may have fulfilled all my procedural epistemic obligations, but failed to fully grasp a crucial concept in the midst of acting. Rosen would say that such a failing is not a culpable one.

Ignorance is only culpable when an agent fails to satisfy her procedural epistemic obligations, and procedural epistemic obligations are only unsatisfied due to "recklessness or negligence in the management of one's opinion" (Rosen, "Skepticism About Moral Responsibility", 302). Returning to the main question posed in the first section — *when, if ever, is ignorance culpable?* — the answer that has received the most support is: *When the agent should have known*. While Rosen does not directly answer the question of *when an agent should have known*, he does suggest that an agent should have known to satisfy her procedural epistemic obligations. If the agent negligently and recklessly abandoned her procedural epistemic obligations, thus leading to her ignorance, then she would be culpable. The agent would have to knowingly fail to satisfy her procedural epistemic obligations, which she knew she had a duty to do. Once again, ignorance can only be derivatively culpable. For instance, perhaps I keep an identical bowl of arsenic beside my sugar bowl in the cupboard. It would be a procedural epistemic

obligation of mine to double check that I am grabbing the bowl of sugar rather than the bowl of arsenic before spooning it into your tea. Although I know of such an obligation and judge that I should, I still do not double check, which produces my ignorance of whether the "sugar" I spoon into your tea is safe for you to drink, and I would be culpable for this ignorance. Rosen concludes that "X is culpable for failing to know that P only if his ignorance is the upshot of some prior culpable act or omission" (Rosen, "Skepticism About Moral Responsibility", 301).

To summarize Rosen's position: 1) actions from ignorance can *only* be derivatively responsible; 2) so the ignorance needs to be directly responsible; 3) so there needs to be some knowing wrongdoing at the source of the ignorance if it is to be blameworthy; 4) to find a knowing wrongdoing at the source of the ignorance involves finding an instance in an agent's causal history of knowingly violating a procedural epistemic responsibility; 5) for Rosen, this can only be explained by a case of clear-eyed akrasia.

I spell his position out here because his final step (the one that stipulates clear-eyed akrasia as a requirement) is controversial. And it is controversial because it depends on the idea that no one can ever form the intention or knowingly do something morally wrong, and this is not a universal view. Rosen sets himself apart from other theorists that also endorse Tracing Views, but who do not make claims about whether it is possible to intentionally do something morally wrong.[5] Thus, I stress that Rosen's skeptical argument relies on this controversial notion that is not commonly held by others in this field. Indeed, FitzPatrick goes on in his argument — which I will explain further in this chapter — criticizes Rosen's skepticism about the frequency at which agents actually do act akratically (FitzPatrick 595-599).

[5] This claim is not one I want to address in this thesis, but the distinction is pertinent here.

Various other tracing theories have been posed that follow Rosen's reasoning (Zimmerman 412-418; Smith 555-563), but none have as radical of a conclusion. Rosen concludes that, for the most part, we, as onlookers, do not know whether an agent acted from clear-eyed akrasia or genuine ignorance, so we are never in a position to judge whether ignorance is culpable. There is extreme difficulty in confidently identifying when someone has acted akratically, even in instances of our own actions. Since it is necessary that culpable ignorance can be traced back to a case of akrasia, and due to the difficulty in attributing akrasia, Rosen claims on epistemic grounds that it is never justified to hold confident positive judgements of responsibility.

2.3. WHAT AGENTS CAN REASONABLY HAVE DONE: THE DEBATE BETWEEN FITZPATRICK AND LEVY

FitzPatrick worries about Rosen's skeptical conclusion, namely that it is too demanding of a conclusion for moral responsibility and that agents can be culpable for their ignorance in the absence of akrasia. It is important that I draw a distinction here between facts that agents can be ignorant of. An agent can exhibit *merely factual ignorance*. Thus far, all my examples concern factual ignorance where an agent of reasonably sound mind and judgement acts wrongfully due to not knowing or believing a mere factual statement. The sugar bowl example pertains to mere factual ignorance (I do not know that the contents of the sugar bowl is arsenic, not sugar). Examples of *merely factual ignorance* are rather clearcut. There is a very obvious piece of truth that the agent is ignorant of, which causes her to act wrongfully. And it is just as obvious to consider how easily the agent could have known it. Some facts are easily knowable, through occupational standards or a Google search or experience. For instance, someone who has completed all the schooling and residency requirements to become a doctor should know that cannot divulge confidential patient information to anyone you please, and a first time dog owner

can easily Google what foods are safe for dogs to eat. Sometimes, facts are trivially easy to discover. Other times, facts are impossible to discover. The same doctor would not be culpable for not knowing how to treat a disease that has never been seen before, and the first time dog owner would not be culpable for not knowing that her dog is allergic to beef if an allergic reaction had never happened before. Mere facts sit at two opposite ends of knowability: Either trivially easily knowable or impossibly knowable. And while discussions of *merely factual ignorance* have far-reaching implications (in business ethics, legal studies, and medical ethics, to name a few), cases of factual ignorance are nevertheless less complicated because how easily an agent can come to know a fact can be determined with more ease.

Conversely, agents can be *normatively ignorant*. I have intentionally neglected to touch on instances of normative ignorance thus far as I set up my argument, but it will become a larger focus in this section and those to follow. Agents are normatively ignorant when they do not know of certain *moral* facts. To briefly illustrate the difference, let me pose a common example in the literature: Consider the Ancient Roman Slaveholder who makes his slaves fight starving lions in large arenas for the entertainment of himself and the people of Ancient Rome. The Ancient Slaveholder and the people of Rome cheer the lion on and rejoice when the slave inevitably loses the fight (because even after watching *Gladiator*, I can admit that it is virtually impossible for an unarmed man to win a fight against an angry, hungry lion). For all the Ancient Slaveholder knows, he is doing a good thing by bringing people joy during these events. Moral theory has not progressed to the point where everyday people consider slavery and sadistic fights against lions to be morally wrong, and if it has, it surely has not been communicated to the Ancient Slaveholder. Then, the obvious question of whether the Ancient Slaveholder is culpable for his ignorance pops

up immediately.[6] If the Ancient Slaveholder is not culpable, then it seems as if his ignorance should excuse him from blame, even though the action of keeping people in slavery in order to have them eaten by lions for the sake of entertainment is profoundly morally wrong. If the Ancient Slaveholder did not know that it was wrong, and if he is not to be blamed for his ignorance, then it might seem as if he should not be blamed for the morally abhorrent actions he performed. But if he is culpable for his ignorance, then we need an account of why.

FitzPatrick's work focuses intentionally on normative ignorance, hence why such a distinction here is needed. And this is because normative ignorance is more difficult to assign a level of ease of knowability. Considering the Ancient Slaveholder example, it is not the case that the Ancient Slaveholder has been told or taught that making slaves engage in depraved fights against lions is morally wrong. But it is also not the case that the Ancient Slaveholder would be entirely unresponsive to an argument as to why it is wrong. An agent could tell the Ancient Slaveholder that the slave he is putting in these fights feels pain just as much as he does, and since he would not like it if he were put in an arena to face off a lion, then he should not do it to others. Although it could be argued that it would be difficult for the Ancient Slaveholder to understand and appreciate such an argument, it would not be the case that it would be *impossible* for the Ancient Slaveholder to understand and appreciate such an argument. The normative structure of why making slaves fight against lions is available to him, and the structure that would work for the Ancient Slaveholder would be rather similar to the structure I would use in contemporary arguments. I stress this similarity because there is a sense in which normative arguments are not

[6] The example of the Ancient Slaveholder has been referenced by various theorists within the field. See: Rosen, Gideon. "Culpability and Ignorance." *Proceedings of the Aristotelian Society*, vol. 103, 2003, pp. 61–84, www.jstor.org/stable/4545386.; Moody-Adams, Michele M. "Culture, Responsibility, and Affected Ignorance." *Ethics*, vol. 104, no. 2, Jan. 1994, pp. 291–309, www.jstor.org/stable/2381578.; Harman, Elizabeth. "Does Moral Ignorance Exculpate?" *Ratio*, vol. 24, Dec. 2011, pp. 443–468.

bound by the same limitations of knowability of, say, science or technology or experience. In a sense, normative arguments are available to all competent agents regardless of their situational binds.

As I mentioned previously, FitzPatrick is uncomfortable with Rosen's conclusion, deeming it too demanding. In arguing against Rosen's *skepticism about moral responsibility*, FitzPatrick begins by denying his claim that it is always difficult to know whether someone performed a wrongful act akratically. Rosen exaggerates the difficulty in attributing akrasia, FitzPatrick argues. From a personal perspective, giving into the gustatory temptation of eating ice cream while on a low-fat diet or procrastinating a project with a looming deadline are common examples of very blatantly acting akratically. Further, he provides the presence of guilt or shame as an indicator of acting akratically. If an agent genuinely believed she was doing the right thing in the moment of acting, only to find out later she was operating ignorantly and did the wrong thing, then she should not feel any guilt or shame for her actions. Alternatively, if the agent does feel shame, then it can be attributed to akratic action since "it will often be clear that one does not take oneself to be in a moral dilemma, that one does endorse standards that condemn the action, and that the guilt or shame in so acting stems precisely from that" (FitzPatrick 596). And finally, he argues that agents will, more often than anticipated, be honest about their failings. He provides the example of how a husband who cheats on his wife would be more likely to tell the truth ("I'm sorry for cheating") rather than claim circumstantial ignorance ("I didn't know it wasn't you!) or factual ignorance ("I didn't know I wasn't supposed to cheat") (FitzPatrick 596). While these are only heuristics and FitzPatrick himself does not fully endorse that every instance of breaking a diet or feeling guilty shows akratic action or cheating spouse is a case of akratic action, these indicators nevertheless illustrate how Rosen is too firm on his claim that akrasia is always too difficult to identify.

With all this said, FitzPatrick has yet to allay his worries about Rosen's position. To do that, he must prove his overall point that agents can be culpable for their ignorance even when it is not the result of an akratic action; that a lot of the time "[b]ad actions resulting from an agent's moral ignorance are among those we typically regard as paradigm cases of moral responsibility, yet upon reflection they seem often to fail the akrasia-involving condition and to do so all the way back" (FitzPatrick 590). He aims to separate akrasia from reasonability in ascribing culpability by presenting what he calls the R Condition: "What, if anything, could the agent reasonably (and hence fairly) have been expected to have done in the past to avoid or to remedy that ignorance?" (FitzPatrick 603). While ignorance is not always culpable, FitzPatrick argues that it usually is in cases that concern us, and the reason why is because the condition for culpability is weaker than what Rosen proposes. Culpable ignorance is identified by what an agent can be reasonably and fairly blamed for, which is a weaker standard because more unwitting ignorant agents would be considered blameworthy for their wrongful acts. The standard is easier to meet because it is often more likely that agents *could have known better* than that they actually *did know better*.

The next step is to then uncover what agents can reasonably be blamed for. To help answer his question, FitzPatrick uses the example of Mr. Potter, the ruthless capitalist from *It's a Wonderful Life*. Potter is:

> "[A] powerful businessman who holds false moral views. He takes certain business practices — such as liquidating Bailey's Building and 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." (FitzPatrick 599-600).

Potter's normative ignorance is not unique by any means — he was not raised in a cult or brainwashed. Presumably, he grew up in a good family household, received a good moral education, and has been exposed to opposing moral views that challenge his ingrained understanding of American capitalism and he understands them as such. Presently, he operates in a world that has various moral views — he lives in a town whose population has varying views, some of which oppose his own; he reads newspaper articles containing arguments for his views and opposing ones; he watches news broadcasts that deliver multiple points of views on a particular situation — but he still dismisses each opposing moral view with claims that they are "liberal" or "socialist."

FitzPatrick argues that it is clear that Potter could have engaged with the multitude of opportunities to remedy his normative ignorance, but does not. Unlike the Ancient Slaveholder, who by hypothesis did *not* have easy access to opposing moral views, Potter could have easily arrived at the correct moral view; after all, many of Potter's friends and neighbours did just that. He suggests that this kind of ignorance is what Aquinas would describe as "when [a person] does not actually consider what he can and should consider (this is ignorance of evil choice, arising either from passion or habit) or when [a person] does not take care to acquire the knowledge which he ought to have" (quoted in FitzPatrick 602). It is not that Potter knowingly and willingly avoids taking other views into careful consideration, but it is that Potter is engaging in some epistemic vice that debilitates his ability to do so. There are no barriers to Potter's social context that prevent him from gathering information and deep reflection. It is only because Potter voluntarily exercises epistemic vices — such as overconfidence, arrogance, or dismissiveness, to name a few — that he

fails to do so. It would then be reasonable to conclude that if Potter did not exercise these vices, he could be reasonably expected to take the steps to remedy his ignorance. Thus, FitzPatrick concludes with a much weaker criteria for culpable ignorance:

> "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 609).

FitzPatrick also takes time to defend his position against anticipated objections, but I believe it will be more worthwhile to address his defenses in direct conversation with Neil Levy's objections, which, of course, will require that I fully lay out Levy's position.

Levy believes that FitzPatrick is asking the right 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?" — but he argues that FitzPatrick arrives at the wrong conclusion. In actuality, if we are to *fairly* answer the R Condition in relation to Potter, for example, then we would have to concur that there was nothing Potter could have reasonably been expected to do to remedy his normative ignorance. To Levy, FitzPatrick's belief that Potter *could* have behaved better had he not voluntarily exercised certain epistemic vices is either tautological or false. It is tautological if FitzPatrick intends on saying that agents are culpable for their ignorance if it is a result of the culpable exercise of epistemic vices. It is false if FitzPatrick intends on saying — which Levy thinks is the case — that "if agents failed to take advantage of opportunities to avoid or remedy ignorance due to their epistemic vices, their failure was culpable because we can

reasonably expect agents not to exercise these vices" (Levy 734). Given Potter's case as laid out by FitzPatrick, Levy argues that he could not rationally have taken advantage of the available opportunities for moral improvement, so we cannot reasonably expect him to. Thus, he concludes that the only way we can hold agents culpable for their unwitting wrongful acts performed from ignorance is if we can expect them to reasonably take advantage of their opportunities. And such a claim entails that akrasia is, once again, a requirement for culpability.

To breakdown Levy's argument further, he argues that Potter cannot be reasonably expected to have improved his normative ignorance because it was not rational for him to do so. He succinctly states that "[w]hile the demands of morality are typically binding on all rational agents, regardless of their ability to comply with them, the demands that underwrite culpability are sensitive to agents' capacities" (Levy 735). Thus, if we are focusing on culpability, it is necessary that we focus on what an agent can *rationally* be expected to do when considering the R Condition. What is *reasonable* for an agent to do is dependent on external factors regardless of whether the agent is aware of them. It is not what is *rational* for an agent to do, which are an agent's *internal* reasons and proattitudes she genuinely believes in that a certain course of action will bring about a state of affairs that satisfies her goals.

The rationality Levy focuses on is an *internal* one. He emphasizes the reasons that explain what is rational for an agent to do, and what is rational for an agent to do is determined by her internal proattitudes. Contrarily, what an agent has the most reason to do is not always something that can be reasonably expected of an agent. To wit:

> "[A]gents have a fair opportunity to avoid performing actions (and so on) only if they could have rationally chosen to omit the action, and what agents can do rationally is a function of their internalist reasons alone. What we can do rationally

> is what we can do by means of a reasoning procedure, and our reasoning procedures can only operate on our actual representations. Only when a potential action is (at least minimally) rationalized by their beliefs can agents perform it rationally. If I do not represent the state of affairs at which an action is aimed as desirable, or as furthering the satisfaction of my goals or ends, then I cannot settle on that action by means of a reasoning procedure; if I perform that action, I will do so only as a consequence of a failure of my reasoning." (Levy 739).

To illustrate his contention, suppose while I am driving home one night, I come across a car that looks like it has been badly damaged from crashing into a tree off the side of the road. The front of the car is crumpled, I see smoke coming from the engine, and it is obvious to me that the airbags went off. I pull over to the side and immediately rush out of my car to see if there are any survivors. To my relief, the car's only occupant, the driver (obviously) is alive, but unconscious. I dial 9-1-1, give them my location, and wait for their arrival. As I wait, I am assessing my surroundings and suspect, on good grounds, that the engine is going to explode. There is a lot more smoke coming out, I see sparks. I conclude that the best thing to do at this moment is to pull the driver out of the car to save them from the guaranteed death that will come from the impending car explosion. However, by the time the EMS arrive on site, the car has not exploded, and, to my shock, they tell me that by moving the driver, I injured their spinal cord and paralyzed them for life.[7] My intended good deed turned out to be wrongful due to my ignorance. Levy would argue that it would be unreasonable to find my ignorance culpable because, while with the benefit of hindsight it is obvious that the course of action I had the most *reason* to do was to leave the driver in the car, it was not *rational* for me to have known better. From my perspective, I could not have

[7] This example was first introduced by Michael Zimmerman. See: Zimmerman, Michael J. "Moral Responsibility and Ignorance." *Ethics*, vol. 107, no. 3, Apr. 1997, pp. 410–426, www.jstor.org/stable/2382323.

known that the car would stay intact and that the safest thing to do for the driver was to keep my hands off of them.

Thus, it is false that merely having his moral views challenged is any reason for Potter to seriously reassess his moral stance. And even if he does take the time to reflect on his views and the criticisms that FitzPatrick points out, it is very difficult for him to conclude that his initial views were mistaken. And if this is true for Potter, then it's certainly true for the Ancient Slaveholder. The R Condition must be considered in an internalist manner — "the agent has to see herself as having reasons to take her opportunities" (Levy 738). It would be unreasonable to blame agents for actions they cannot remedy or avoid through some reasoning procedure; they can only be appropriately blamed for actions they had a fair opportunity to avoid, which requires rational capacities.

Levy's argument relies on two claims: 1) the Reasonable Expectations Claim and 2) the Capacity Claim. The Reasonable Expectations Claim states that "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" and the Capacity Claim states that "an agent has the capacity to form her beliefs in the expected way only if she can conform rationally to the expectation" (Robichaud 141). According to Levy, an agent should only be culpable for an unwitting wrongful act performed from ignorance if she could be reasonably expected to have followed some procedural epistemic obligation, which she did not do prior to acting and so acted from ignorance. And further, agents can only be reasonably expected to follow some procedural epistemic obligation only if they were able to. I believe that FitzPatrick would not have any issue with either of these claims. In the question that both FitzPatrick and Levy start with, there is an assumption that agents should only be culpable for things that they could reasonably have been expected to do or avoid. However,

notice that these two claims do not necessarily entail that culpable ignorance must require akrasia. He could have simply made the claim that it would be plainly rational for Potter to scrutinize his beliefs, since unscrutinized beliefs are more likely to be false (Robichaud 141). It is only in Levy's definition of "rationality" that we arrive at the akrasia requirement (Robichaud 141-142). Since he defines "rationality" in the Capacity Claim in an internalist way — i.e. the agent must have internal reasons to do or avoid doing a certain action in order for it to be "rational" (Levy 739) — then we return to the akrasia requirement where agents can only be culpable for their ignorance only if they could rationally (internally) think that they are obligated to do or avoid doing something, but do not.

Levy concludes his article with the striking note that "Rosen was right all along: only if we can rightfully attribute akrasia to agents can we hold them responsible at all because we cannot reasonably blame them for actions they could intentionally omit only through a failure of their practical rationality" (Levy 741).

Although the debate goes beyond the interaction I have outlined up to this point[8] (Robichaud 137-151, Levy, "Culpable Ignorance: A Reply to Robichaud", 263-271), I do not believe that introducing any of the further arguments adds to my purposes. So, now I want to now bring in the nuance that is in the anticipated objections FitzPatrick poses to himself and attempts to address and Levy's response to them.

FitzPatrick briefly acknowledges that perhaps the character traits that Potter has led to his epistemic vices were formed in early childhood, so it is unreasonable to say that he could have known better; "given his background and his actual beliefs, it really wasn't possible for him to

[8] For more on this discussion, see: Robichaud, Philip. "On Culpable Ignorance and Akrasia." *Ethics*, vol. 125, 2014, pp. 137–151, doi:10.1086/677139.; and Levy, Neil. "Culpable Ignorance: A Reply To Robichaud." *Journal of Philosophical Research*, vol. 41, 2016, pp. 263–271, 10.5840/jpr201672778.

behave other than he did, and it would therefore be unreasonable to expect this" (FitzPatrick 608). He dismisses this worry on the grounds that it falls back on deterministic arguments that undermine moral responsibility, and is, hence, irrelevant to his discussion. However, Levy presents a closely related argument: Given Potter's background, it *might* be the case that he could not have acted other than he did, but it is true that, given his background, he could not have *rationally* done other than he did. Levy proposes that:

> "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." (Levy 737).

I am inclined to agree with Levy. The idea that Potter could not have been expected to rationally do better seems to be a critical criticism of FitzPatrick's position. Returning to the main tenets of moral responsibility, it is inappropriate to blame someone for doing something wrongful if they did not know what they were doing and they were unable to control it. If Potter — and it seems to be the case — genuinely did everything he could which he deemed to be right and could not seriously control his actions, then it seems intuitively and logically unreasonable to blame him for his wrongdoing. However, I am hesitant to lend full acceptance of Levy's criticism. Levy intends for his criticism to entail that Rosen was right all along. Agents, it turns out, cannot be blamed for

a wrongful act committed out of ignorance unless it is a case of or can be traced back to a case of akrasia. I do not endorse such a demanding position for moral responsibility.

As an aside, I do see an opportunity to agree with Levy's overall Internalist position *without* endorsing Rosen's view by separating Internalism from Tracing. While both arguments conclude that akrasia is a necessary requirement for culpability, they arrive at their conclusion in different ways. Tracing views rely heavily on an agent's causal history — looking back on actions and determining when or if an act was ever performed out of akrasia. Such backwards-looking has intuitive reasonability and is enlightening in many aspects, but it seems to be neglecting an important aspect of an agent's decision making process. Namely, reasonability. Here is where Internalist views fill in the gaps. Internalism, at least the version Levy puts forth, begins from the R Condition that focuses on what agents can reasonably do. Internalism takes into account what agents consider to be rational, and does not rely on tracing their causal history. Further, since Internalism does not look backwards into an agent's causal history, it avoids the regressive structure of Tracing Views. While Tracing Views can look back on an agent's causal history to find an instance of akrasia ad infinitum, Internalists do not. Rather, they only have to show whether it was reasonable for an agent to have known that what she did was wrong. That does, of course, involve attention to the past, but it can include counterfactual questions about what the agent could have reasonably done, rather than specific questions about the exact original cause of the agent's ignorance. Internalists would not say that culpability *traces* back to akrasia, only that it seems to require it. Although Levy agrees with Rosen in the end, I think they are ultimately proposing differing views with different starting points and answering different questions.

2.4. CONCLUSION

Returning to the main issue at hand, I am uncomfortable with Levy's conclusion that brings the akrasia requirement back. To allay these worries, I intend to defend FitzPatrick's view in the following chapter.

First, I propose that Levy's argument is resting on two false assumptions, which make his argument for moral responsibility untenable. The first is that he mistakenly conflates the criteria for culpable factual ignorance for culpable normative ignorance. The second is that he mistakenly assumes that upbringing and cultural influences can "blind" an agent from being able to rationally (internally) conform her beliefs to what is "right."

By pointing out these concerns that I have with Levy's argument, I intend to show two things. The first is that Levy's criticism of FitzPatrick's Epistemic Vice View and his subsequent theory of Internalism cannot stand under scrutiny when applied to cases of normative ignorance. The second is that FitzPatrick's position holds more plausibility than Levy's. All of this is working towards my more substantive conclusion that some actions from normative ignorance are never excusing. I will proceed to explain this in further detail in the next chapter.

CHAPTER 3

3.1. INTRODUCTION

In this chapter I will introduce my most substantive claim, which is that some actions from (specifically) *normative ignorance* are rarely, if ever, excusing. And these actions from normative ignorance are more commonplace than perhaps anticipated. While they include villains and evildoers, they also include everyday people who are mistaken about the morality of their actions. In essence, more unwitting wrongful agents are on the hook for their actions performed from ignorance than supposed by Levy.

To support this conclusion, I will argue that Levy's argument relies on two falsely made assumptions. The first is that what an agent can fairly be blamed for in cases of factual ignorance can translate directly to what an agent can fairly be blamed for in cases of moral ignorance. The second is that an agent's upbringing and cultural influences can effectively "blind" them from the possible immorality of their actions. I will use the work of Michele Moody-Adams and Charles Mills on culpable ignorance to support this second idea.

3.2. FACTUAL IGNORANCE AND NORMATIVE IGNORANCE

Firstly, from the outset of the argument, beginning with Rosen, there is, I argue, a falsely made assumption. Namely, that what an agent can be fairly blamed for in a case of *merely factual ignorance* can be translated to cases of *normative ignorance*. Rosen structures his argument to contemplate actions from ignorance of *mere fact* first. Here is where he presents his main conclusion that actions from ignorance are only blameworthy if they are the causal upshot of an instance of akrasia. However, he does not go on to make a distinction between actions from ignorance of *mere fact* and actions from *normative ignorance*. He clearly states that:

> "Nonetheless, I maintain that our principles [for action from ignorance of mere fact] hold in full generality, no matter what the content of the agent's ignorance may be. It might be ignorance of law; it might be ignorance of moral principle. It might be ignorance of the rational force of moral principles or ignorance about what is to be done all things considered. In each case, whenever the agent acts from pertinent ignorance, he is culpable for the act only if he is culpable for the ignorance from which he acts, and he is culpable for the ignorance only if it is the upshot of prior culpable negligence or recklessness in the management of his opinion." (Rosen, "Skepticism About Moral Responsibility", 304).

From here he proceeds to argue "that the principles [he has] articulated hold good even when the normative error in question is not the upshot of factual error" (Rosen, "Skepticism About Moral Responsibility", 304). Rosen, himself, notes that this is a substantive claim. And this is the claim that I find to have fallen victim to his false assumption that normative ignorance may be treated the same as ignorance of mere fact. He contends that in common, everyday acts, agents may have full knowledge of all the relevant factual details before acting, but be mistaken about the moral weight that each one of these considerations should take. To put it in a concrete way, Rosen presents the example of Bill, who knows that it is wrong to lie to his wife about where he was at night, but tells a self-serving lie because if he told the truth, then he would suffer (Rosen, "Skepticism About Moral Responsibility", 305). The strongest moral reason (let's assume) points towards telling the truth, but Bill mistakenly does not give the moral considerations of this strong reason as much weight as he should (for reasons that are irrelevant to my current exposition), so he concludes that telling self-serving white lies and flouting the requirements of morality is acceptable in certain cases, such as his present one. So he tells a self-serving lie.

Rosen calls examples that mimic this — instances where the agent 1) does something wrong without knowing that it was wrong, but 2) knew all the relevant *merely factual* details, and 3) was adequately reflective prior to acting — bare ignorance of the rational force of moral considerations (Rosen, "Skepticism About Moral Responsibility", 305). And in these cases, we are in no position to hold the unwitting wrongful agent culpable for their ignorance. Thus, Bill's wife would only be able to blame Bill for holding a view that telling self-serving white lies is morally acceptable from time to time, and she can only blame Bill if this was a decision that was within Bill's control. But since Bill had considered all the relevant details of his situation carefully, but only made a calculation error, he should not be to blame for his belief or his lie (it was a mistake!). So, Bill's wife should not blame Bill, under Rosen's argument.

However, this argument relies on Rosen's claim that actions done from normative ignorance may be treated the same as actions from ignorance of mere fact. FitzPatrick catches the assumption, but does not take issue with it. Witness:

> "There is no doubt something compelling in the intuition at the core of [Rosen's] argument, that it is unfair to blame someone for an action done out of ignorance that he cannot fairly be blamed for having. This is something we already accept when it comes to circumstantial ignorance, and Rosen is on solid ground in extending the point to include normative ignorance as well: if circumstantial ignorance can excuse when it is not one's fault, then it seems that the same should be said for normative ignorance…." (FitzPatrick 601-602).

Rather, FitzPatrick believes that there is inherent merit in Rosen's assumption. And although I agree with FitzPatrick to the extent that I can see where this assumption is coming from — that, for continuity's sake, we should treat all cases of ignorance the same, and we would have some

logical reasons for doing so — I do not agree that this is something to be merely assumed. Instead, this assumption should be scrutinized. But when put under scrutiny, as I will proceed to show, the assumption cannot. I would like to take a moment now to make a note that this false assumption does not make its way into FitzPatrick's argument, but it does permeate into Levy's. How it occurs will be explained further in this section. For now, I will only aim to put Rosen's assumption under scrutiny. To do so, I intend to show that the nature of ignorance of mere facts and normative ignorance are different, and should be treated with this consideration in mind.

As support for this contention, I want to return to a point I made in the second chapter that there is a spectrum of ease of knowability between matters of mere fact and normative principles. Matters of mere fact can be both trivially easy to know and impossible to know. It is trivially easy to know, for instance, that you should not feed your dog chocolate or give water to your two week old infant through a simple Google search. There are very few barriers to coming to know these facts, if you do not already. Some things are very easily knowable. On the other end of the spectrum, it would perhaps be impossible to know for certain what the inside of a black hole looks like or what sound a triceratops made. Astronomy and astrophysics is unable to successfully bring humans in and out of a black hole intact (and perhaps it never will) and we do not have the ability to time travel (and, again, perhaps we never will) to the Cretaceous period to hear a triceratops. Some things are almost impossible to know because we do not have access. Merely factual ignorance, then, can be separated between things that are easily knowable and things that are impossible to know. It occupies clear spots on the spectrum of knowability.

Further to that, matters of mere fact can usually entail understanding and acceptance. Supposing that I have no ill will towards my dog and then come to know that dogs should not eat chocolate, then I would understand and accept that I should not share my chocolate bar with her.

If I am concerned about my infant's fluid intake and then come to know that infants should not drink water, then I would understand and accept that I should not give her water. And to flip the examples, perhaps I am a deeply malicious person, so when I come to know these facts I act according to my intention and poison my dog and my infant. My understanding of a matter of mere fact can (usually) entail that I will act accordingly.

In the normative case, it is more complicated. To illustrate the complication, I will refer back to the Ancient Roman Slaveholder who does not know that owning slaves and making them fight lions for entertainment is wrong. I can accept that the Ancient Slaveholder has no inkling that what he is doing is wrongful, but that is not the same as saying that the argument against slaveholding was unavailable to him.

Now, the unavailability of the argument against slaveholding to the Ancient Slaveholder is a complicated contention. I could (assuming that I had both a time machine and spoke Latin!) explain to the Ancient Slaveholder why slavery is morally wrong. And he may *understand* it, but that does not entail that he will necessarily *accept* it. Suppose this is pre-Christian Rome. The egalitarianism of Christianity would not make any sense to the Ancient Slaveholder, nor would the concept of sin or of one's moral duty deriving from God's law. So that is one argument that would not hold any persuasion. Further, any argument that I would pose *now* would include a central appeal to autonomy, and the idea that everyone's autonomy was equally valuable would likely also be entirely foreign and unpersuasive to the Ancient Slaveholder. I could go on trying to find arguments that I would use to convince the Ancient Slaveholder that he is doing the wrong thing and persuade him to reconsider his actions, and I may find approximations of certain concepts (maybe even the ones I just mentioned!) that might be translatable. But there would be significant differences as well. These translations are only approximate, and they will not ensure the Ancient

Slaveholder's acceptance. If I were given a time machine and a couple years to study Latin, I would accept these statements to be true before embarking on any mission to change the Ancient Slaveholder's mind.

However, what I am trying to get at here is that the Ancient Slaveholder's problem is not about *access* to evidence. I accept what Rosen, FitzPatrick, and Levy all accept, which is that non-culpable ignorance involves non-culpable lack of access to evidence. Culpable ignorance requires that the ignorance came about *culpably* (whatever that may mean for each theorist). But that is not obviously present in the case of the Ancient Slaveholder. He can have all the evidence to suppose that his slaveholding activities are wrongful. What he is missing is the conceptual framework for evaluating it. The conceptual framework I am thinking of is one where the Ancient Slaveholder could reason differently, to self-scrutinize his ideas. The distinction between lack of access and lack of conceptual framework is important for me here because the Ancient Slaveholder *could* think differently. And although the mere opportunity to think differently does not entail that an agent will think differently, I want to emphasize that there are cases where an agent's ignorance is culpable due to a lack of sufficient self-scrutiny. And that this is mostly the case in instances of normative ignorance.

Having illustrated the difference between matters of mere fact and normative principles, let us return to why Rosen is mistaken for conflating the two. The central question in the culpable ignorance debate is when is ignorance culpable? All three of the theorists I have focused on agree that there are instances where ignorance is culpable. Further, all three of these theorists would agree that agents need to satisfy their procedural epistemic obligations. Failure to satisfy these procedural epistemic obligations either through recklessness or negligence would be inculpating. However, where these theorists separate in their thinking is over the question of how strong our

moral epistemic obligations are. In other words, what moral obligation do we have to properly manage our moral beliefs?

Rosen does not believe that there should be any difference between our moral epistemic obligations and our factual epistemic obligations. But he fails to recognize that there is an important difference between factual beliefs and moral beliefs. Namely, that, for the most part, there is little risk that we are going to harm or wrong someone by having a mistaken merely factual belief or being ignorant of some mere fact. For instance, the risk is *very* low that my sugar bowl actually contains poison, and so there is little reason for me to double check it because there is little risk of wrongdoing if I do not. But: say I have two identical sugar bowls in my cupboard, one with sugar and one with arsenic, then the risk of wrongdoing is higher, and because it is higher and I ought to know this, my obligation to confirm I grabbed the right bowl is also higher. One difference between factual beliefs and moral beliefs is that, in the moral case, the risk is *always* high that I will wrong someone if I have a mistaken moral belief or if I am morally ignorant. Because of this, the beliefs in question when it comes to cases *about* morality almost always have implications for how I treat people. So I always have a relatively higher obligation to investigate and examine my moral beliefs because the risk of being wrong is high. And I would only have a similar obligation when it comes to matters of mere fact when the moral risk of being wrong is also high. Thus, the procedural epistemic obligations must be different for cases of normative ignorance.

As I stated earlier, I believe that Rosen's false assumption is also taken up by Levy. However, the point in which he introduces this false assumption is more concealed than Rosen. As a reminder, Levy's argument relies on his definition of "rational," which is understood in an internal sense. What is rational for an agent to do — which is by way of reasoning procedure —

is determined by her actual representations and proattitudes (Levy 735-736). Rationality is what an agent can decide to do as a result of engaging in reasoning, which is settled by her *internal* mental state (Levy 735). And, as his argument goes, we are only justified in blaming agents for unwitting wrongful actions if we can reasonably expect them to have rationalized (internally) differently.

Based on this definition of rationality, Levy further proposes that a moral outlook is an internal perspective on how we view the world — "It constitutes the viewpoint from which we assess moral claims and in terms of which we assess whether a particular objection to our views should be regarded as an annoyance or a serious challenge" (Levy 736-737). Accordingly, he argues that, in the case of the Potter example, Potter gives the moral arguments that challenge his position as much scrutiny as he deems necessary because his outlook (that any appeal for him to refrain from socializing risk for his personal financial gain is "socialist" and "liberal" and wrong) is constituted by his (internal) moral outlook. In this sense, Potter manages his normative views adequately, and we cannot blame him for doing otherwise because it would not be rational for him to do so.[9]

Levy further strengthens his claim by claiming two things. The first being that even if Potter's views were challenged, that does not entail that he would find any internal reason to scrutinize them. He claims that "[o]ur moral views are challenged all the time, without us feeling challenged. We dismiss the views of Holocaust deniers and racists without pausing to consider whether there is a grain of truth in their claims (and we are right to do so)" because they are

[9] For the purposes of my argument, I will assume that Potter is wrong in his judgement that socializing risk for his personal financial gain is acceptable. I understand that there may be arguments to support his position, but that is irrelevant to my current discussion. And if the reader finds my assumption to be incorrect, I ask that you flip the example. My argument will stand just as well in the opposite sense.

egregiously mistaken (Levy 737). And the second claim is that even if Potter were to have scrutinized his beliefs, it is very difficult for him to come to the conclusion that he was mistaken.

I contend that Levy's false assumption that what an agent can be blamed for in a case of ignorance of mere fact can be translated to what an agent can be blamed for in a case of normative ignorance is embedded in his view of a moral outlook and his subsequent strengthening claims. I accept Levy's claim that a moral outlook — which determines how we judge our morals and any challenges they face — is an internal matter. All things considered, it is possible that Potter is presented with and understands all the evidence for why privatizing profit and socializing risk is morally wrong, while still not accepting it. Just as with the Ancient Slaveholder example in this section, the issue at play here (since it pertains to arguments for moral claims) is not whether Potter has access to arguments that challenge his view. Potter is not prevented from knowing the views that challenge his own. Rather, the issue he faces is with his conceptual framework. He cannot rationally (internally) accept that the challenges to his view should be given any significant uptake; and I accept this as true.

However, I contend that Levy is wrong about how we determine when objections to our views are an "annoyance or a serious challenge." We do not dismiss the arguments of Holocaust deniers and racists without hesitation because they do not pose real challenges to our moral views, as Levy proposes. We dismiss them instantly because we have done the work (or accept the work that has been done for us) to determine that they do not stand up under scrutiny. We determine this for reasons beyond the challenge that they pose to our beliefs. Self-scrutiny of our moral views needs to match the risk of our moral views. We typically understand that the Holocaust and racism pertain to moral claims, and these have serious effects on real people. And since the risk of wrongdoing in cases of moral claims is always higher than in cases of mere facts, then we must be

expected to self-scrutinize our moral beliefs to a greater degree. So it is not that we immediately decide that the arguments of Holocaust deniers and racists are wrong. It is that we can see the fallacies in their argument that make them implausible.

Thus, Levy is incorrect in asserting that an agent's culpability relies on her ability to rationally (internally) conform her beliefs to an expectation. An agent has expectations to self-scrutinize her beliefs to a higher degree in cases of normative claims.

But what Levy supposes is rational (internal) can apply to cases of mere fact. Perhaps I do not accept the challenge to my belief that I forgot to say "happy birthday" to my friend because it is August and I know that she is a Pisces. Or that I do not accept the challenge that my brake pads need replacing because I just got them replaced a month ago. Or that I do not accept the challenge that I sent an email containing confidential information to the wrong person because I later received a call from the intended receiver telling me that they got my email. Granted, I may still be incorrect in my judgement (I could be wrong that my friend is a Pisces and that every brake pad installation is done successfully and that the intended email recipient was referring to the email containing confidential information). In cases of mere fact, I can accept Levy's conclusion. Since I did not consider these challenges to my views to be *serious challenges* to my view because of my outlook on the world, I could not have been reasonably expected to conform my beliefs. Thus, I should not be held culpable.

And I can accept this position easier when framed in cases of ignorance of mere fact because of the level of risk. It may be true that my friend is truly devastated that I missed her birthday and that my shoddy brake pads result in an accident that injures people and that the leak of confidential information causes an intense legal dispute involving my company. But these cases of mere fact are less *certainly* harmful in comparison to cases of normative beliefs. It may just as

well be true that my friend understands that sometimes people forget birthdays and that my brake pads hold up until my next vehicle service and that the leak of confidential information was deleted by the unintended recipient. However, the same cannot be said about normative ignorance. Potter's normative ignorance has a greater risk to people's financial health and wellbeing, and, in a way, entails it. Just as the normative ignorance of a Holocaust denier or racist has a greater risk to people who were (and still are) affected by the Holocaust and racism.

Potter is required to self-scrutinize his beliefs to a higher degree because of the risks he takes on by not. It is a procedural epistemic obligation that he does so because his actions pertain to morality. Levy's requirement of him is too weak for matters of normative ignorance. And while there are surely genuinely hard cases of moral issues where even after careful deliberation and self-scrutiny an agent can still act wrongly due to some unforeseen ignorance, these are not the cases that we often come across, and they are not the ones of my present concern.

It can surely be agreed upon that Potter is an evil character. He is cast as the villain in *It's a Wonderful Life*, and we can find it easier to blame him for his actions due to this characterization. However, he, and characters like him, are not the only ones subject to this blame. In criticizing and amending Levy's argument, I open avenues for agents who are generally good people, but who make wrongful normative judgements. For instance, an agent who tells a sexist joke in the workplace that she thought was funny, without knowing the effects of her saying it, could also be found culpable. And she would be found culpable because 1) (assuming she is working in a contemporary temporal setting and she has no limitations in her environment that would prevent her from accessing information about the dangers of sexist rhetoric) she does not have any issue with access to information, and 2) she should have self-scrutinized her actions to a higher degree.

In anticipation of an objection, I do not doubt that the agent may have self-scrutinized her normative beliefs to a sufficiently high degree, but still came to the wrong answer. She may believe that telling a funny joke to impress her coworkers is in her best interest. But I do not accept that her self-scrutiny is enough to exculpate her. As I mentioned, agents who are generally good, going about their day-to-day lives, may also be found culpable because of their insufficient self-scrutiny. The agent in question cannot claim that she has not heard arguments against telling sexist jokes. Further, she should have an understanding that her act of telling a sexist joke falls in the category of morality. Thus, she should be expected to self-scrutinize her joke beyond thinking it is merely funny.

With all this said, I anticipate that Levy may respond to all of this by arguing that perhaps all this may be true, but if an agent — such as Potter — was raised to have such a worldview that privatizing profit was the right thing to do in all or most cases, then he should not be blameworthy for his actions. He should not be blameworthy because he would not rationally (internally) know better. And if he were to do anything other than privatize profit — perhaps because he felt remorse or guilt — then he would have done it by luck or chance. And actions that can only be done out of luck or chance cannot be blameworthy because the agent would not have been able to conform rationally to that expectation.

3.3. CULTURAL "BLINDNESS"

In response to that criticism, I want to introduce the work of Michele Moody-Adams and Charles Mills on culpable ignorance to argue that Levy is making a second false assumption that upbringing and culture can "blind" agents from questioning the morality of their culture's practices. The idea here is that it is in fact possible for the ancient slaveholder or for Potter to know that their beliefs are wrong and so that their actions are mistaken, and so to the extent that they *do*

not know they have failed to fulfill their epistemic obligations. To illustrate this point, I aim to show that much of the normative ignorance that we are typically concerned with are cases of *affected ignorance*.

First I will give my attention to Moody-Adams. Moody-Adams proposes that there is a commonsense view that cultural influences can, and often do, cause serious impediments to responsible agency, and thus our attitudes toward blame should acknowledge this. She calls such an impediment the inability thesis. The inability thesis about cultural limitations claims that "sometimes one's upbringing in a culture simply renders one unable to know that certain actions are wrong" (Moody-Adams 293). It claims that "it is a commonplace of everyday moral reflection that it is 'only fair' to judge people by the standards of their day…[so] we ought to judge participants in past cultures as we would want to be judged by future generations" (Moody-Adams 298). The thesis is bound up in the claim about diminished responsibility. While there are many variations of this view — of which I think Levy's is one of, but I will return to this — they all "accept the standard bifurcation of excuses into coercion or nonculpable ignorance, and they then argue that cultural influences can excuse wrongdoing because they produce nonculpable ignorance" (Moody-Adams 293). These arguments propose that culture can sometimes "blind" agents to alternative courses of thinking and acting, hence rendering agents unable to question the morality of the cultural practices they live with.

However, Moody-Adams wants to reject the inability thesis on the grounds that, while appearing to be an empirical claim, it has no empirical support, and that it dangerously misconceives the ease at which wrongdoing occurs. To support her first claim, she argues that a culture cannot persist without the internal conformity of individuals. Cultures are "a way of life of a given social group that will be shaped by more or less intricate patterns of normative expectations

about emotion, thought, and action" (Moody-Adams 295), but its survival depends on the emergence of support from the participants of the culture. It is only from an internal perspective that those subject to the demands for conformity and criticisms of society's rules that the culture can be seen as justified. In these internal perspectives, a culture's participant may even come to self-criticize their conformity of or lapses in following the culture's rules. Any appeal to the commonsense view that we ought to judge people by the standards of their time or cultural limitations vehemently ignores that our own failings as human beings may be the ones that perpetuate cultures and that there is a common — and culpable — tendency "simply to affect ignorance of the possibility that some cultural practice might be morally flawed" (Moody-Adams 298).

To support her second claim, she introduces the idea of the banality of wrongdoing,[10] which suggests that there is a tendency for "self-deceptive complacency about the potential each agent has to support and engage in morally culpable conduct" (Moody-Adams 299). Here, she introduces affected ignorance, which is "choosing not to know what one can and should know" (Moody-Adams 296). She proposes there are many forms affected ignorance can take, but the most commonplace, and the most dangerous, is the tendency to "avoid acknowledging our human fallibility" (Moody-Adams 301). She cites an example of a university administrator who refuses to investigate a colleague's charge of sexual assault because "he couldn't possibly be guilty of such a thing." Such a tendency is pervasive and proves that there is evidence for the banality of wrongdoing. Thus, everyday moral reflection should not be so oblivious to such evidence.

[10] The banality of wrongdoing is inspired by the banality of evil, an idea that ordinary people are capable of extraordinary evil, introduced by Hannah Arendt. A full understanding of Arendt's idea of the banality of evil is not relevant for my purposes and is beyond my scope.

Her position shows that "no culture is perpetuated without some modification of cultural patterns in the lives of individuals," and claims that culture can pose serious impediments to people's judgements dangerously distorts the connection between culture and agency, and prevents any careful understanding of what a culture is (Moody-Adams 306).

I contend that Levy's argument that if people are not capable rationally to conform their ideas to an expectation, then they should not be held responsible for their non-conformity and any wrongful acts that result from it appeals to the inability thesis. The validity of his argument relies on the assumption that cultural influences can effectively "blind" people to alternatives to their cultural practices, and such "blinding" produces nonculpable ignorance in relation to the morality of their actions. However, as Moody-Adams contends, the very act of perpetuating ignorance that a cultural act may be morally flawed is a culpable act. Culture relies on individual participants to justify its functioning. If people were to question and criticize the morality of some cultural influence, that cultural influence would lose its power and conform to the culture's participants.

Further, there are instances of everyday morally questionable behaviours that are accepted culturally. Potter is a good example. In the United States, where *It's a Wonderful Life* takes place, ruthless capitalism is upheld by the American Dream. Privatizing profit and building a fortune for yourself is considered a morally acceptable thing to do. It may even be praiseworthy and enviable. And since Potter was raised in such a culture — perhaps, as FitzPatrick suggests, he was born into a family that praises these values and he surrounds himself in his adult life with people who also share these values — he is not considered morally reprehensible or wrong for his actions. However, this is an example of the banality of wrongdoing, which is a form of affected ignorance, and affected ignorance is culpable. Levy's appeal to the inability thesis makes his assumption that we cannot hold people morally responsible for wrongful acts performed from ignorance unless it can

be traced back to an ability to rationally (internally) conform false. Thus it is not out of luck that Potter *could* have done differently. Culture and upbringing does not provide an excuse for wrongful acts performed out of ignorance.

Rather, Potter's ignorance is affected. The argument that Potter could not have known better because he lives in a culture that supports his ambitious capitalist actions does not entail that he is unable to know better. As Moody-Adams suggests, the most dangerous kind of affected ignorance is where agents blindly follow cultural norms without confronting the possibility of human fallibility. Potter may have been assured that his actions are acceptable, but his unquestioning acceptance of these views, which perpetuates the very culture that supports his wrongdoing and reifies it, is culpable. He is not unaware of the effects of his actions. He is only assured by the cultural judgements that his actions are acceptable. The banality of wrongdoing is pervasive, and Potter would be complicit in it. If he accepts the argument that his ruthless capitalism is acceptable based on the approval from his culture, then he is refraining from properly self-scrutinizing his beliefs, which is an instance of affected ignorance. Thus, Potter *could* have done better. Levy is wrong in suggesting otherwise. And since Potter *could* have done better, then it is reasonable to expect that he *should* have done better.

Now I want to turn my attention to Mills' Racial Contract. Mills presented the groundbreaking idea of the Racial Contract that sparked interest in the study of the epistemology of ignorance, which this current conversation is a subset of. He describes the Racial Contract as a social, political, and epistemological system that places white people hierarchically above nonwhite people, allowing nonwhites to be subjugated and allowing whites to perform the subjugating acts without moral guilt because the Contract is hidden in the eyes of the Whites (Mills 11-39). The Racial Contract is a central to the polity that invades most of the world, and hence,

racism finds its way into almost all social spaces. The Contract also plays the role of hiding itself, which produces "an inverted epistemology, an epistemology of ignorance, a particular pattern of localized and global cognitive dysfunctions (which are psychologically and socially functional), producing the ironic outcome that whites will in general be unable to understand the world they themselves have made" (Mills 18). Mills states that "[w]hites will then act in racist ways while thinking of themselves as acting morally. In other words, they will experience genuine cognitive difficulties in recognizing certain behavior patterns as racist, so that quite apart from questions of motivation and bad faith they will be morally handicapped simply from the conceptual point of view in seeing and doing the right thing" (Mills 93).

Mills makes various strong claims to support his idea of the Racial Contract in his book, but I will only focus on one for my purposes. Namely, I will focus on his idea of white moral cognitive dysfunction. Mills begins by asking "[h]ow were people able consistently to do the wrong thing while thinking that they were doing the right thing?" (Mills 94). His answer is that the Racial Contract establishes white moral cognitive dysfunction that both 1) forms global white solidarity that continues shaping white cognition and 2) hides the Racial Contract from white cognition. And building on this, he suggests that silence regarding the Racial Contract is complicit and "evasion and self-deception thus become the epistemic norm" (Mills 97). White moral cognitive dysfunction, which effectively blinds White people from seeing and understanding their own wrongful acts of racism.

There are three causes to this. First, "[b]ecause of the intellectual atmosphere produced by the Racial Contract, whites will (in phase one) take for granted the appropriateness of concepts legitimizing the racial order, privileging them as the master race and relegating nonwhites to subpersonhood, and later (in phase two) the appropriateness of concepts that derace the polity,

denying its actual racial structuring" (Mills 95). Second, "[b]ecause of the reciprocally dependent definitions of superior whiteness and inferior nonwhiteness, whites may consciously or unconsciously assess how they're doing by a scale that depends in part on how nonwhites are doing, since the essence of whiteness is entitlement to differential privilege vis-a-vis nonwhites as a whole" (Mills 95). And third, "[b]ecause the Racial Contract requires the exploitation of nonwhites, it requires in whites the cultivation of patterns of affect and empathy that are only weakly, if at all, influenced by nonwhite suffering" (Mills 95).

In essence, there is an epistemology that prioritizes whiteness in hierarchical form, but then denies the prioritization, causing ignorance amongst the general white population that they are privileged and expressing (or merely benefitting) from racist ideals. White people will act with the assumption that they are actually doing well morally because they are positioned differently than nonwhites. And the white moral cognitive dysfunction perpetuates Racial Contract, replicating it among generations and populations, and influences the abilities of white people to recognize the suffering of nonwhite people. But just because the racism exhibitted by some white people is hidden does not mean that it is exculpating.

He posits that missing knowledge is not the issue. Rather, it is the "courage to understand what we know and draw conclusions" (Mills 105). While nonwhites have always recognized the realities of race in morality and politics, it has been absent in the epistemology of whites (Mills 109). Until the Racial Contract is seen and understood, there can be no "serious theoretical appreciation of [its] significance" (Mills 125). Thus, it is not enough to say that when agents are unaware of their wrongdoings performed from ignorance, they are excused. There is a call for action to dispel the Racial Contract. And to sign onto the Racial Contract makes an agent complicit in the harms that it causes and justifies.

Although Mills focuses on racism, I believe he is pulling on a notion that is transferable to many other aspects of culture with the idea of moral cognitive dysfunction. Using the example of Potter again, Potter may be unaware that he is behaving immorally because he is in a culture that supports the efforts of ambitious capitalists. Even more, Potter may have a plausible set of moral arguments that could justify his conclusions. He has even rehearsed them! Maybe he believes strongly (with some good evidence) in the overall benefits of the free market in maximizing overall welfare, and thinks that this justifies fairly self-interested behaviour within a well-governed system. It is not *his* job, but the government's, to take care of other people — that is why he pays taxes. Potter of course would agree that it is morally wrong to unjustifiably cause someone harm. But he would *disagree* about when in fact we are responsible for the harms that others experience.

However, it is his moral cognitive dysfunction that produces this idea. Much like *white* moral cognitive dysfunction, the moral cognitive dysfunction present here privileges capitalist views. Those that are ambitious, even ruthless, in their capitalist pursuits are presented with privileges that those who are not ambitious or who have failed are not allowed. Further, this cognition that privileges ambitious capitalism creates, supports, and perpetuates the idea that ambitious, ruthless, capitalism is morally acceptable. And at the same time as it is being perpetuated, it is also being hidden from the eyes of the privileged, which in this case includes Potter. Although he is unaware that he is signing onto it, he is not off the hook yet. The privilege that allows ambitious and ruthless capitalists to benefit from people who are less financially well off will *always* harm others. While it is not Potter's job to overturn the system, he has a procedural epistemic obligation to recognize the system that he is complicit in is a harmful one. And in order to fulfill this procedural epistemic obligation, he must understand his actions and consider his own normative beliefs that support his complicit actions.

By not considering the effects of his actions and self-scrutinizing his beliefs, I argue that he is displaying affected ignorance. Mills, as a pioneer in this field, did not talk about affected ignorance in his book. But I contend that he would agree that the "blindness" of white people — and in my case, others who are privileged and experience moral cognitive dysfunction — is culpable because they exhibit affected ignorance. Potter's ignorance is affected because he is not unaware of what he is doing. As an ambitious capitalist, he is probably very strategic with what he does. He also does not lack access to the knowledge that what he is doing is harmful to others. Living in the same town as the very people he exploits, he has opportunities to see the effects of his business decisions. However, he acts with the confidence that his normative beliefs are correct and do not need any further scrutinizing. Once again, it is not Potter's lack of access that is an issue. What he lacks is the courage to understand that what he is doing is harmful. By failing to make a connection between his acts and the harms that others face, he is culpable. And his culpability, according to Mills' general idea of moral cognitive dysfunction, is grounded in his failure to self-scrutinize.

Having broken down the Potter example with the Moody-Adams and Mills' ideas, I further contend that it is not a big step to conclude that, more often than not, when we talk about unwitting wrongful agents who genuinely mean well but made a mistake in their normative beliefs, are also culpable. Both Moody-Adams' idea of the banality of wrongdoing and Mills' widespread moral cognitive dysfunction applies to a wide range of people, who may otherwise be moral people. I contend that, with this in mind, most of the normative ignorance that is of concern in this literature is culpable. Culture, and the subsequent argument using moral luck, is a poor excuse for culpability. Just because a culture engages in and supports an act does not mean that agents are "blinded" from knowing better. These very agents and their unwitting wrongful acts are what

perpetuate the culture. And agents who refuse to consider and self-scrutinize their normative ideas that were adopted from their upbringing and enhanced by their cultural experience display affected ignorance. Instances of these indeed do include banal, every day occurrences of normatively ignorant wrongdoing.

3.4. WHAT THIS MEANS FOR EPISTEMIC VICE VIEWS

Now that I have shown that Levy's argument relies on false assumptions, I will briefly discuss what this means for FitzPatrick's argument. In any way, just because Levy's argument is flawed in certain ways, does not mean that FitzPatrick's argument is correct. However, FitzPatrick's argument is not subject to the same criticism as Levy's.

First, FitzPatrick does not conflate normative ignorance with factual ignorance. Although he does not make any explicit mention between the two, he intends to focus on normative ignorance. And his criteria for normative ignorance is strong enough to withstand the objection I posed against Levy. Since he focuses on epistemically vicious practices, and views them as culpable, he is considering the possibility of egregious harms of unscrutinized normative beliefs.

Secondly, FitzPatrick does not use cultural blindness as support in his argument. His argument is perfectly compatible with cases where agents were raised and live in a culture that supports their immoral actions. The epistemic vices that are central to his argument determine culpability based on the culpable exercise of them. In cases where agents who do not recognize their wrongdoing because of cultural norms, FitzPatrick could still argue that they are culpable because they exercise some epistemic vice.

Although FitzPatrick's argument is not subject to the same criticisms as Levy's argument, I do believe that his argument could be strengthened. I propose that further work could be done on relating the exercise of epistemic vices to intellectual irresponsibility, as Montmarquet does

(Montmarquet 43). But even then, what actually makes intellectual irresponsibility culpable in cases of wrongful acts performed from ignorance needs to be done.

3.5. CONCLUSION

In this section I proposed my most substantive claim, which is that most unwitting wrongful agents who act from normative ignorance are culpable. I introduce this claim by objecting to Levy's position that, once again, requires akrasia. I contend that Levy makes two false assumptions about what agents can rationally do to dispel their ignorance. The first false assumption is that factual ignorance may be treated the same as normative ignorance. In reality, normative claims ought to require a higher level of self-scrutiny because of the risk level associated with these beliefs. There is always a higher risk of harming others when it comes to normative claims, and agents ought to be careful about how and why they hold certain normative claims. Just because they were unable to rationally (internally) self-scrutinize, does not mean that it is not expected of them. The second false assumption is that culture and upbringing may effectively "blind" agents from proper self-scrutiny of their normative beliefs. Using the work of Moody-Adams and Mills, I contend that it is not the case that culture is a viable excuse against culpable ignorance. Agents *can* self-scrutinize their normative beliefs, but if they accept that culture is a reason to refrain, then they are exhibiting affected ignorance, which is culpable.

My ambition in this project was to refute the skepticism about moral responsibility present in Levy's argument that was used to deny FitzPatrick's Epistemic Vice View. While I see an avenue to strengthen FitzPatrick's view directly, further work is necessary that is beyond the scope of this project. My main goal was to deny the need for akrasia in determining culpability, which rarely takes the seriousness of normative ignorance into account. In summary, I introduced the topic of culpable ignorance, the main questions the topic tries to answer, and the main general

answers that have come up in the field. Then, I introduced the specific debate between William FitzPatrick and Neil Levy on what is reasonable for us to expect of agents. Finally, I place myself in the discussion by supporting FitzPatrick's view and defending it against Levy's criticism. I do so by arguing that Levy's criticism and subsequent argument relies on false assumptions. I prove that Levy is working on false assumptions, but the work of strengthening FitzPatrick's argument remains. Strengthening FitzPatrick's argument is beyond the scope of this project, but I outline possible avenues to do so.

Having delineated my position, I want to leave on the note of apt displays of blame. Although I conclude that most people in the case of interest are blameworthy, that does not entail that they are all worthy candidates for blame. Blame is a reactive attitude characterized by reproach and resentment. Not all unwitting wrongful agents may be deemed to be a proper recipient of this emotion. Perhaps the agent who told a sexist joke to her coworkers realized her mistake after speaking with her Human Resource team, and decided to genuinely apologize for it. Although she was culpable for her act, her coworkers may not judge that she needs to be blamed for it. Perhaps she was racked with guilt after finding out her wrongs, so her coworkers judge that she is showing remorse and blaming her would only be piling up on her. There are many reasons to suspend feelings of blame, but mere suspension does not suggest that an agent is not blameworthy.

Works Cited

Arpaly, Nomy. *Unprincipled Virtue: An Inquiry Into Moral Agency*. Oxford University Press, 2002.

Fischer, John Martin, and Neal A. Tognazzini. "The Truth about Tracing." *Noûs*, vol. 43, no. 3, Sept. 2009, pp. 531–556, www.jstor.org/stable/40267354.

FitzPatrick, William. "Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge." *Ethics*, vol. 118, no. 4, July 2008, pp. 589–613, www.jstor.org/stable/10.1086/589532.

Haji, Ishtiyaque. "An Epistemic Dimension of Blameworthiness." *Philosophy and Phenomenological Research*, vol. 57, no. 3, Sept. 1997, pp. 523–544, www.jstor.org/stable/2953747.

Harman, Elizabeth. "Does Moral Ignorance Exculpate?" *Ratio*, vol. 24, Dec. 2011, pp. 443–468.

Levy, Neil. "Culpable Ignorance and Moral Responsibility: A Reply to FitzPatrick." *Ethics*, vol. 119, no. 4, July 2009, pp. 729–741, www.jstor.org/stable/10.1086/605018.

Mills, Charles Wright. *The Racial Contract*. Cornell University Press, 1999.

Montmarquet, James A. "Culpable Ignorance and Excuses." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, vol. 80, no. 1, Oct. 1995, pp. 41–49, www.jstor.org/stable/4320619.

Moody-Adams, Michele M. "Culture, Responsibility, and Affected Ignorance." *Ethics*, vol. 104, no. 2, Jan. 1994, pp. 291–309, www.jstor.org/stable/2381578.

Peels, Rik. "Tracing Culpable Ignorance." *Logos and Episteme*, vol. 2, no. 4, 2011, pp. 575–582.

Robichaud, Philip. "On Culpable Ignorance and Akrasia." *Ethics*, vol. 125, 2014, pp. 137–151, doi:10.1086/677139.

Rosen, Gideon. "Culpability and Ignorance." *Proceedings of the Aristotelian Society*, vol. 103, 2003, pp. 61–84, www.jstor.org/stable/4545386.

Rosen, Gideon. "Skepticism about Moral Responsibility." *Philosophical Perspectives*, vol. 18, no. Ethics, 2004, pp. 295–313, www.jstor.org/stable/3840937.

Rudy-Hiller, Fernando. "The Epistemic Condition for Moral Responsibility." *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, Winter 2022, plato.stanford.edu/archives/win2022/entries/moral-responsibility-epistemic/.

Smith, Holly. “Culpable Ignorance.” *The Philosophical Review*, vol. 92, no. 4, Oct. 1983, pp. 543–571, www.jstor.org/stable/2184880.

Strawson, P F. “Freedom and Resentment.” *Proceedings of the British Academy* , vol. 48, 1963, pp. 187–211.

Talbert, Matthew. “Moral Competence, Moral Blame, and Protest.” *The Journal of Ethics*, vol. 16, no. 1, Mar. 2012, pp. 89–109, www.jstor.org/stable/41486948.

Timpe, Kevin. “Tracing and the Epistemic Condition on Moral Responsibility.” *Free Will and Moral Responsibility*, vol. 88, no. 1/2, Jan. 2011, pp. 5–28, doi.org/10.5840/schoolman2011881/22.

Zimmerman, Michael J. “Moral Responsibility and Ignorance.” *Ethics*, vol. 107, no. 3, Apr. 1997, pp. 410–426, www.jstor.org/stable/2382323.