

Revisiting the Relationship Between Moral  
Responsibility and Reasonable Expectations

Adam Karpman

A Major Research Paper in  
The Department of Philosophy

Presented in Partial Fulfillment of the  
Requirements for the Degree of Master of Arts  
(Philosophy) at Concordia University,  
Montreal, Quebec, Canada.

August 2023

© Adam Karpman 2023

## **ABSTRACT**

### Revisiting the Relationship Between Moral Responsibility and Reasonable Expectations

Adam Karpman

In this paper I build upon the capacitarian account of moral responsibility, whose main tenet is that in order to hold someone blameworthy for a wrong they commit, it must be the case that the agent could and should have known better and done otherwise. I first suggest, on grounds of fairness, that the capacitarian account is more plausible than its strongest competitors, notably the attributionist and volitionist account. Proceeding this, I argue that many of the most prominent capacitarian accounts fall short of encompassing all wrongs an agent can fairly avoid, and so proceed to offer my own amended capacitarian account that formulates an expanded conception of reasonable expectations that I argue is vital to our understanding of moral responsibility. This new formulation of reasonable expectations considers four variables specific to the agent in question: Capabilities, Environmental Influences, Access to Information, and Professionally and Socially Assumed Roles. I use these variables, in conjunction with a focus on effort and risk to define the condition for moral responsibility of a wrong. I conclude by considering some of the practical consequences of assenting to this theory of moral responsibility, and argue that despite certain uncertainties in practice, this account nonetheless leads to the humanization of those around us, even those who have wronged.

## **Introduction**

The act of blaming is a central practice to our social lives as humans; we blame friends, family, and acquaintances when they are responsible for certain wrongs, we write papers and articles blaming public figures for different transgressions. We may agree or disagree with attributions of blame made by people, but what makes blame and the hot emotions that come with it justifiable in certain cases and not in others? When we examine blame in the normative sense<sup>1</sup>, it is widely agreed that it is appropriate when an agent can rightfully be held morally responsible for the wrong they have committed. But just as quickly as we have answered that first question, we find that we have only reformulated the question to: *When is an agent morally responsible for a wrong they have committed?* Addressing this will be the central focus of this work. Within the literature, there are many approaches different camps of philosophers have taken to address this question, and I will be examining the landscape of the literature in the following section. Following which, my central focus will be on refining the *capacitarian* position. In general, capacitarians believe that to assign moral responsibility to an agent, we must consider their capacities in order to determine whether they could have and should have done differently. Fundamental to this position is the distinction between what one can do, in the very broad metaphysical sense of possibility, from what they can reasonably be expected to do. I will aim to refine the capacitarian position by providing a criticism of some of the most recent contributions to the capacitarian literature. My work, building off the ideas of FitzPatrick (2008, 2017), will develop an expanded idea of reasonable expectations that will be foundational to the assignment of moral responsibility. In doing so I will work to propose a capacitarian-inspired approach to moral responsibility that incorporates the insights of various authors (ex: Ayars 2021; FitzPatrick 2008, 2017; Rudy-Hiller 2017; Smith 2006) while dealing with some of the

internal issues, most notably that the scope of moral responsibility is either too narrow, too broad, or is inconsistent with the principles which ground the capacitarian position. In §I I will be exploring three of the most prominent theories of moral responsibility, these being the volitionist, attributionist, and capacitarian accounts. Then, on grounds of *fairness*, I will argue that the capacitarian account is the most favorable position. In §II I will build off many of the arguments presented in the previous section to argue that although it is the most favorable position, it must be slightly amended from the most prominent capacitarian accounts in the literature to be able to encompass all instances of what we can fairly expect someone to do. In §III I will propose my own capacitarian account that provides the revisions needed to strengthen the capacitarian position. Finally, in §IV I will explore some of the implications of this amended capacitarian account of moral responsibility.

## **I - A Survey of the Literature: Different Theories of Moral Responsibility**

### **I.i - A Quick Preliminary Discussion of Blame**

Following the trend in the literature, I will hold that the negative instantiation of moral responsibility and blameworthiness are fundamentally intertwined, such that one logically implies the other (see: Ayars 2021, 56; FitzPatrick 2008, 591; Rosen 2004, 296; Talbert 2017, 47)<sup>ii</sup>. Rosen, whose work will be examined shortly, summarizes this clearly when he writes “Since I will be focusing exclusively on bad acts and their consequences, we may simply identify moral responsibility with culpability or blameworthiness” (Rosen 2004, 296). I will not be significantly concerned with the positive instantiation of moral responsibility (praiseworthiness) in this work. An agent can be an appropriate target of blame when and only when the agent is

blameworthy. So, regarding blame itself, there are two distinct aspects of it that are worth identifying right at the outset, the *emotional aspect* of blame, and the *functional impact* of blame. Regarding the emotional aspect, I follow the trend of recognizing blameworthiness as being disposed to certain ‘negative’ emotions from others, most notably anger, resentment, and indignation (ex: Ayars 2021, Strawson 1962). When these confrontations do occur, we can call this the manifestation of blame. But blame also plays a relevant social role. The act of blaming often helps to convey our values to those around us, and aids in the convergence of values with those in our social communities (ex: Fricker 2014, Shoemaker and Vargas 2019). I acknowledge that both the emotional and functional role are important features in our understanding of blame.

### **I.ii - The Attributionist Account and its Shortcomings**

Let us now begin by exploring the attributionist account. When a wrong has been committed, attributionists like Talbert (2017) maintain that if an agent is responsive to the reasons and consequences of their wrong actions, then that wrong is attributable to them and they are morally responsible for those actions. Talbert writes that “for attributionists, an agent is open to moral blame when her actions, omissions, attitudes, patterns of awareness, and so on reflect objectionable evaluative judgments” (Talbert 2017, 48). Therefore, it is not necessary that the agent understand their action as *wrong*, but simply that they know why they are doing it and what the consequences of the act will likely be (Talbert 2017, 48). The attributionist position is seemingly the simplest, with the smallest room for gray areas:

If a wrong X is attributable to agent Y, then agent Y is morally responsible (and hence blameworthy) for wrong X<sup>iii</sup>.

All we have to worry about is the antecedent. Given that attributability only demands responsiveness to the reasons and consequences of an action, without any concern for how easy or difficult alternative courses of action would be, in most cases it should be unproblematic to determine if the antecedent does in fact hold.

This simplicity comes at a cost, however. Consider the case of someone having grown up in a sheltered and dogmatic community where they were constantly instilled the value that the environmental degradation that is occurring is to be of no concern to them and their community, and consequently this agent shows no concern for any environmental initiatives. We may see their lack of concern as wrong on some objective level, but we would recognize that the attitude is near inevitable until they have been confronted with alternate ideas. Holding someone morally responsible for an inevitability is a result of the attributionist disregarding the question of *was the wrong avoidable* and instead asking only the question *was the wrong attributable*<sup>iv</sup>. This is problematic because it removes a significant degree of agency in receiving/avoiding blame.

If we are sympathetic to this line of reasoning, it is because we believe that moral responsibility must be grounded in an ideal of *fairness*, in that to hold someone morally responsible for an action or belief, it is fair to have expected them to do otherwise<sup>v</sup>. It is this ideal that the attributionist fails to satisfy. This ideal is essential because without it, blame boils down to a game of pointing at the manifestation of wrong attitudes that express objectionable evaluative judgments. In doing so the social impact of blame is severely damaged, as it will cause more frustration on the part of the blamed agent who couldn't have done otherwise, rather than aiding in the convergence of moral values.

It should be noted that the attributionist can simply deny that moral responsibility is grounded in fairness, and instead argue that it is grounded in some other way (perhaps accepting

a more retributive ideal). Both on a practical and theoretical level, this seems like a highly problematic stance to take, as excuses such as ‘I couldn’t have possibly done anything else’ would no longer be a valid deterrent of blame. However, due to the foundational nature of the ideals that ground moral responsibility, and its arguably somewhat intuitive appeals, I will not endeavour to further dissuade the attributionist who denies the ideal of fairness. Nonetheless, I hope to have shown why those who do accept that moral responsibility is grounded in fairness must look elsewhere for a coherent account.

### **I.iii - The Volitionist Account and its Skeptical Upshot**

With fairness in mind, let us consider the volitionist account. The volitionist account I will explore is presented by Rosen in his article *Skepticism About Moral Responsibility* (2004) and is primarily concerned with actions done from ignorance. Although this section and the following one (§I.iv) focus largely on culpable ignorance (i.e. blameworthy ignorance), I will ultimately aim (in §III) to propose an account that encompasses all wrongs, not just those that stem from ignorance<sup>vi</sup>. Nevertheless, a high degree of focus on culpable ignorance will be fruitful since this is one of the most difficult instantiations of moral responsibility to understand.

Volitionists believe that to hold an agent morally responsible for an act or belief X, they must have acted akratically in the present or in a previous moment leading up to X, where *akrasia* is defined as the act of going against one’s better judgment (Rosen 2004, 301). Such that, even if an agent is *currently* ignorant of a pertinent truth (whether normative or factual) which causes them to commit a wrong, they may still be morally responsible for that wrong if there was an akratic moment that explains their current ignorance or negligence.

Let's take another simple example. If Alex was tired and decided to skip class on Tuesday despite knowing full well that this was not the prudent thing to do, and because of this he failed to receive proper instructions on his responsibilities for a group project, his classmates might blame him for failing to hold up his end of the stick. Although he might respond, 'how could you blame me, I didn't even know what I was meant to do', his classmates could respond, 'well you decided last Tuesday to skip class, knowing full well that you should have shown up particularly because we were trying to organize this project'. To a volitionist, this would be a satisfactory attribution of blame. Although Alex's ignorance doesn't stem from a present instance of akrasia, it can very easily be traced back to a past instance of akrasia. If, when we trace back, there is no identifiable moment of akrasia, that is when the agent should not be considered morally responsible for a wrong.

Interestingly however, although in his 2004 article Rosen argues in favor of the volitionist approach to moral responsibility, he is skeptical that it can be used much in practice. As it concerns judging others and oneself, Rosen writes that "I suggest that given the opacity of the mind—of other minds and even of one's own mind—it is almost always unreasonable to place significant confidence in such a judgment [of responsibility]" (Rosen 2004, 308). With regards to introspective attributions of moral responsibility, Rosen expands on this by saying that:

When I consider my own case and ask whether some weakish act of mine amounts to genuine akrasia as opposed to ordinary moral weakness, I have no trouble identifying tolerably clear cases of the latter; but I confess that I cannot identify clear examples of the former with any confidence. (Rosen 2004, 309)

Given the difficulty in identifying akratic acts both in oneself and in others, Rosen argues that this should lead us to accept a significant degree of skepticism about moral responsibility.



This argument should be of concern to anyone who thinks that attributions of moral responsibility are often justified.

#### **I.iv - The Capacitarian Account as a Response to the Volitionist**

This volitionist account leaves us with two things to ponder: do we find the account satisfactory at a theoretical level, and if so do we think these skeptical worries follow when this account is put into practice. In his article *Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge*, FitzPatrick (2008) addresses both these thoughts. With regards to the latter FitzPatrick questions the sweeping empirical claims Rosen seems to be basing his arguments on. Regarding identifying one's own acts of akrasia, Fitzpatrick thinks there are many clear cases where one knowingly goes against their better judgment, examples of which would include procrastinating work or eating unhealthily (FitzPatrick 2008, 594-595). In such cases, we typically go against what we ourselves see as the prudent thing to do. Moreover, FitzPatrick points out that certain emotional responses are evidence that we have acted akratically. He writes that "The presence of guilt or shame at the time of acting is therefore often good evidence of akrasia, which diminishes the plausibility of general skepticism about attributions of akrasia" (FitzPatrick 2008, 595).

But regarding the first and more fundamental question, FitzPatrick argues that the volitionist account is lacking at a theoretical level. To begin with, he believes the question of blameworthiness for ignorance, for both himself and the volitionists can be understood as the 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?" (FitzPatrick 2008, 603). Whereas Rosen identifies akratic acts as being the line in the sand, Fitzpatrick argues that we can make claims of reasonable expectations without resorting to something as strong as akrasia (FitzPatrick

2008, 603). To help illustrate where the volitionist account is lacking, FitzPatrick discusses the fictional case of Mr. Potter, from the beloved Christmas special *It's a Wonderful Life*. The example follows that due to an unideal upbringing where faulty moral norms were pushed upon him, Mr. Potter never went against his better judgment in becoming the morally dubious businessman he is today. Henceforth, according to the volitionist view, we would have to concede that in this case he wouldn't be to blame for his current faulty behaviors and beliefs (FitzPatrick 2008, 599-601). But can this be so? Can we truly conclude that Mr. Potter is not to blame because when we trace back we cannot identify an akratic act? I follow FitzPatrick in rejecting this conclusion. FitzPatrick believes that it is enough that Potter could reasonably have been expected to correct his moral ignorance, particularly due to his high degree of intellect made evident by his successful business practices, but instead chose to act in ways which strengthened his faulty moral compass (FitzPatrick 2008, 608; FitzPatrick 2017, 41). In short, FitzPatrick sees Rosen's focus on akrasia as overly narrowing the scope of moral responsibility, and so tries to extend it, leading to his final formulation of culpable ignorance, which in the context of our discussion is worth quoting in full:

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

Although FitzPatrick does not explicitly classify himself as a capacitarian, his account can appropriately be labelled as such. When an agent commits a wrong, at the most general level

capacitarians agree that if an agent “could have and should have known better” (Ayars 2021, 64), then this is enough for blameworthiness. With the stark difference being that volitionists hold that for an agent to be morally responsibly they *did* know better, either now or at some point in the past.

## **II - Engaging with the Capacitarian Literature**

### **II.i – Rational Capacitarianism**

FitzPatrick makes a compelling case that fairness in holding people accountable for their ignorance does not demand an instance of akrasia. He instead extends it to a disjunctive view allowing either akrasia or the voluntary (and hence culpable) exercise of nonakratic vices to be the explanatory factor in identifying instances of culpable ignorance, where these vices were understood to include things such as overconfidence, arrogance, dismissiveness, laziness, etc. (FitzPatrick 2008, 609). From here on out, given their nature in hindering our epistemic activities, I will use the term *epistemic vices* to capture the vices FitzPatrick is here referring to.

So, we must now consider whether this new disjunctive account fully satisfies all instances of culpability. I believe it does not, and I would like to argue this by providing the counterexample of a class of actions that we can label *avoidable honest mistakes*. This is a class of action that does not stem from an epistemic vice, nor does it trace back to an instance of akrasia. FitzPatrick’s account denies that an agent could ever be morally responsible for this type of action, and hence not liable to blame for it. But both on an intuitive and theoretical level, it seems highly doubtful that this could be so. This is in part because we often have no problem judging that an agent could have fairly avoided such a mistake. To illustrate this, it is worth

exploring Alisabeth Ayars' account of *Rational Capacitarianism* from her article *Blaming for Unreasonableness* (Ayars 2021).

Ayars' work focuses on pushing back against the *quality of will condition* which holds that ill will is a *necessary* condition for someone to be held morally responsible for a wrong that they commit. Ayars understands "Ill will [as], roughly, an objectionable lack of concern for morality or the morally significant interests of others", although she acknowledges that the exact definition is controversial in the relevant literature (Ayars 2021, 56)<sup>vii</sup>. Ayars argues that the ill will condition does not hold true to everyday moral practice, nor is it compelling at a theoretical level. For example, in many cases of ignorance, ill will is absent, but it may still be reasonable to have expected the agent to avoid such a wrong and hence be morally responsible for it. To illustrate this point, Ayars begins with a few examples, and she asks us "to imagine the characters [in the examples] as being as goodwilled as they could possibly be—as the sort of people who really are concerned with morality and others' interest as much as they should be, but who suffer from bouts of ordinary unmotivated irrationality" (Ayars 2021, 61). Although Ayars' expressed goal is to push back against the ill-will theorists, we will see how it equally pressures the account from FitzPatrick that we have just explored.

Ayars titles one of her examples *Bridge Collapse* (Ayars 2021, 62). The example states that an engineer, Marie, is conducting a stress test on a local bridge, and must make a judgment call as to whether the bridge should be closed during this time as a precaution. This decision rests on the judgment of the engineer, and Marie is a responsible engineer who in the past has always been prudent and closed the bridge when she thought the particular stress test would be too risky (Ayars 2021, 62). However, in this particular instance she makes a grave error of judgment, "Given the magnitude of the stress test being conducted and the age of the bridge, it would in

fact be *highly* imprudent not to close it... She is fully capable of drawing this inference; but she does not” (Ayars 2021, 62). As a result of this poor decision, the bridge collapses, and a pedestrian is killed. Ayars goes on establishing that Marie did not display an improper pattern of concern, and she explicitly denies that the act can be attributed to the kind of epistemic vice FitzPatrick (2008) spoke of<sup>viii</sup>. Now we arrive back at our question, is Marie blameworthy for this act? Ayars argues that, yes, because Marie *should* have known better. And the relevant sense of should, for Ayars, is that “‘X rationally should have known better.’ Her epistemic performance fell below some operative epistemic standard” (Ayars 2021, 65). It is building off of these considerations that Ayars gives her final formulation of *Rational Capacitarianism*, which again, is worth quoting in full:

*Conditions of Blameworthiness:* X is blameworthy for a wrong act A iff X’s doing A expressed an unreasonable attitude that is attributable to her.

*Theory of Attributability:* An attitude is attributable to X iff X has the capacity to assess the reasons for and against the attitude and modulate her attitude in light of this assessment during the period in which the attitude was formed and sustained. (Ayars 2021, 71)

Ayars states that both unreasonable ignorance and ill will qualify as ways of acting from an unreasonable attitude attributable to the agent in question (Ayars 2021, 68). By isolating *unreasonableness* in unison with *rationality* as together being the centerpiece for attributions of blame, Ayars manages to put forward a compelling theory of blameworthiness that manages to account for *avoidable honest mistake* while also being mindful not to make the scope of moral responsibility overly far-reaching.

## II.ii – Blame in Degrees

Before we proceed, I think a note regarding blame for wrong acts that do not stem from ill will, akrasia, or an epistemic vice would be helpful. It would be easy to feel that such attributions of blame would be too harsh. If a well intended friend makes an honest yet grave mistake, would it not be a little excessive to blame them for their action? This thought becomes less problematic when we acknowledge that blame comes in *degrees* and is not an all or nothing act. Undoubtedly, the degree of blame will usually rise sharply when ill will or akrasia is established. A friend accidentally and inattentively walking into your dog may be a relatively insignificant transgression, where a quick apology would amend any tension on the relationship. However, if ill will is established and you know your friend acted deliberately, the degree of blame will inevitably be significantly higher and may cause irreparable damage to the relationship.

A complete examination of the degrees of blame extends beyond the scope of this work, but I would like to briefly state at the outset a few preliminaries about this in relation the account I will shortly propose. The blame that results from being deemed morally responsible should come in degrees based on three primary factors: how easy or difficult the alternative action was, the severity of the situation, and finally whether the agent displayed ill-will, akrasia, or an epistemic vice.

Regarding the first point, even if we make the judgment that it is fair to have expected the agent to do otherwise, there will be situations where the alternative was much easier to do, in the sense of less effort and risk, and so the failure to realize this alternative would call for higher degrees of blame.

Second, we must consider the seriousness of a situation. When the stakes are higher, we generally expect a higher degree of attentiveness on the agent's part, and so the failure to perform some relevant act would call for higher degrees of blame. Therefore, a doctor forgetting about a coffee date with a friend would be less blameworthy than forgetting about a medical procedure they were scheduled to perform.

Lastly, as was just discussed, the manifestation of ill-will, akrasia, or an epistemic vice induces more blame because it conveys something more fundamentally problematic about the moral integrity of the agent in question. This holds particularly true for ill-will or akrasia, whereas some epistemic vices might be more easily overlooked (such as a brief bout of laziness).

These considerations will be more easily understood after reading the full account in §III.

### **II.iii – Response to Ayars**

Returning to Ayars' work, I agree insofar as the examples she provides of honest mistakes do in fact demonstrate cases of a blameworthy moral agent, but my disagreement lies at the theoretical explanation that justifies those intuitions. Let me follow the same approach Ayars takes in her work by first showing that her account fails to account for many instances of blameworthiness that occur in day-to-day life. Let's take a hypothetical example.

Consider the case of a competent and smart individual, Bob, who is completely unsympathetic to the struggles of the less fortunate. When we trace back in time, however, we find that this attitude is rooted in a failure to sympathise with the homeless and impoverished that he had encountered throughout his life, in the sense that he never went through the emotional process of understanding the experience and struggles of those people. Now the question

becomes, can we blame Bob for his morally dubious attitudes towards the poor? Assuming there is no relevant additional information regarding Bob's capacities or environment that might sway us (ex: arguably and somewhat controversially, a sociopathic or psychopathic diagnosis), I believe the answer is yes. Since Bob has nondeficient emotional capacities, there seems to be no good reason we could not have fairly expected him to utilize those capacities. This, however, does not come down to a purely rational mistake. The target of our blame is primarily the emotional failure, rather than a rational one. This is but one example, and I think it is not too difficult to think of a whole range of similar day to day cases of blameworthiness that are rooted in emotional failures.

Of course, the strength of these intuitions will ultimately come down to judging the merit of theory that grounds them, and I will proceed to put forth my theory. But first I want to mention the common mistake (as I see it) that is present in all three of the previously examined theories, as well as many of the other most prominent accounts in the literature, which ultimately leaves them open to counterexamples. Rosen, FitzPatrick, and Ayars, to different extents, recognize the importance of *fairness* in attributions of moral responsibility, and so try to make an account that can accommodate this grounding principle<sup>ix</sup>. But they all proceed by identifying very specific classes of attitudes that they believe are the ones an agent can fairly avoid. For Rosen, it was instances of akrasia. For FitzPatrick, it was instances of akrasia or failures that stem from an epistemic vice. For Ayars, it was unreasonable attitudes that stem from a rational failure. The inclination to highlight a narrow set of attitudes is understandable, as it makes for a well-knit theory. But as soon as we focus in on a chosen set of attitudes, the counterexamples start to mount. These counterexamples hold weight because, in day-to-day moral practice, they seem to present clear cases of blameworthiness, and additionally when we ask ourselves 'could



the moral agent have fairly avoided this wrong', the answer seems to be yes. So, we need a theory that can account for this.

### **III - An Expanded Conception of Reasonable Expectations**

#### **III.i - The New Account**

As I have showcased in the previous two sections, there are many factors that are traditionally overlooked when making judgments of when one can fairly be expected to avoid a wrong, and hence be morally responsible for that wrong. As such, there is a significant need for an expanded theory of moral responsibility which is not concerned solely with identifying specific attitudes, but rather aims to understand the wider context of the agent in question. I will propose such an account, and I will identify four relevant variables that need to be considered in judgments of moral responsibility. I take this proposed account to encompass all wrongs, including those that stem from ignorance. Let me begin by identifying the condition for moral responsibility.

*Condition for Moral Responsibility of a Wrong:* When an agent commits a wrong, whether manifested in an action or belief, they can be held morally responsible (and hence blameworthy) when and only when they can reasonably have been expected to do otherwise<sup>x</sup>.

The formulation above immediately brings up the question of reasonable expectations, which can be expanded on as follows:

*Reasonable Expectations:* We can reasonably expect a moral agent to do otherwise when and only when, considering the four relevant variables particular to them, the agent could have avoided the wrong without an amount of effort or risk that exceeds what the situation demands. These four variables are:

- 1- Capabilities
- 2- Environmental Influences
- 3- Access to Information
- 4- Professionally and Socially Assumed Roles

I take this list of variables to be an exhaustive list of what we must consider when making judgments of reasonable expectations, and this is possible in part because the categories are broad. Moreover, this formulation of reasonable expectations is sensitive to time, such that a past failure may make an agent morally responsible for a present faulty action or belief. Let us now move on to examine each of these variables in more detail.

***Capabilities<sup>vi</sup>:*** This variable concerns the physical and mental conditions of the agent. If the person in question has deficient or proficient physical attributes, then this surely shapes what can be expected of them. For example, an Olympic swimmer would be expected to be able to help a child drowning at a small local pool, whereas someone with difficulty walking would not be. The same holds true for parallel situations that involve mental abilities. Here we are concerned with those things which are more specific to the agent themselves, rather than the environmental factors that allows them to fully exercise those capabilities, which we shall cover in this next point.

***Environmental Influences:*** Under this heading I include all environmental influences that may impact an agent's beliefs and actions therefore affecting our evaluation of their blameworthiness. This is an extensive list that includes things such as: location, socioeconomic status, exposure to alternate ways of thinking, biases held by close family and friends, etc. All these factors may afford one greater or less ability in making good decisions and avoiding moral wrongs, in other words, they can determine whether alternatives to the moral wrong are indeed accessible to one. Considering the environment of the individuals in question helps us move away from the assumption that we can universally expect people to avoid certain problematic actions and beliefs. This allows us to recognize that even two agents with similar capabilities may face different challenges in order to avoid the same wrong based on the influences that surround them<sup>xii</sup>.

***Access to Information:*** Related to the previous point, it is essential that we consider what information is available to the agent in question to enable them to make good decisions. This is relevant in two primary ways. Firstly, we must consider things concerning the agent's whereabouts, such as: access to internet, smartphones, libraries, teachings etc. The second way in which this is relevant is we must consider the state of knowledge in the world more generally at the time when the agent was expected to act. This is particularly relevant when trying to assign blame retroactively. Consequently, certain faulty judgments that could have blamelessly been upheld previously, such as an agent in the 1950s believing there are no risks associated with smoking following their doctor's advice, would become blameworthy if still upheld today.

***Professionally and Socially Assumed Roles:*** Our roles, such as doctor, pilot, friend, parent, also inevitably affect what can reasonably be expected of us in a given situation. We assume roles both professionally and socially. The former type concerns roles that we take up in a professional

setting (ex: doctor, pilot), and the latter concerns roles we take up interpersonally, outside of any professional commitment (ex: friend, parent). In the previous *bridge collapse* example for instance, the agent had assumed the professional role of being hyper vigilant in their calculations when approving a stress test for a bridge, and so their failure to meet these responsibilities may make them rightfully blameworthy. This would be different from, say, an intern who was allowed to look at the plans in passing but was not tasked with reviewing them.

Roles can be explicitly or implicitly assumed. In the bridge example, the engineer had explicitly assumed the role of checking the safety of the stress test. But there are many instances where responsibility falls upon an agent implicitly. Consider the case that you are visiting your friend who is babysitting their younger cousin that evening. Your friend injures themselves while cooking and has to be rushed to the hospital, leaving you at home with their cousin. In such a case it would be assumed that you will look after the child, even if you hadn't explicitly committed to that role. This would be an example of an implicitly assumed social role<sup>xiii</sup>.

After understanding the four variables, let us return to the formulation of reasonable expectations where, as I defined it, we must consider these variables to see if an *agent could have avoided the wrong without an amount of effort or risk that exceeds what the situation demands*. This line is essential in differentiating mere possibility from reasonable expectations. If someone lives in a highly sheltered and dogmatic environment, with no access to internet, it may be *possible* that they are able to move past certain faulty norms that were pushed upon them, but the amount of effort they would have to put in to achieve that may be beyond what we can reasonably expect someone to invest in educating themselves. This would differ from someone who lives in a multicultural metropolis with access to good education, internet and alternate

ways of thinking. Even if this individual had been taught certain faulty norms growing up (perhaps xenophobic attitudes towards co-citizens), the amount of effort it would take to correct their judgments would fall in the scope of effort appropriate to what the situation demands<sup>xiv</sup>.

Here the idea of *heroic acts* becomes relevant as well. We can define *heroic acts* as actions in which the agent does something good that exceeds the amount of effort and risk that the situation demands. Consider someone jumping onto the subway tracks to aid someone who had fallen, and consequently saving them while avoiding an oncoming train by a split second. This kind of heroic act is highly virtuous in part because the amount of effort and risk assumed by the agent exceeds what the situation demands of an onlooker. Such an act is rightfully praiseworthy, but the failure to perform such an act should not induce blame.

Different situations call for different levels of effort and risk; I believe we must always consider these four variables in relation to what the situation demands to gauge what can reasonably be expected of an agent.

### **III.ii – The Relation Between Reasonable Expectations and Empirical Data**

Some may ask the question of whether statistical data regarding the likelihood one will perform an act would affect our judgement of whether they are morally responsible for that act. Indeed, the question of the relationship between reasonable expectations and empirical data now becomes a central one. Let's say we see someone commit a wrong that we take to be fairly avoidable and thus liable to blame. Immediately after, however, it comes to our attention that we have statistical data showing that similar agents in similar situations, more often than not, commit the same wrong. So, if people seem to be prone to make this mistake, does that give us

some insight into how avoidable certain wrongs are? Reformulating this question, we can ask whether this data of people's likelihood of committing a similar wrong: completely negates blame, holds some amount of weight in our judgments of blame, or does it remain irrelevant to our judgments of blame.

I follow other capacitarions like Rudy-Hiller<sup>xv</sup> (2019) in maintaining that it is the second of these options that holds true, namely that empirical data about people's likelihood to commit a similar wrong holds some weight in our judgements of blame. Such data cannot completely negate attributions of blame for two primary reasons. For one, no amount of empirical data alone taken from previous situations can ever *fully* inform us on the amount of effort and risk that would be required by a particular agent to avoid a wrong in a new manifestation of a similar situation. This becomes evident when we consider the vast list of relevant variables that impacts reasonable expectations. Secondly, even if this data was accurate enough, we have no reason to suppose that there aren't situations where, more often than not, moral agents tend to perform a wrong even when the amount of effort and risk it would take to avoid that wrong given their situation falls within the scope of what we can reasonably expect of someone. The data may make this less likely, but what is essential to understand here is that it does not remove the logical possibility. Take the example of a particular professor whose students consistently submit their assignments late. This professor kept a record of the timeliness of submitted assignments by all their students throughout the class they are currently teaching, and it turns out that students have been late submitting their assignments roughly 70% of the time. Does this data alone tell us that the students cannot reasonably be expected to submit their assignments on time, and hence are not blameworthy for their fault? It might lend some plausibility to this conclusion, as the data might be explained by a flaw in the professor's teaching methods (perhaps a lack of discipline, or

unrealistic deadlines), but it may also very well just be the case that this professor's students are particularly lazy, causing them to fail to complete assignments they have ample opportunity to complete. The data alone cannot answer this question, such that if a particular student proceeded to argue that they are in fact not blameworthy for submitting an assignment late, we would have to consider the specific situation of that student (namely the four variables in conjunction with effort), not merely the data that had been amassed.

Although empirical data regarding the probability of certain behaviors cannot fully dictate attributions of blame, it is equally implausible that it is irrelevant to our judgments of blame. Rudy-Hiller makes this point salient when bringing forth the example of situationism, which is the study of how people's behavior is influenced by their surroundings (Rudy-Hiller 2019, 2951). As a particular subset of situationism, he takes the example of the bystander effect, first discovered by Bibb Latané and John Darley (1970). Rudy-Hiller summarizes the findings as follows:

whenever a person witnesses a potentially dangerous situation for herself or others, she is more likely to correctly interpret it as an emergency and to react appropriately to it—by either providing assistance herself or asking for help—when she is alone than when she is surrounded by other bystanders. In fact, the likelihood that a member of a group witnessing an emergency will provide assistance decreases as the number of people in the group increases. (Rudy-Hiller 2019, 2953)

There are two leading theories that explain this phenomenon. First, people rely on the reaction of others to gauge the seriousness of a situation, and thus if it is ambiguous and others don't respond, it may convey to oneself that the situation is not an emergency (Rudy-Hiller 2019, 2953). Secondly, it is possible that the increase in bystanders diffuses the degree of responsibility

people feel, since people recognize that they may just need someone's help, and given that there are so many people, not necessarily theirs (Rudy-Hiller 2019, 2953).

So, whereas we might at first be more content in attributing a high degree of blame to those involved in this sort of bystander situation, we can come to understand that there is a psychological phenomenon at play that could make it the case that the effort required to avoid the wrong (failure to act, in this case) is higher than what we might have originally thought. Our judgments of blame are therefore influenced by the situation itself and the *explanation* of why the wrong happened. I am therefore not claiming that the bystander effect, for example, would always exculpate the agent(s) in question. I am simply claiming that under certain conditions it could, and in others it might not fully excuse but merely reduce the degree of blame. Rudy-Hiller summarizes it perfectly when he writes "Base rates and predictive expectations, *in tandem with relevant explanations*, do affect normative expectations" (Rudy-Hiller 2019, 2957 my emphasis).

This, however, is not the end of the story. Returning to the example of situationism, even if we have good evidence to show that a wrong is virtually unavoidable in a given situation, this may not in itself exempt one from blame. This is because, although we may not be able to reasonably expect the agent to avoid the wrong in that situation, we may be able to reasonably expect them to be aware of this fact and avoid that situation in the first place. This is an instance where looking at the history of an event instead of merely the spotlight moment would be important. In the situationism literature, choosing which situations to put ourselves in is sometimes referred to as the *seek and avoid strategy* and is summarized nicely by Hagop Sarkissian when he writes "Individuals should seek situations that strengthen or support virtuous behavior, and avoid situations that tend toward vice or moral failure. In choosing situations, one chooses to embrace the behavioral tendencies they elicit" (Sarkissian 2010, 4). This, of course is



not always possible. As Sarkissian<sup>xvi</sup> points out, there may be situations where the behavioral tendencies that accompany it are either unknowable before hand, or there may be situations that are simply unavoidable in the first place (Sarkissian 2010, 5). However, considering the relevant variables specific to the agent in question, we can try and arrive at a judgment regarding whether the particular situation was reasonably avoidable, even if the wrong that the situation elicited once there was not.

#### **IV - Implications of My Approach to Moral Responsibility**

One of the virtues of my proposed theory of moral responsibility is that it is highly adaptive to every particular situation. The four variables in unison with the focus on effort and risk make it such that an agent can be deemed morally responsible (or not) based on a plethora of considerations. As such, it can adequately explain many situations that other theories cannot. However, if we assent to this theory of moral responsibility that builds upon reasonable expectations, it will readily become clear that many real-world situations will end up in a gray zone<sup>xvii</sup>. Absent of some form of new technology that allows us to better understand the lived experience of others, it will be very difficult in certain situations to judge how the capabilities of an agent and the way they are shaped by their environment influences what can reasonably be expected of them. It is also worth mentioning that people could accept the general schema while disagreeing on the application of it. Such that there can be complete theoretical agreement, while there being some disagreement on the application of parts of the theory (such as disagreeing on the amount of effort and risk a situation demands).

These difficulties may seem to some as a major downfall of this approach to moral responsibility. After all, if there are so many uncertainties in practice, how much good can it do at a theoretical level. But I alternatively hold that far from being a slight to this theory, it is in fact a virtue of it. When it is not immediately clear whether an agent could reasonably have been expected to do otherwise, this calls on us to exercise extra effort to really try to understand the perspective of the agent in question, which encourages modesty and thoughtfulness in our attributions of blame. And the mere act of exercising this effort humanizes the people we are trying to judge<sup>xviii</sup>. Rather than quickly judging them to be *idiots, cowards, selfish, lazy*, or any other such slander that we so commonly hear used, we are instead encouraged to see them as people with specific capabilities, unique histories, and particular responsibilities. Of course, this is not to say we will always be dissuaded from attributing blame, but I believe that even when we do, the intellectual process of considering this extensive list can have its benefits in the way we see and treat others.

In a world that many believe is becoming more and more divided, and where social media creates the perfect stomping grounds for quick uncritical attributions of blame that will immediately be visible to likeminded individuals, this account can help us resist this urge to belittle and stereotype those we disagree with. Even on substantive issues where we have a large degree of certainty that the opposing view is faulty, we may nonetheless come to understand that in many cases, we cannot reasonably expect certain agents to have avoided the dubious belief. Of course, we must also be humble and recognize that, in many situations, we should not overextend the certainty of our own beliefs and also understand that there may be many contradicting beliefs that are reasonable to hold<sup>xix</sup>.

To sum up, I have argued in this work that the capacitarian account is the most promising account of moral responsibility but is not without its own shortcomings. As such, I have offered my own expanded account of moral responsibility that draws its inspiration from many authors, most notably FitzPatrick (2008, 2017). With a focus on the many variables relevant to the agent in questions, as well as recognizing the level of effort and risk that a particular situation demands, I have endeavoured to offer an expanded conception of reasonable expectations that is both theoretically sound and practically useful. Minimally, even if the reader does not agree with all the substantive arguments I put forward, I hope they can benefit from the idea that judgments of moral responsibility require an understanding of others and the situation they are in, and that this understanding alone has the potential to do a little good.

## **Bibliography**

Aquinas, Thomas. *Summa Theologica*. Thomas More Pr 1981 (1485).

Ayars, Alisabeth A. "Blaming for Unreasonableness: Accountability without Ill Will." *Journal of Ethics and Social Philosophy* 19, no. 1 (2021): 56-79.

Faraci, David, and David Shoemaker. "Insanity, Deep Selves, and Moral Responsibility: The Case of JoJo." *Review of Philosophy and Psychology* 1, no. 3 (2010): 319-32.

FitzPatrick, William J. "Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge." *Ethics* 118, no. 4 (2008): 589-613.

- FitzPatrick, William J. “Unwitting Wrongdoing, Reasonable Expectations, and Blameworthiness.” In Robichaud and Wieland *Responsibility: The Epistemic Condition*. Oxford: Oxford University Press, 2017, 29-46.
- Frazer, Michael L. *The Enlightenment of Sympathy: Justice and the Moral Sentiments in the Eighteenth Century and Today*. New York: Oxford University Press, 2012.
- Fricke, Miranda. “What’s the Point of Blame? A Paradigm Based Explanation.” *Noûs* 2014: 165-83.
- Hume, David. *A Treatise of Human Nature*, London: Penguin Books, 1969 (1739).
- Latané, Bibb and John Darley. *The Unresponsive Bystander: Why Doesn't He Help?* New York: Prentice-Hall, 1970.
- Rosen, Gideon. “Skepticism about Moral Responsibility.” *Philosophical Perspectives* 18, no. 1 (2004): 295–313.
- Rosen, Gideon. “The Alethic Conception of Moral Responsibility.” In Clark, McKenna and Smith *The Nature of Moral Responsibility: New Essays*. Oxford: Oxford University Press, 2015, 65-88.
- Rudy-Hiller, Fernando. “A Capacitarian Account of Culpable Ignorance.” *Pacific Philosophical Quarterly* 98 (2017): 398-426.
- Rudy-Hiller, Fernando. “Reasonable Expectations, Moral Responsibility, and Empirical Data.” *Philosophical Studies* 177, no. 10 (2019): 2945–68.
- Sarkissian, Hagop. “Minor Tweaks, Major Payoffs: The Problems and Promise of Situationism in Moral Philosophy.” *Philosophers' Imprint* 10, no. 9 (2010): 1-15.

- Shoemaker, David. "Attributability, Answerability, and Accountability: Toward a Wider Theory of Moral Responsibility." *Ethics* 121, no. 3 (2011): 602-32.
- Shoemaker, David. "Response-Dependent Responsibility; Or, A Funny Thing Happened on the Way to Blame." *Philosophical Review* 126, no. 4 (2017): 481-527.
- Shoemaker, David, and Manuel Vargas. "Moral Torch Fishing: A Signaling Theory of Blame." *Noûs* 55, no. 3 (2019): 581–602.
- Smith, Angela M. "Control, Responsibility, and Moral Assessment." *Philosophical Studies* 138, no. 3 (2008): 367-92.
- Strawson, P. F. "Freedom and Resentment." In *Freedom and Resentment and Other Essays*. Abingdon, UK: Routledge, 2008 (1962), 1-28.
- Talbert, Matthew. "Akrasia, Awareness, and Blameworthiness." In Robichaud and Wieland *Responsibility: The Epistemic Condition*. Oxford: Oxford University Press, 2017, 47-63.
- Tognazzini, Neal and D. Justin Coates, "Blame." *The Stanford Encyclopedia of Philosophy* (Summer 2021 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/entries/blame/>.
- Wallace, R. Jay. *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press, 1994.
- Wieland, Jan Willem, and Philip Robichaud. *Responsibility: The Epistemic Condition*. Oxford: Oxford University Press, 2017.
- Wolf, Susan. *Freedom within reason*. Oxford: Oxford University Press, 1990.

---

<sup>i</sup> The type of blame that will be of concern in this work is *agential normative blame*, which is the type of blame directed at a moral agent for a failure to meet some relevant moral norm. In everyday usage of the word *blame*, we may use it to refer to other things. It may be used in a very broad causal sense (I blame the bad weather for my mood today), and it can often be directed at an inanimate object (I blame this dull knife for my difficulty cooking today). These other common usages of *blame* will not be relevant to what is discussed in this work.

<sup>ii</sup> FitzPatrick makes this equivalence between the negative instantiation of moral responsibility and blameworthiness apparent when he opens the first section of his work writing “The argument applies only to *responsibility or blameworthiness* for wrong acts, not to responsibility or praiseworthiness for right acts” (FitzPatrick 2008, 591 my emphasis).

Similarly, Ayars makes the equivalence implied when she opens her work stating, “Many theorists of moral responsibility endorse the principle that *X* is blameworthy for a wrong act *A* only if in doing *A*, *X* expressed ill will” (Ayars 2021, 56).

Talbert also notes, on his very first page that “I will be concerned in this chapter with moral responsibility in the sense of blameworthiness” (Talbert 2017, 47).

<sup>iii</sup> Attributability of the wrong in this case is a sufficient condition for blameworthiness. I understand Talbert to hold this position as in the passage just quoted, he writes “for attributionists, an agent is open to moral blame *when* her actions...” (Talbert 2017, 48 my emphasis). Equally, in a surrounding passage he writes “while I am not prepared to say that circumstantial awareness is *necessary* for responsibility, I will argue that, in many garden-variety cases, the blameworthiness of agents who lack awareness is at least called into question” (Talbert 2017, 48 emphasis from text).

<sup>iv</sup> To further appreciate the complexities the attributionist fails to account for, see David Shoemaker’s work *Attributability, Answerability, and Accountability: Toward a Wider Theory of Moral Responsibility* (2011). In this work Shoemaker compellingly argues that attributability, answerability, and accountability are three distinct sorts of responsibility, and should not all be conflated with moral responsibility and blameworthiness more generally.

<sup>v</sup> This definition is seemingly circular, but I follow FitzPatrick in maintain that what we mean when we say an agent can fairly have been expected to avoid a wrong is that it is reasonable to have expected them to avoid a wrong

---

(FitzPatrick 2008, 603). Understanding what it then means to be able to reasonably expect someone to avoid a wrong will ultimately be a primary focus of this work, and my view will be apparent after reading §III.

<sup>vi</sup> Although, depending how we define *ignorance*, one could argue all wrongs stem from some form of ignorance.

Whether this be ignorance of some relevant moral norm; ignorance of an empirical fact; or ignorance of the strength that a norm puts on one's actions. This semantic issue is of no consequence to the account I will ultimately propose since I aim to encompass all wrongs.

<sup>vii</sup> It is worth noting here that this concept of ill will has overlap with the previously discussed concept of *akrasia*, although they are not logically equivalent.

<sup>viii</sup> It is only in footnote 16 of Ayars's work that she explicitly refers to FitzPatrick's 2008 article. In that footnote Ayars writes "Nor is Marie's ignorance the result of an epistemic vice. FitzPatrick argues that we can be responsible for ignorance that was the result of 'voluntary exercise of vices;'... this is not one of these cases" (Ayars 2021, 63).

<sup>ix</sup> Ayars work is the least explicitly concerned with fairness, but near the end of the work she addresses how her theory is compatible with it. She writes "Agents can avoid being the target of resentment by exercising their capacity to recognize reasons for believing and desiring and to be responsive in their beliefs and desires to their assessment of reasons. This is enough, it seems, for a fair opportunity to avoid" (Ayars 2021, 77).

<sup>x</sup> The use of the word *otherwise* refers to avoiding or significantly reducing the scale of the wrong.

<sup>xi</sup> It should be noted that I am deliberately not using the term *capacities* here, or anywhere in this section. Many authors use the term capacity in the more general sense of ability, which can extend beyond the internal conditions of the agent and be made to include specific surrounding influences. To avoid confusion with these other uses of the term *capacity*, I have omitted it.

<sup>xii</sup> A significant subset of environmental influences is *cultural norms*. Cultural norms have a dual impact, in that they can both impact what can be expected of a given agent, and, arguably, it can affect normative standards more generally. I am not concerned here with defining a wrong, so this second point falls outside the scope of what is relevant in this work. However, this former impact of cultural norms is immediately relevant to what is here being discussed.

<sup>xiii</sup> Some roles are egregious, such as being a willing militant in a tyrannical society. An agent may be blameworthy for holding a morally reprehensible role and the wrongs that the role elicits by judging whether the agent can reasonably have been expected to avoid it. In the last paragraph of §III.ii I consider this in relation to the environment someone finds themselves in. Those same considerations apply here.

<sup>xiv</sup> FitzPatrick puts forth a similar example in his work *Unwitting Wrongdoing, Reasonable Expectations, and Blameworthiness* (FitzPatrick 2017, 34).

<sup>xv</sup> Fernando Rudy-Hiller's presents his own capacitarian account of culpable ignorance in his 2017 article, *A Capacitarian Account of Culpable Ignorance*. Capacitarian culpable ignorance is defined by Rudy-Hiller as follows:

An agent S is culpably ignorant of some pertinent truth p if, and only if, S's ignorance of p is either: i) the foreseeable upshot of having disregarded an epistemic duty of care he had the abilities and fair opportunity to comply with; or ii) results from his failure to exercise an awareness-related ability despite having a fair opportunity to do so. (Rudy-Hiller 2017, 414)

Here awareness-related abilities are a set of psychological abilities, such as the capability to be aware of moral and nonmoral considerations, to be able to question, decide, and remember (Rudy Hiller 2017, 405).

A complete comparison of this account from the one I have offered extends beyond the scope of what I can cover in this work. However, my account has significant differences, ones which I believe increase its plausibility. For one, my account doesn't rely on the concept of 'duties of care' which I see as overly narrowing the scope of moral responsibility. Additionally, I believe that my expanded conception of reasonable expectations makes my account better equipped to handle the varying degrees to which an agent may possess certain abilities. However, a full defence of these claims would have to be explored at length, something I will not endeavour to do here.

<sup>xvi</sup> Hagop Sarkissian, in his article *Minor Tweaks, Major Payoffs* (2010), is not only concerned with how a situation affects an agent but also how the agent can impact the situation they are in. Therefore, there is a big stress on moderating one's own behavior because "influencing how situations unfold begins with minding the cues arising from one's person" (Sarkissian 2010, 9).

<sup>xvii</sup> FitzPatrick recognizes these same difficulties in the context of his account. He writes that:



---

The real question thus becomes this: Where do we draw the line between such genuine hard cases, where we must withhold blame for false beliefs and resulting wrong actions, and cases where people ought to know better and are culpable for not doing better at finding out the truth? This still presents a deep challenge insofar as we lack any clear or general criteria for drawing that line. And if we are not sure where to draw that line, then it will sometimes be hard to say where we should be confident in attributing blame and where we should withhold it. (FitzPatrick 2008, 612-613)

Although I have endeavoured to expand on the relevant criteria, it far from removes all such cases of uncertainty.

<sup>xviii</sup> Under certain sentimentalist accounts, most famously that of David Hume, sympathy involves trying to put ourselves in the shoes of another to understand what they are feeling. Michael Frazer, a Humean scholar, puts it nicely when he writes “Mere association with the highly vivid sense of ourselves, for example, can enliven the ideas we relate to ourselves into impressions. Sympathy occurs when the idea of another’s passion is, by this means, enlivened into the experience of an analogous passion” (Frazer 43, 2012). *Passions*, here, is a broad term that scholars of Hume’s time used to refer to affective phenomenon. If we take this kind of definition of sympathy, the very act of considering the relevant variables, most notably the first two (capabilities, environmental influences), involves us trying to understand the experience of the agent in question and so requires a degree of sympathy. So, it follows that sympathy for the agent we are considering is a prerequisite to arrive at a judgment of moral responsibility. This, I think, is another virtue of the theory I am here offering.

Hume also makes the empirical claim, which seems to widely hold true, that our manifestation of sympathy will be higher towards those we identify more easily with (T 2.1.11.6). Frazer again summarizes this nicely when he writes “And the primary prediction of Hume’s account of sympathy is that, the greater degree of similarity between two individuals or their passions, the easier and stronger is the transmission of sentiments between them. Contiguity in space and time, as well as preexisting ties of blood or affection, can also give us a greater sense of closeness to the candidate objects of our sympathy” (Frazer 43, 2012). Of course, these ideas have been expanded upon in modern times, and are encompassed by the wide range of personal biases that have been identified as prevalent in our species. Nonetheless, it is just another reminder that, when we are judging someone we have a strong bias against (perhaps due to political, cultural, or moral disagreements), we must recognize that we will be inclined to find them

to be blameworthy, and so we must make the extra effort to conclude that this blame is the product of the morally relevant criteria, and not our personal bias.

<sup>xix</sup> FitzPatrick makes this point salient when he writes “there are plenty of genuinely hard cases in morality where it would be naive to suppose that everyone who tried could reasonably be expected to come to the correct answer, even on the assumption that there is one” (FitzPatrick 2008, 612), and he gives the examples of deontologist and consequentialist philosophers who are still divided on fundamental questions in ethics. Equally I think we can apply the same acknowledgment of uncertainty to some empirical matters when the evidence is inconclusive.