

Degrees of Responsibility

1. Introduction

The traditional debate about moral responsibility concerns the question of *where* to draw the line between responsible and non-responsible agency: which agents can be praise- and blameworthy for their actions and which cannot be? Young children, people suffering from certain kinds of cognitive, affective, or behavioral disorders, and non-human animals are all agents who are plausibly thought fall below the minimal threshold of responsible agency. These agents altogether lack or only have insufficiently developed versions of the agential capacities that are required to be morally responsible for one's actions. As such, they do not deserve praise or blame for what they do. On the other hand, mature adults are widely thought to be above the threshold.¹ And so, the thought goes, most psychologically and developmentally normal adults *are* deserving of praise or blame for how they conduct themselves.²

There are, of course, more difficult cases. Early adolescents, psychopaths, addicts, and those laboring under great stress are all “close calls” in the sense that it's just not clear as to whether agents of these sorts meet the minimal threshold conditions on responsible agency. These sorts of agents seem in some ways to be morally responsible for their actions. However, they also seem in some other ways to not quite meet the minimal standards on responsible agency.³ But how we regard someone, what attitudes we'll adopt towards them, and what modes of interaction we'll open ourselves up to crucially depends on whether or not they are in fact morally responsible. The question of what separates responsible from non-responsible agency—a question made all the more vivid by the fact how we answer it has significant normative implications for

¹ Because they are *generally* skeptical of anyone's being morally responsible, responsibility skeptics deny that adults typically meet this threshold. For some skeptics (see Derk Pereboom 2001), it's because the bar for morally responsible agency is too high for any *actual* human agents to meet. For other skeptics (see Galen Strawson 1994), it's because the

² Throughout the paper when I speak of agents' actions or conduct, I mean to be as inclusive as possible such that some omissions and even some mental attitudes can count as being part of this class.

³ See Chapters 27 (Gary Watson on psychopathy), 28 (Jeanette Kennett on addiction), and 30 (John Doris and Dominic Murphy on duress in wartime) of this volume for some of the difficulties that (some of) these classes of agents pose for theories of moral responsibility. The question of the moral responsibility of children and young adolescents is not one that has been taken up in any systematic way, though philosophers working on these issues have tended to follow P.F. Strawson (1974) in thinking that children at least are exempt from moral responsibility.

how we should treat the millions of agents who fall into one of these difficult cases—is thus an important one.

Yet even if we are able to ascertain just where the threshold between responsible and non-responsible agency lies, we can't leave it there.⁴ After all, agents who meet the minimum standards on responsible agency are not all the same with respect to what they deserve. This is true, to take one obvious kind of case, when comparing two morally responsible agents, one of whom performs a qualitatively worse action than does the other. If Randall steals \$20 from you and Pearl steals \$2000 from you, then these two agents are deserving of very different responses, since the content of Pearl's morally responsible action is significantly worse.

But cases of this sort surely do not exhaust the scenarios in which we have reason to adjust the degree to which we hold someone responsible. Nor do they really get to the heart of the matter, since it's perfectly sensible to say of both Randall and Pearl that they are each *fully morally responsible* for their actions. Indeed, the explanation of why these agents deserve different things is a function of *how bad* the thing they are morally responsible for doing is and *not* how morally responsible they are for doing it. The real problem that *degrees of responsibility* present for general theories of moral responsibility thus lies in cases in which we hold fixed the action type (and also, perhaps, hold fixed the circumstances in which the action is performed and the foreseeable consequences of the action). For even in such cases, it still seems possible that two morally responsible agents can differ significantly in what they deserve.

To see this, let's first consider Blaire, a friend who's suffering from depression but who has promised to help you move.⁵ Since a general lack of motivation—even in the face of a judgment that the act in question really is what you should do—is symptomatic of clinical depression, it's natural to think Blaire has done something quite impressive simply by keeping her promise. Indeed, Blaire seems to be *more* praiseworthy for mustering what little energy she has in an effort to help you than another friend, Aida, would be for doing the same thing, since helping you would be no trouble at all for Aida. Now in this case, the fact that it's easy for Aida to help you doesn't mean that Aida is praiseworthy only to a small degree for her action. She need not be

⁴ In addition to the difficult cases I've already mentioned, the notoriously thorny debates between compatibilists and incompatibilists concerning the significance of causal determinism provide a further reason to think that *this* will prove to be no small thing.

⁵ These cases are similar to ones I developed with Philip Swenson (Coates and Swenson 2013), but unlike those, I here consider the implications for praiseworthy as well as blameworthy agency. For thoughtful engagement with those cases, see Nelkin (2016) and Tierney (n.d.).

downgraded in the degree to which she's praiseworthy simply because she finds no difficulty in keeping her promises. Instead, facts about Blaire—in particular, that keeping the promise took a lot of effort on Blaire's part and that although her effort wasn't likely to succeed, she overcame the odds—lead us to regard her as being an exceptional case. It's in light of these facts that we come to regard Blaire as an agent who really does deserve more praise than is ordinarily made fitting by praiseworthy behavior of this sort.

Alternatively, consider a teenager, Emma, working her first job. Frustrated by the tedium of her job, Emma finds herself sorely tempted to make some easy money. It's plausible that Emma would be *less* blameworthy for stealing customers' credit card numbers than Mila, an adult with lots of customer service experience but who also finds herself deeply bored with her job, would be for engaging in the exact same kind of fraud. We mitigate blame in Emma's case not because we're skeptical that she's morally responsible but because she's "got a lot of growing up to do." Only then, can we reasonably hold her to the normative standard in question with the *same* alacrity and the *same* intensity that's found in expressions of blame directed towards Emma's adult counterparts (like Mila) for their type-identical actions.⁶ In other words, the fact that she's still developing the very capacities that underwrite her status as a morally responsible agent doesn't get Emma off the hook completely. But it does seem to rationalize importantly different responses in otherwise similar cases of wrongdoing.

On a plausible theory of act-individuation, Blaire doesn't do anything better than Aida does, and Mila doesn't do anything worse than Emma does.⁷ The fact that these agents deserve different responses doesn't obviously depend the relative badness of their actions (as it does in the case of Randall and Pearl). It seems, then, that agents can be *more or less* responsible for their

⁶ In "Freedom and Resentment" (1974), P. F. Strawson considers pleas of the sort that we (or third parties) avail ourselves of when confronted with blame that we don't regard as altogether fitting. Strawson's own treatment of these pleas—excusing and exempting conditions as they've come to be known (see, Gary Watson 1987)—leaves out the complexities that arise when thinking about diminished or enhanced moral responsibility. This is an unfortunate omission, since the thought that someone has "a lot of growing up to do" is a genuine basis for somewhat altering one's attitudes towards agents like Emma even though it doesn't fit nicely into either of Strawson's categories. A full account of these pleas must supply not only explanations of excusing and exempting conditions but also of mitigating and enhancing conditions as well.

⁷ One should not, I think, offer such a fine-grained account of act-individuation so as to include what an agent must overcome in order to perform the act in question into one's account of what action has been performed. If one rejects this advice, then precisely the same issues that I'll consider in this paper arise in the context of assessing what exactly it was that the agent did rather than in assessing the degree to which an agent is morally responsible. There is no escape.

actions—that they can be morally responsible to different degrees. It’s of course important to know the minimal threshold for morally responsible agency, but it is, I submit, just as important to know what determines the degree to which agents are morally responsible for their actions.

What this means is rather straightforward. Any adequate theory of moral responsibility must account for the fact that agents can be more or less morally responsible for their actions. To do this, I think that we have to answer two distinct questions. The first of these questions concerns what exactly it means to say that one agent is more (or less) responsible for her actions than some other agent who has also perform an action of that type *or* than she would be for performing the same actions in some relevantly similar set of alternative circumstances. Put differently, it’s just the question of what we mean when we say that Blaire is *more* praiseworthy for keeping her promise than Aida is or that Mila is *more* blameworthy for committing fraud than Emma is. The second of these questions concerns the conditions that can affect or alter the degree to which an agent is more (or less) morally responsible for what she’s done. This is ultimately the question of *what explains* why Emma is less blameworthy than Mila for committing credit card fraud. This won’t be all there is to say about degrees of responsibility, but it is the bare minimum that adequate theories of moral responsibility owe us.

2. Being More (or Less) Morally Responsible

Before going any farther, we need to address the question of what exactly it is to be morally responsible for some action. Only then can we hope to understand what exact property it is that we’re ascribing to someone when we say of her that she is more (or less) morally responsible for her action.

2.1. The Desert-entailing Conception of Moral Responsibility

There is no one thing people are referring to when they discuss moral responsibility. Perhaps the most common conception of moral responsibility, however, is the desert-entailing conception of moral responsibility.⁸ Concerning this conception of moral responsibility—a conception that traces back to (at least) the work of Spinoza—Derk Pereboom offers the following statement:

For an agent to be *morally responsible for an action* is for this action to belong to the agent in such a way that she would deserve blame if the action were morally wrong, and she would

⁸ Here I’m specifically focusing on what Gary Watson (1996) has called the “accountability face” of moral responsibility.

deserve credit or perhaps praise if it were morally exemplary. The desert at issue here is basic in the sense that the agent, to be morally responsible, would deserve the blame or credit just by virtue of having performed the action, and not, for example, by way of consequentialist considerations (Pereboom 2001, xx).⁹

So what separates the desert-entailing conception of moral responsibility from other conceptions is that on this view an agent is not morally responsible simply because it's good or beneficial to praise or blame her for what she does; rather she is morally responsible because she *deserves* to be praised or blame, even if praise or blame aren't themselves good or beneficial in the circumstances.

Of course, Pereboom's characterization of what it is to be morally responsible is not without its critics.¹⁰ But for present purposes, we can set these worries aside, since Pereboom has quite clearly adequately characterized *something* that many of us care about, even if in so doing, he's characterized the phenomenon too narrowly.

2.2. Two Ways to be More (or Less) Morally Responsible

The desert-entailing conception of moral responsibility fits nicely with the idea that agents can be more or less morally responsible for their actions. A simple extension of this conception of moral responsibility might hold that agents are *more* morally responsible for their actions when they are more deserving of praise or blame and that they're *less* morally responsible for their actions when they're less deserving of praise or blame. This seems right, as far as it goes, but this simple extension is ambiguous in a significant way.

To get a better handle on precisely the way that this sort of talk is ambiguous, let's return to the case of Emma and Mila. Recall that Emma and Mila each commit credit card fraud, but that because Mila is an adult and Emma is a teenager, it seems that Mila is in some sense more morally responsible for her crime than Emma is for hers. But how is this connected to what they each deserve? One possibility is that in attributing to Emma and Mila different degrees of responsibility, we're saying that Mila is *more deserving of blame* for what she's done than Emma is. Another possibility is that we're saying that Mila is *deserving of more blame* for her action than

⁹ For further elucidation of these ideas, see Michael McKenna (2012) and Derk Pereboom (2014). Of note is that in his most recent restatement of this view, Pereboom adds that not only is basic desert something that obtains independently of consequentialist considerations, it also obtains independently of contractualist ones as well.

¹⁰ See James Lenman (2006) and Manuel Vargas (2013) for two examples.

Emma is for hers. And despite the apparent similarities in these two possibilities, the properties of *being more deserving* and of *being deserving of more* are importantly different.¹¹

To say that someone is more deserving of something suggests that it is intrinsically better to give that person whatever it is that they deserve. What this means is that solely from the point of view of desert it's better to give Mila what it is that she deserves than it is to give Emma what it is that she deserves. Alternatively, we could say that in such a case there are weightier desert-based reasons to blame Mila than there are to blame Emma.

On the other hand, to say that someone is deserving of more of something (be it praise or blame) is to say something about *how much* praise or blame is deserved. Now unlike monetary rewards or prison sentences, deserved praise and blame resist easy quantification. But this is not to say that we can't say something meaningful about deserving differing amounts of praise or blame. For example, the blame we bestow on someone for acting poorly can vary in intensity or duration. If we regard someone as being deserving of a great deal of blame, then we might be more reticent to extinguish our blaming attitudes. We might also experience the blaming attitudes more intensely. And this, in turn, might more powerfully motivate us to express these attitudes. What this means is that if an agent is deserving of *more blame* then she is deserving of more intense blame, or of being targeted with blame for longer, or having more direct expressions of blame being leveled at her. On the other hand, if an agent is deserving of *less blame* then she is deserving of less intense blame, or blame that doesn't last as long, or blame that's less intrinsically motivating.

The degree to which an agent is morally responsible is therefore fixed by these two independent variables: how weighty the desert-based reasons there are to praise or blame and how much praise or blame the agent deserves. These two properties can combine in ways that explain what exactly it is that people deserve in light of their actions.¹² Recall Blaire, who helped you even though it was quite difficult for her to overcome temptation that had its source in her depression. Blaire is very praiseworthy in the sense that you should be more effusive with your praise (i.e., she is deserving of more praise), but she is also very praiseworthy in the sense that the desert-based reasons to praise her in the way that she deserves are quite weighty. To see this, suppose you're talking to a mutual friend about your move. If you fail to praise Aida, who helped you, but for

¹¹ For more on this distinction, see Robert J Hartman (2017) and D. Justin Coates (forthcoming).

¹² I expand upon this point with the help of desert graphs—an apparatus for understanding desert developed by Shelly Kagan (2012)—in greater detail in Coates (forthcoming).

whom helping was very easy, then you've missed a chance to give credit where credit's due. This is bad, given the fact that your conversation was about your move. Yet this failure seems amplified if you were to fail to praise Blaire. What explains this, I think, is that the desert-based reasons to praise Blaire are weightier than the desert-based reasons to praise Aida. As such, failure to act on those reasons absent weighty countervailing reasons is correspondingly greater in Blaire's case than it is in Aida's.

Having these two variables to play with allows us to conceptually map the different degrees to which agents might be morally responsible for their actions. That is, this framework (at least when it's suitably worked out) helps us to clearly articulate what exactly we're attributing to an agent when we claim that she is more (or less) responsible. However, it doesn't yet tell us the conditions under which agents are more (or less) morally responsible for their actions. I turn to this question now.

3. When are Agents More (or Less) Responsible?

What, then, are the conditions under which agents are more (or less) morally responsible for their actions? In particular, how do differences in the intrinsic properties of an agent *or* in the circumstances that agents find themselves affect the degree to which they are morally responsible?

3.1. Three Proposals

There are three common ways of explaining why one agent is more (or less) morally responsible than another (or than she herself would be were circumstances different) for performing type-identical actions. At first glance, these explanations appear to be quite distinct from one another. As we'll see however, they actually rely on a shared assumption about the connection between how hard it is to comply with legitimate normative standards and the degree to which an agent is morally responsible for her compliance (or non-compliance, as the case might be).

The first proposed explanation appeals to the quality of agents' wills. On this proposal, agents with better (or worse) qualities of will can differ in the degree to which they are morally responsible for performing some action. If this is right, then we'll need an account of what "quality of will" comes to exactly. One such account appeals exclusively to the intrinsic desires that are manifested in an agent's action. Another way of understanding quality of will emphasizes the quality of the reasons that move an agent to action. On these accounts, the degree to which an agent is morally responsible depends, *inter alia*, on the degree to which one's action manifests

intrinsic desires of the relevant sort *or* on the degree to which one's action manifests better (or worse) reasons. Or both!¹³

A second proposal appeals to an agent's epistemic position with respect to the moral status of her action. Many agents are ignorant of the moral status of what they do. Some of these agents are *culpably* ignorant. Ignorance of this sort plausibly doesn't excuse, but it does seem that an agent who knowingly does evil would be more responsible for what they've done than someone who does evil due to culpable ignorance (compare, e.g., the reactions one has to Iago's villainy to Othello's tragic rage). So perhaps culpable ignorance mitigates responsibility.¹⁴ Other agents, however, are apparently *non-culpably* ignorant. What should we say about these individuals' degree of moral responsibility? Does non-culpable ignorance excuse, as some have suggested? Or does it merely mitigate?

A third proposal posits that the degree to which an agent is morally responsible for her action depends at least in part on the degree to which that agent has control over what she does. On one prominent theory, having control over one's action amounts to the capacity to *guide* one's behavior in light of what reasons one recognizes oneself to have. If this right, then improvements to or deficiencies in one's ability to recognize reasons or to translate those reasons into action will affect the degree to which you control that action. However, if control depends more broadly on the quality of one's opportunity, where this is determined not only by the agent's own rational capacities but also by the circumstances she finds herself in, then the degree to which she is morally responsible depends on a more holistic assessment of her situation. In either case, differences in the degree to which agents have control over their actions do seem relevant to the degree to which those agents are morally responsible for what they do.

3.2. Unifying the Proposals: The Significance of Difficulty

Rather than viewing these proposals as being radically at odds with one another, I'm inclined to think that each provides us with a *partial* account of the conditions under which agents are more (or less) morally responsible for their actions. This might be taken to suggest that there's no one thing underlying the conditions that affect the degree to which an agent is morally responsible.

¹³ As we'll see these proposal are, for the most part, consistent with one another. So it might be that only when they're each given their due that we'll have a suitably fleshed out explanation of why it is that some agents is more (or less) morally responsible for their actions.

¹⁴ Or alternatively, perhaps knowingly doing evil enhances the degree to which an agent is morally responsible.

However, I don't think that's right. The explanations for why agents are more (or less) responsible that these proposals offer are more unified than it appears at first glance.

The common thread in each of these proposals is that they each get their appeal in light of the following thought: how responsible you are for what you've done depends on how difficult it was for you to comply with legitimate normative expectations. To wit: if it's easier for me to do the right thing, then I show worse quality of will for doing the wrong thing than I would if it had been hard to comply with legitimate norms. That's what explains why I'm more blameworthy in the former case than I would've been in the latter. Or, if I'm ignorant of some important moral fact, but it would have been very hard for me to have come to the truth, then I am less blameworthy for acting wrongly than if I could have very easily arrived at the moral facts but didn't concern myself to do so. Similarly, if circumstances make it harder for me to control myself in the way that's expected of me, then I seem to be more praiseworthy for doing the right thing precisely because in so doing, I overcame a significant difficulty.

This straightforward explanation of the connection between degrees of responsibility and praise- and blameworthiness seems to rely on what Holly Smith (1991) has called the *battle citation* model of moral praise and on a related model of moral blame. On the battle citation model, "an agent is creditable for performing a right action if and only if a morally good desire won a hard battle in the war against temptation to perform the wrong act," (Smith 1991, 281-82). Since this model posits that what it is in virtue of which an agent deserves moral credit or praise is that he or she has won a *hard* or *difficult* battle, it seems plausible on the very same grounds that *how* deserving of praise the agent is depends (at least in part) on just how difficult it was for her to win the battle. If beating the temptation in question constituted defeating "overwhelming odds" then it stands to reason that the agent accrues more credit- or praiseworthiness than if it was merely a hard-fought-but-still-manageable battle against some moderately recalcitrant desire. After all, mustering the effort required to overcome an extremely powerful temptation reveals that you are deeply committed to the good. In other words, overcoming a difficult temptation *reveals* that you have a good will. Difficulty thus seems to *enhance* praiseworthiness.

By parity of reasoning, how difficult it is to comply with normative standards is also relevant to how blameworthy an agent is. Suppose that you fail to comply with legitimate standards in a case in which it would've been quite difficult for you to have done the right thing. Here, difficulty seems to *mitigate* rather than enhance blameworthiness precisely because when it's very hard to comply with legitimate norms—norms that exist at the edge of what we can reasonably expect of one another—that agents seem less blameworthy when they fail to comply with them. One

explanation of this is that if it's easy for me to comply with a normative standard, then failure to do so will typically reveal that I do not care about meeting that standard. On the other hand, failure to comply with difficult but legitimate demands does not necessarily show me to be indifferent about meeting the standard in question.

Smith's *battle citation* model of moral praiseworthiness and the corresponding model of moral blameworthiness provide us with a conceptual basis for the slogan that *difficulty determines degrees*. In other words, it provides us with a framework for understanding how the difficulty of complying with legitimate norms regularly explains our judgment of others as being more (or less) morally responsible for their conduct.¹⁵ The three evidently disparate proposals that I started with all seem to agree on this basic idea. Indeed, they might even be *based in* this basic idea. As such, to properly evaluate these proposals, we'll need to see exactly how difficulty affects the quality of an agent's will, how culpable she might be for ignorance, or the degree to which she can be said to control her action.

3.3. Quality of Will

One explanation of why two agents who perform type-identical actions could be morally responsible for doing so to different degrees is that one agent's action expresses more good or ill will than did the other's. This seems to follow from the widely held view that praise- and blameworthiness track, *inter alia*, quality of will.¹⁶ After all, if it's the case that I must display good or ill will in order to be praise- or blameworthy for my action, then *how much* good or ill will I display will render me more (or less) deserving of praise or blame for that action. And since what it is to be morally responsible for some action is to be deserving of praise or blame for that action, then it seems like the degree to which I display good or ill will to in my actions is directly relevant to the degree to which I am morally responsible for those actions.

Nomy Arpaly and Tim Schroeder (2014) have recently proposed an account along these lines. On their view, (complete) good will is "an intrinsic desire for the right or the good, correctly conceptualized," and (complete) ill will is "an intrinsic desire for the wrong or the bad, correctly

¹⁵ This does not mean, of course, that it's *only* difficulty that determines the degree to which an agent is morally responsible for what she does. But appeals to difficulty do a lot of explanatory work in this domain. So for that reason I focus on it in this chapter.

¹⁶ In the contemporary debate, most trace this thesis to P. F. Strawson (1974), but the general thought goes back at least to Adam Smith (1759/1976).

conceptualized,” (Arpaly and Schroeder 2014, 162).¹⁷ From this they build scalar theories of praise- and blameworthiness. An agent is praiseworthy for their action, on their view, to the extent that that action *manifests* good will (as they’ve defined it in terms of intrinsic desires for the right or the good). On the other hand, an agent is blameworthy for some action to the extent that that action *manifests* ill will (again, as they’ve defined “ill will”).¹⁸

There’s something to this thought. Agents who more fully manifest intrinsic desires for the good do seem more praiseworthy for what they’ve done than agents who only partially manifest desires for the good. And agents who more fully manifest intrinsic desires for the bad seem correspondingly more blameworthy than agents who manifest a mixture of desires or who appear to be reluctant in going along with something bad. This might not provide a general explanation of the conditions under which agents are more (or less) morally responsible, but it does appear to have explanatory power in a circumscribed range of cases.

Yet despite its initial plausibility, Dana Kay Nelkin (2016) has raised an important objection to accounts like Arpaly and Schroeder’s that take degrees of responsibility to be explained by more (or less) good or ill will. Nelkin invites us to:

Consider a case in which a person is blamelessly drunk at a party. In this state, she gleefully and loudly shares a scandalous secret, told to her in confidentiality about another guest. Suppose that this person would never have done this if she had been sober, and she values keeping promises and respecting others’ privacy. Still, she harbors an intrinsic desire for the other guest to suffer embarrassment [i.e., she has ill will], and it is this desire on which she acts. This seems a case in which the action manifests significant ill will. She also has good will, which normally would overcome the ill. But she is drunk and it does not, (Nelkin 2016, 363-64).

From this case Nelkin goes on to argue for a more general point, which is simply that your action can manifest ill will even if you’re *excused*, such that you’re not morally responsible *at all* for that action.¹⁹ As a result, it *cannot* be the quality of the blameless drunk’s will that explains the degree

¹⁷ In this same passage, Arpaly and Schroeder also define partial good will and partial ill will, but for our purposes, that level of detail is unnecessary.

¹⁸ Arpaly and Schroeder (2014): 170.

¹⁹ Nelkin frames the point slightly differently, since what she’s after is an account of difficulty of the sort that might affect the degree to which an agent is praise- or blameworthy for her action. But in the case of the blameless drunk, she’s

to which she is responsible, since in this case, she has both a poor quality of will and is not morally responsible at all.

This is a powerful challenge to Arpaly and Schroeder, and I think Nelkin is right in her assessment of the case as she describes it. But I also think she's wrong (or at least, too quick) to draw the general conclusion that she does. That is, Nelkin's right to conclude that differences in good and ill will *as described by Arpaly and Schroeder* are insufficient to account for the degree to which agents are morally responsible in cases of this sort. But this, I submit, should lead us to doubt that Arpaly and Schroeder's account of good and ill will *rather* than the more general claim that the quality of an agent's will affect the degree to which she is morally responsible for the actions that manifest that will. After all, on one natural interpretation of case that Nelkin describes, the blameless drunk doesn't actually display ill will at all. It's true, of course, that she's mean-spirited and maybe even malicious. But young children can be mean (and maybe even malicious) without manifesting ill will in the sense at stake.²⁰ So manifesting good and ill will in one's actions thus requires some further connection between the intrinsic desires for the good or the bad and the mechanisms that translate these desires to action. It's only when these mechanisms have some further property—e.g., that they are suitably reasons-responsive—that the manifestation of one of the relevant intrinsic desires in an agent's action counts as manifesting good or ill will. It's therefore possible to maintain the initial thought that differences in the degree to which agents are morally responsible are explicable by appeal to differences in the quality of their wills.

But to clearly see how quality of will might still matter, let's reconsider the cases of Blaire and Aida. Unlike the blameless drunk or a young child, these two agents plausibly meet the minimum threshold on morally responsibility. This means that the mechanisms issuing in their actions suitably instantiate whatever property (e.g., reasons-responsiveness) that's necessary for the manifestation of one of the relevant intrinsic desires to count as good or ill will. In these cases,

less responsible because it would've been difficult for her to comply with legitimate expectations even though *that* difficulty doesn't seem to affect the degree to which she displays ill will in this case.

²⁰ This thought follows from a roughly "Strawsonian" conception of quality of will. Although P. F. Strawson (1974) never explicitly defines what he takes good or ill will to be, it's clear that he regards young children as being unable to display good or ill will of the sort that the reactive attitudes respond to (they are *exempt* from moral responsibility, after all). But no one with young children would deny that they could be mean—after being very gently scolded for continually ignoring a day care employee, my three-year-old daughter recently informed my wife that she was, "just a little mean sometimes, mommy."

Nelkin's worries won't obviously apply. Both Blaire and Aida are connected to their actions in the "right" way; they both possess the kind of control that's necessary for moral responsibility. Yet Blaire seems more responsible for keeping her promise than Aida does. The explanation for this is just that Blaire has manifested an even better quality of will towards you than Aida has (or perhaps any agent could reasonably be expected to), since given her severe depression, the alternatives were more (rationally) tempting for Blaire than they were for Aida.²¹ More (rationally) tempting, and yet, Blaire overcomes that temptation. As Smith helpfully puts it, Blaire *won* the battle, and this triumph speaks especially well of her quality of will.²²

3.4. The Epistemic Condition

Whether I am morally responsible at all for performing an action depends not only on whether I meet the control condition on moral responsibility, but also on whether I meet some further epistemic condition on responsibility. It stands to reason, then, that differences in agents' epistemic statuses can affect the degree to which those agents are morally responsible for their actions.

²¹ Along these lines, Hannah Tierney (n.d.) has recently argued that the degree to which an agent is morally responsible depends in part on "the quality of the reasons" for which she acts. Tierney takes this to show that the degree to which an agent is morally responsible depends on more than how difficult it is (or would be) for an agent to comply with legitimate normative standards. In some cases, Tierney tells us, it depends the "quality of [an agent's] reasons," where this concerns the moral, epistemic, or "popular" status of those reasons. Yet although I think that Tierney's right to think that degrees of responsibility depend (at least in part) on the quality of an agent's reasons, I'm not sure that the quality of reasons can be understood without appeal to *some* conception of difficulty. For example, on Tierney's interpretation of cases like that of Blaire and Aida, part of what explains why Blaire is so praiseworthy for keeping her promise is that in light of her depression she could have, in some sense, been reasonable to prioritizing her mental health over keeping her promise. But to do so, she would have had to make the difficult decision of prioritizing herself over her promise to a friend. Instead, she opted for the difficult decision of prioritizing her promise to a friend over her short-term health. The difficulty here isn't a causal force, like the difficulty of overcoming a compulsive desire. Rather it's the kind of difficulty that arises in the context of ordering one's own priorities. Blaire, we're imagining, ordered them correctly in this case, so she's deserving of more praise than someone like Aida, who faced no such burden.

²² This isn't to say that there are *no* control-relevant differences between the degree to which Blaire is morally responsible and the degree to which Aida is morally responsible. I'll argue in §3.5 that there are. But whatever one thinks about the connection between control and the degree to which an agent is morally responsible, once we correct for the problem in Arpaly and Schroeder's view, there's room left for the quality of an agent's will to matter for the degree to which she is morally responsible.

One factor that affects whether an agent meets the epistemic condition on moral responsibility is her *attention to* and *awareness of* morally relevant facts.²³ These are the building block on epistemic and moral conscientiousness, which serves to protect agents from the epistemic and moral mistakes that we're readily prone to make. Yet agents can attend to moral considerations to varying degrees. And unfortunately, even when an agent is especially attentive to morally relevant considerations she might nevertheless fail to "put it together" in a way that leads her to recognize the moral truth. In such a case, she seems (at least) *less* blameworthy for acting in a way that's contrary to how she ought to act than she would be had she simply ignored those considerations. Conscientiousness of this sort seems to *mitigate* blameworthiness because genuinely conscientious agents display good will even when they get it wrong; and this tells against them being deserving of blame to the same degree that non-conscientious wrongdoers are deserving of blame.²⁴

But if conscientiousness mitigates in ordinary cases, then maybe in exceptional cases—cases in which even maximal conscientiousness wouldn't lead the agent to the right conclusion—fully exculpates. Gideon Rosen, to name one example, takes up this thought, claiming that "[g]iven the intellectual and cultural resources available to a second millennium Hittite lord, it would have taken a moral genius to see through the wrongness of chattel slavery," (Rosen 2003, 66). Such an agent, Rosen concludes, is blameless: "in my view it makes no sense to hold this injustice against the perpetrator when it would have taken a miracle of moral vision for him to have seen the moral case for acting differently," (Rosen 2003, 66). Here Rosen seems to be thinking that because no amount of careful reflection (and indeed, nothing short of a miracle) would have led the Hittite lord to sensible views about slavery, he's excused for his wrongdoing. This is startling, since Hittite lords knowingly and intentionally treated their slaves with brutal cruelty. But if Rosen's exculpatory conclusion simply follows from the thought that some degree of conscientiousness can mitigate, then maybe we should regard that idea to be dubious.

In fact, however, there's much to be said in response to Rosen's claims that moral ignorance is an excusing (rather than a merely *mitigating*) condition. First, it's doubtful that most Hittite lords were even minimally conscientious with respect to the question of chattel slavery. As

²³ This thought calls to mind Neil Levy's (2013) "consciousness thesis," which holds that an agent is morally responsible for her action only if she is conscious of the facts that give that action its moral significance."

²⁴ This part of the explanation of why epistemic conscientiousness mitigates fault is related to Tierney's quality of reasons thesis, which concerns not only the moral reasons for which agents act but also the epistemic reasons that they regard as action-guiding.

such, if they really are blameless (or weaker: if they are significantly less blameworthy than contemporary slaveholders would be²⁵), it can't be because they're working hard to arrive at the truth, if only to an imperfect degree. More generally, however, it's not clear why we should be so quick to reject the Hittite lords' responsibility in cases like this. Just because it's difficult—nearly impossible!—for them to know that slavery is wrong doesn't mean they cannot be morally responsible for that action, at least to some degree. Difficulty might determine *degree*, but not in the limiting case. That is, difficulty isn't *per se* exculpatory.

Alex Guerrero (2017) develops a version of this thought by distinguishing between three types of difficulty: *skill-related difficulty*, *effort-related difficulty*, and *difficulty in trying*. On this taxonomy, we can say that it might be difficult for the ancient slave owner to have gotten it right due to some skill-related difficulty that's inherent to solving the moral problem at issue. Here we're imagining that the figuring out that slavery is wrong is somehow akin to figuring out a complicated mathematical problem. Alternatively, it might be difficult for the ancient slave owner to have gotten it right due to some effort-related difficulty that arises in the psychology of the slave owner. Maybe in this case, it would've just been very hard for the ancient slave owner to give the question the attention required in order to figure out the answer, even though, were he to have done so, arriving at the correct answer would have been easy. A third possibility is it's difficult for the slave owner to get it right because it would have never even occurred to him to try to figure out the moral status of slavery.²⁶

However, Guerrero claims that whether one thinks that difficulties of these kinds mitigate ultimately depends on the theory of moral reasoning one accepts. If you accept a theory according to which moral reasoning is relatively straightforward, then it's doubtful that skill-related difficulties are really present in the case at hand. If, on the other hand, you accept a theory according to which successful moral reasoning typically requires a great deal of cognitive sophistication and expertise, then, Guerrero claims, it's doubtful that skill-related difficulties would mitigate. On this point he does admit that it's plausible that coming to believe novel moral truths might have skill-related difficulties, but he counters this with the equally plausible claim that this can't what explains why Hittite lords fail to believe what's true, since they don't even try

²⁵ See Erik Loomis (2015) for an up-to-date account of the role of slave labor in contemporary supply chains. Unlike ancient slave owners, contemporary ones inhabit a world where egalitarian values are widespread (if not in practice, at least in principle). Accordingly, they can't pretend not to know that the people they're enslaving matter and should be treated accordingly.

²⁶ Guerrero rightly notes that it also could've been difficult in all three of these ways.

to figure things out. Yet what Guerrero says here is consistent with the view that failing to try something that's hard to do mitigates *less* than trying and failing would. As a result, if arriving at moral truths does involve skill related difficulties, then Guerrero hasn't yet given us a reason to deny that it mitigates (at least to some small degree).

Later, Guerrero claims that to the degree the Hittite lord faces effort-related difficulty or difficulty in trying, those forms of difficulty only arise as a "result of the agent's moral attitudes/character," (213) and so don't mitigate the degree to which the agent is morally responsible.²⁷ But if Rosen goes too far in denying that the Hittite lord is blameworthy for owning slaves, Guerrero overcorrects on this point. The fact that some effort-related difficulty or difficulty in trying results from an agent's moral attitudes or character will fail to mitigate only if that agent is *fully* morally responsible for those attitudes or character. But plausibly, if an agent is less than fully responsible for those attitudes or character, then difficulties that arise from them would mitigate the degree to which she is blameworthy for failing to arrive at the right conclusion.

But this is just the position that the Hittite lord is in. After all, if he is *fully* blameworthy for his attitudes, then we'd have reason to treat him in precisely the same way that we have reason to treat a contemporary who argues for the moral permissibility of chattel slavery. Yet it's not clear that we do have reason to treat these two agents in exactly the same way. If I were to meet (roughly) a peer who apologized for slavery, I'd be inclined to simply blame him. On the other hand, if I met a time-traveling Hittite lord I'd feel a strong pull towards at least *trying* to educate him.²⁸ This, of course, doesn't mean that he's not blameworthy to some degree, but it does suggest that the desert-based reasons to blame the Hittite lord are weaker than the desert-based reasons to blame the contemporary defender of slavery.

There is a final reason to doubt that the Hittite lord's blameworthiness isn't to some degree diminished. Suppose that the Hittite lord did engage in the kind of reflection that epistemically and morally conscientious agents are prone to, and that after a lot of careful thought, he came to believe that slavery was immoral. If the difficulty of this undertaking didn't mitigate

²⁷ Guerrero admits that this claim rests on a controversial conception of agency and responsibility. I agree with this, though I suspect Guerrero and I come down on different sides of the controversy.

²⁸ To make this point a bit differently, consider what would be required for you to give these two agents the benefit of the doubt. Plausibly, each of them would exhaust your generosity fairly quickly, but it also seems plausible that you'd have reason to withdraw your generosity even more quickly in the case of the contemporary slavery apologist. The explanation for this: there are weightier desert-based reasons to blame him than there are to blame the Hittite lord. But as I argued in §2.2, this just is what it is for the former agent to be more blameworthy than the latter.

blameworthiness in the case in which he *failed* to reflect, it's doubtful that it could *enhance* praiseworthiness in the case in which the Hittite lord did reflect and after a grueling bout of self-assessment came to believe that he was doing something quite horrific. Yet it seems clear to me that in this case the Hittite lord is more praiseworthy for getting it right (and acting on that belief) than I am for doing the same. And the explanation for this is just again, that difficulty determines degrees. So although I heartedly agree with Guerrero that Rosen too easily lets ignorant wrongdoers off the hook, I also think that how hard it can be to get the facts right (in any of the senses of difficulty Guerrero articulates) is directly linked to the degree to which agents are responsible.

3.5. Control

Agents are morally responsible only if they control their actions.²⁹ But what does control of the relevant sort consist in? Minimally, it requires that an agent is suitably reasons-responsive. After all, if an agent is incapacitated in a way that renders the mechanisms that issue in her actions non-receptive to reasons or non-reactive to those reasons, then she is not normatively (or rationally) competent in the way that is the basis for well-functioning moral agency.³⁰ But it can be more or less difficult to recognize reasons for (or against) some course of action. So too, it can be more or less difficult to translate that recognition into action. And the degree to which an agent controls her action plausibly depends on how easy or hard it is for the mechanisms issuing in her actions to issue in a right action. So far, so good. But if we want we want to secure this explanation, then we'll need to say more about what it is in virtue of which it might be more or less difficult for the mechanisms to issue in actions that are suitably responsive to sufficient reasons for action.

In recent work with Philip Swenson (2013), I've argued that we can do just this by building on extant accounts of control. We start with an influential account of *guidance control* due to John Martin Fischer and Mark Ravizza (1998). According to this account, agents possess the kind of control required for moral responsibility only if the mechanisms that issue in their actions are "moderately reasons-responsive." For Fischer and Ravizza, this amounts to the claim that:

²⁹ The notion of control at issue here is connected to freedom of the sort that is putatively required for moral responsibility.

³⁰ Here I'm following John Martin Fischer and Mark Ravizza (1998) by talking about the control-grounding properties of the mechanisms that issue in an agent's action rather than the agent herself. If this worries you, then I think almost everything I say can be translated into an idiom that appeals only to the control-grounding properties of agents themselves without remainder.

A mechanism of type *K* is moderately responsive to reason to the extent that holding fixed the operation of a *K*-type mechanism, the agent would *recognize* reasons (some of which are moral) in such a way as to give rise to an understandable pattern (from the viewpoint of a third party who understands the agent's values and beliefs), and would *react* to at least one sufficient reason to do otherwise (in some possible scenario). That is, a mechanism is moderately responsive to reason insofar as it is "regularly" receptive to reasons (some of which are moral), and at least weakly reactive to reasons (Fischer and Ravizza 1998, 243-244).

Of particular note here is the idea that a mechanism is weakly reactive to reasons just in case that mechanism would issue in the action for which there is sufficient reason in at least one possible world.³¹ This seems plausible as a bare minimum. After all, the mechanisms issuing in my actions—mechanisms that include practical reasoning, habit, instinct, etc.—can't count as being responsive if there's no possible world in which, given sufficient reason to do otherwise they do so. If they failed to have this property, then we would evidently lack the very capacity that makes us rational agents.

Since Fischer and Ravizza's analysis works well as an account of the bare minimum that's required for morally responsible agency. Swenson and I extend it in the following key way. If one is weakly reactive to reasons in virtue of what happens in one possible world, the *degree* to which an agent is weakly reactive to reasons depends on the *comparative similarity* of the possible world in which the mechanism reacts to the sufficient reason to do otherwise. If that world is *closer*, the thought goes, then there is more overlap in what it and the actual world share. So if we can hold *more* things fixed between these two worlds, then it's *easier* in some meaningful sense for the agent's mechanism to issue in the right action. And if it's easier to get it right, then an agent seems more blameworthy for getting it wrong.

We motivate this by considering a case not unlike that of Blaire and Aida. However, in our original case, the two agents, Marcia and Thomas, *fail* to keep their promise to help you. Marcia fails because she is suffering from severe but non-debilitating depression. And even though

³¹ Swenson and I also discuss reasons-receptivity and offer a model of how that might vary in a way that explains variances in the degree to which agents are morally responsible. For our current purposes, however, I think it's fine focus exclusively on extending Fischer and Ravizza's account of weak reasons-reactivity, since what we say about receptivity is, to a large degree, isomorphic to what we say about reactivity.

it would've been as easy for Thomas to keep his promise, he fails simply because at the time he needed to leave to go help you, he was moved to sit around at home and watch reruns instead. In light of these cases, Swenson and I claim:

In our view, to say that it is “more difficult” or that it is “harder” for Marcia to keep her promise is to say that in the relevant sense, the world in which she does so is less accessible from the actual world. Of course, what makes worlds more or less accessible is a matter of comparative similarity . . . And we reductively analyze the notion of “difficulty” in terms of comparative similarity, (Coates and Swenson 2013, 638–39).

It's more difficult for agents like Marcia to respond to good reasons because the worlds in which they do so are more distant. They therefore have less control over their actions, and are, accordingly, less responsible for those actions. Here again we see difficulty determining degrees. Here it does so because how easy or hard it is for these agents to comply with legitimate normative standards affects the degree to which they possess the kind of control that's necessary for morally responsible agency.

The account that Swenson and I offer seems to get it right in a wide range of cases in which, intuitively, agents are more (or less) morally responsible for their actions. Despite that, it can't serve as an explanation of why agents have differential degrees of control. Nelkin (2016), in particular, has pointed out two important problems for our view. The first is that it's generally a mistake to appeal to modal properties to explain things like difficulty. More plausibly, the modal properties merely co-vary with difficulty because, e.g., the fact that it's hard for a depressed agent to be moved to act on the reasons she judges herself to have is what *grounds* the modal properties in question (viz., that it's a *distant* world in which given sufficient reasons to do otherwise, the mechanism in question issues in an action guided by such reasons). In other words, it's the difficulty of being moved by good reasons that explains the modal properties, not the other way around. This point is almost surely correct. We should have been content to identify the modal differences that track difficulty of the sort that's relevant to how well an agent controls her action, as even this would provide us with what we would need to figure out when to mitigate or enhance the degree to which we should hold someone responsible.

The second objection Nelkin has to the view is that it's overly narrow. Swenson and I explain difficulty of the sort that's relevant to control wholly in terms of the modal properties of agent's own reasons-responsive mechanisms. However, even if we hold fixed all the intrinsic

properties of these mechanisms, the control-related difficulties an agent faces might vary according to the circumstances the agent finds herself in.³² And if this is right—and I think it’s plausible that it is—then the degree to which an agent is morally responsible might depend, at least in part, on the quality of the opportunity she has to exercise her control.

For Nelkin, this means that difficulty depends on *two* variables: the control-grounding mechanisms and the control-grounding circumstances.³³ So if one’s control-grounding mechanisms are more developed, or if one finds oneself in circumstances that better facilitate good decision-making, then it will be easier for you to recognize and respond to good reasons for action. If, on the other hand, one’s control-grounding mechanisms are diminished or impoverished in some way, or if one finds oneself in circumstances that undermine good decisions, then it’ll be harder for you recognize and respond to good reasons. In the former case, you have more control over what you do—you have, in Nelkin’s terminology, a better quality of opportunity. In the latter case, you have less control over what you do, or again, a worse quality of opportunity. This more comprehensive account of control is better equipped to explain why having more or less control affects the degree to which an agent is morally responsible: the quality of opportunity an agent has grounds difficulty-facts of the sort that determine degrees.

4. Conclusion

Agents can be more (or less) morally responsible for their actions. As a result, an adequate theory of moral responsibility owes us (i) an account of what this means and (ii) an explanation why some agents are more (or less) morally responsible. In this chapter I’ve sketched how one might start going about each of these tasks. But it would be foolish (or maybe, foolishly self-deceived) to pretend that I’ve worked out the details all the way here.³⁴ More work is therefore needed.

³² Perhaps the best recent treatment of how an agent’s circumstances can enhance or diminish her control is due to Manuel Vargas (2013). See especially Vargas’ treatment of an agent’s “moral ecology,” (Vargas 2013, 243ff.).

³³ Nelkin would resist characterizing both these elements in terms of control, but I do so because I think, in light of what’s clearly correct about this objection, we should conclude that the nature and degree of control and agent possesses depends at least in part on facts that are extrinsic to the agent.

³⁴ There are at least two glaring omissions in the treatment of these issues here. The first of these simply concerns the nature of difficulty itself. For the most part, I’ve left difficulty unanalyzed in this paper. Elsewhere, Alex Guerrero, Dana Nelkin, and Gwen Bradford (2015) have done some good work in the service of this goal, but more work is needed. The second omission is that I’ve ignored other determinates of degrees. It’s doubtful that *only* difficulty serves as the basis for an agent being more (or less) morally responsible. So we need an account of these other determinates, whatever they are.

Works Cited

- Arpaly, Nomy and Tim Schroeder. 2014. *In Praise of Desire* (Oxford: Oxford University Press).
- Bradford, Gwen. 2015. *Achievement* (New York: Oxford University Press).
- Coates, D. Justin. forthcoming. "Being More (or Less) Blameworthy," *American Philosophical Quarterly*.
- Coates, D. Justin Coates and Philip Swenson. 2013. "Reasons-Responsiveness and Degrees of Responsibility," *Philosophical Studies* vol. 165, no. 2, pp. 629-645.
- Fischer, John Martin and Mark Ravizza. 1998. *Responsibility and Control* (Cambridge: Cambridge University Press).
- Guerrero, Alex. 2017. "Intellectual Difficulty and Moral Responsibility," *Responsibility: The Epistemic Condition* eds. Philip Robichaud and Jan Willem Wieland (Oxford: Oxford University Press), pp. 199-218.
- Hartman, Robert J. 2017. *In Defense of Moral Luck: Why Luck Often Affects Praiseworthiness and Blameworthiness* (New York: Routledge).
- Kagan, Shelly. 2012. *The Geometry of Desert* (New York: Oxford University Press).
- Lenman, James. 2006. "Compatibilism and Contractualism: The Possibility of Moral Responsibility," *Ethics* vol. 117, no. 1, pp. 7-31.
- Levy, Neil. 2014. *Consciousness and Moral Responsibility* (Oxford: Oxford University Press).
- Loomis, Erik. 2015. *Out of Sight* (New York: The New Press).
- McKenna, Michael. 2012. *Conversation and Responsibility* (New York: Oxford University Press).
- Nelkin, Dana K. 2016. "Difficulty and Degrees of Moral Praiseworthiness and Blameworthiness," *Nous* vol. 50, no. 2, pp. 356-378.
- Pereboom, Derk. 2001. *Living Without Free Will* (Cambridge: Cambridge University Press).
- _____. 2014. *Free Will, Agency, and Meaning in Life* (New York: Oxford University Press).
- Rosen, Gideon. 2003. "Culpability and Ignorance," *Proceedings of the Aristotelian Society* vol. 103, no. 1, pp. 61-84.
- Smith, Adam. 1759/1976. *The Theory of Moral Sentiments* (Oxford: Oxford University Press).
- Smith, Holly M. 1991. "Varieties of Moral Worth and Moral Credit," *Ethics* vol. 101, no. 2, pp. 279-303.

Strawson, Galen. 1994. "The Impossibility of Moral Responsibility," *Philosophical Studies* vol. 75, no. 1-2, pp. 5-24.

Strawson, P. F. 1974. "Freedom and Resentment," *Freedom and Resentment and Other Essays* (London: Methuen & Co. Ltd), pp. 1-25.

Tierney, Hannah. n.d. "Reasons and Responsibility."

Vargas, Manuel. 2013. *Building Better Beings* (New York: Oxford University Press).

Watson, Gary. 1996. "Two Faces of Responsibility," *Philosophical Topics* vol. 24, no. 2. pp. 227-248.