



Quality of Reasons and Degrees of Responsibility

Hannah Tierney

University of Sydney

ABSTRACT

Traditionally, theories of moral responsibility feature only the minimally sufficient conditions for moral responsibility. While these theories are well-suited to account for the threshold of responsibility, it's less clear how they can address questions about the degree to which agents are responsible. One feature that intuitively affects the degree to which agents are morally responsible is how difficult performing a given action is for them. Recently, philosophers have begun to develop accounts of scalar moral responsibility that make use of this notion of difficulty [Coates and Swenson 2013; Nelkin 2016]. In this paper, I argue that these accounts, although innovative, are incomplete. The degree to which agents are morally responsible is determined not only by the difficulties that agents face but also by the quality of the reasons for which they act.

ARTICLE HISTORY Received 5 February 2018; Accepted 30 September 2018

KEYWORDS moral responsibility; Coates; Swenson; Nelkin; quality of reasons

1. Introduction

Traditionally, theories of moral responsibility tend to feature only the minimally sufficient conditions for moral responsibility. While these theories are well-suited to account for the threshold of responsibility—they can tell us when an agent is morally responsible and when she is not—it's less clear how these theories can address questions about the degree to which an agent is responsible or about when an agent is more or less responsible than other agents are. But many of our practices of praising, blaming, and holding others morally responsible come in degrees: they have both threshold and scalar features. There are also philosophically compelling reasons to adopt a scalar view of moral responsibility. For example, in earlier work [2013, 2014], I argue that compatibilists can overcome a variant of manipulation arguments by arguing that moral responsibility comes in degrees.¹

But in virtue of what does moral responsibility come in degrees? One feature that intuitively affects the degree to which agents are praiseworthy, blameworthy, and morally responsible are facts about how difficult performing actions can be for agents. The more difficult it is for an agent to do the right thing, the less blameworthy she is for failing to do so (and the more praiseworthy she is if she

¹ Justin Capes [2013] also provides a response to this particular variant of manipulation arguments that relies on a scalar view of moral responsibility.

succeeds). Recently, philosophers have begun to develop accounts of scalar moral responsibility that use this notion of difficulty. Justin Coates and Philip Swenson [2013], expanding on John Martin Fischer and Mark Ravizza's [1998] reasons-responsiveness account of moral responsibility, argue that the degree to which agents are morally responsible for their actions depends on how difficult it is for their actual-sequence mechanisms to react and to recognize sufficient reason to do otherwise. And Dana Nelkin [2016] argues that the difficulty relevant to scalar praiseworthiness and blameworthiness is a function not only of agents' actual-sequence mechanisms but also of their circumstances.²

These accounts vary significantly, although they are alike in one respect: they both focus on difficulty as the underlying feature that affects the degree to which agents are praiseworthy, blameworthy, and morally responsible.³ But our scalar practices of praise and blame are also informed by judgments about the reasons for which agents act. Agents can have better and worse reasons to perform (or not perform) actions, and the quality of agents' reasons—independently of how difficult it is for them to act on those reasons—is directly relevant to the degree to which agents are praiseworthy, blameworthy, and morally responsible. Thus, I argue that scalar theories of responsibility that focus only on difficulty are incomplete, because they ignore the role that the quality of our reasons plays in determining the degree to which we are responsible.

Before proceeding, I'd like to note that in this paper I focus on the conditions under which agents are more or less responsible, and not on the nature of scalar moral responsibility itself.⁴ Of course, questions about the factors that mitigate and enhance responsibility are intimately related to questions about what it means to be more or less responsible, but addressing both sets of questions in a single essay would become unmanageable. For the purposes of this paper, I'll rely on the (fairly) standard view that moral responsibility is closely related to blameworthiness and praiseworthiness [Fischer et al. 2007: 86]:

for an agent to be morally responsible for an action is for it to belong to her in such a way that she would deserve blame if she understood that it was morally wrong, and she would deserve credit or perhaps praise if she understood that it was morally exemplary.

My discussion focuses almost entirely on cases in which agents are blameworthy and praiseworthy to varying degrees, in order to illustrate that moral responsibility is scalar. I'll also assume that, all else being equal, our judgments about these cases can be guides to the degree to which agents are praiseworthy, blameworthy, and morally responsible. Of course, the relationships between blameworthiness, praiseworthiness, and moral responsibility are far from simple. For example, many argue that it's possible to be morally responsible but not praiseworthy or blameworthy for actions that are morally neutral [Fischer and Ravizza 1998]; McKenna 2012]. And some contend that moral responsibility is a purely threshold concept, while

² Manuel Vargas [2013] makes a similar argument in his discussion of moral ecology and blameworthiness.

³ Interestingly, most discussions of scalar responsibility focus on difficulty. Patricia Greenspan, after arguing that difficulty affects freedom, goes on to contend that 'freedom shades into unfreedom in degrees' [1978: 231] and Alex Guerrero, focusing on the epistemic condition, analyses mitigated responsibility in terms of intellectual difficulty [2017].

⁴ For work that focuses on the property of being more or less blameworthy, and not on the conditions under which agents are more or less blameworthy, see Coates [forthcoming].

only praise and blame can come in degrees [Fischer and Ravizza 1998]. Rather than enter these debates, I'll proceed under the assumption (following Coates and Swenson [2013]) that moral responsibility, praiseworthiness, and blameworthiness come in degrees and that the degree to which an agent is blameworthy or praiseworthy is indicative of the degree to which she is morally responsible. However, if it turns out that moral responsibility really is best conceived of as a purely threshold concept, or that the degree to which an agent is blameworthy has no bearing on the degree to which she is morally responsible, then the discussion below can still inform accounts of scalar praise and blame.⁵

2. Difficulty and Degrees of Responsibility

Coates and Swenson [2013] motivate their account of scalar moral responsibility by comparing two cases of mitigated responsibility. In the first case, readers are asked to imagine that Marcia promises to pick up her friend from the airport, but fails to do so. When Marcia is confronted about this failure, she reveals that she is suffering from depression and had trouble getting off the couch on the day when she was supposed to meet her friend. Coates and Swenson argue that, while Marcia's depression doesn't entirely excuse her, it does mitigate the degree to which she is responsible [ibid.: 630]. In the second case, readers are asked to imagine that Thomas also promises to pick up a friend from the airport and fails to do so. However, when confronted, Thomas admits that, when the time came for him to go to the airport, he realized that one of his favourite movies was on television and he didn't want to miss it. In this case, Coates and Swenson argue that Thomas's explanation does nothing to mitigate the degree to which he is responsible (if anything, it may elicit an increase in resentment). While Marcia and Thomas are both morally responsible for breaking a promise to a friend, it seems intuitive that Marcia is less responsible than Thomas are both. According to Coates and Swenson, the driving force behind this intuition is that it would have been more difficult for Marcia to keep her promise. But most theories cannot make sense of cases of moral responsibility and blameworthiness (and praiseworthiness) that come in degrees, for they provide only the minimally sufficient threshold conditions for moral responsibility. Coates and Swenson expand on Fischer and Ravizza's [1998] reasons-responsiveness account of moral responsibility to accommodate these cases.

On Fischer and Ravizza's account, in order for an agent to be morally responsible for a given action, her action must be the product of a moderately reasons-responsive mechanism, which requires that it be both regularly reasons-receptive and weakly reasons-reactive [ibid.: 243–4].⁶

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

⁵ If this is the case, then my use of 'moral responsibility' can be treated much like Nelkin's: she uses 'degrees of moral responsibility' to refer to both degrees of blameworthiness and praiseworthiness, but does so only for ease of exposition and acknowledges the distinction between moral responsibility and praiseworthiness and blameworthiness [2016: 375n1].

⁶ A mechanism is, roughly, a process for making decisions and producing actions [Fischer 2012: 186], and an agent's actual-sequence mechanism is the actual process that she uses to make a given decision or to produce a given action.

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.

Both Marcia's and Thomas's actual-sequence mechanisms meet the thresholds for regular reasons-receptivity and weak reasons-reactivity. Marcia, although she is depressed, and Thomas, although he is distracted, can recognize a host of reasons (including moral reasons) to pick up their friends from the airport. In addition, both of their actual-sequence mechanisms are weakly reasons-reactive because, given sufficient reason to fulfil their promise, there is at least one possible world in which they would each do so.⁷ Thus, both Marcia's and Thomas's decisions to break their promises are the products of moderately reasons-responsive mechanisms, and both meet the threshold for moral responsibility.

Coates and Swenson [2013] argue that the reasons-responsiveness approach can be expanded to account not only for the minimally sufficient conditions for moral responsibility but also for the degree to which agents are morally responsible. In this way, a reasons-responsive view can make sense both of the intuition that Marcia and Thomas are morally responsible and of the intuition that Thomas is more responsible than Marcia. The authors argue further that a difference in reasons-reactivity can ground this intuitive difference in difficulty. It's harder for Marcia to keep her promise, because the possible world in which her actual-sequence mechanism reacts to a sufficient reason to do so is further away than the possible world in which Thomas's actual-sequence mechanism reacts to a sufficient reason to keep his promise.⁸ They argue [*ibid.*: 639]:

On this view, it is difficult for Marcia's actual-sequence mechanism to react to sufficient reason to do otherwise because the nearest possible world in which the mechanism reacts according to a sufficient reason to do otherwise is distant (in terms of comparative similarity) from the actual world. And it is comparatively easier for Thomas's actual-sequence mechanism to react to sufficient reason to do otherwise because the nearest world in which the mechanism reacts according to sufficient reason to do otherwise is close (in terms of comparative similarity) to the actual world.

Presumably, little about the actual world would need to be altered in order for Thomas to decide to go to the airport; he would need only to be slightly less lazy or to have slightly more affection for his friend. But the actual world would have to undergo significant changes in order for Marcia to be able to keep her promise. She would have to be significantly less depressed or to have significantly weightier reasons in favour of driving to the airport. In short, Marcia is less responsible than Thomas for breaking a promise because it was harder for her to keep it, which is to say that the possible world in which her actual-sequence mechanism reacts to sufficient reason to keep her promise is less similar to the actual world than is the

⁷ Many philosophers have argued that weak reasons-reactivity is too low a bar for control and moral responsibility (e.g. McKenna [2005]). Recently, in response to several critics (e.g. Mele [2000]), Fischer [2012] has argued that the threshold to be appropriately reasons-reactive may be higher than mere weak reasons-reactivity. So, while there's controversy over where the threshold should be for reasons-reactivity, the debate has focused on exactly that—a threshold.

⁸ Coates and Swenson admit that determining comparative similarity and distance between worlds is difficult, but, for the purposes of this paper, I would like to grant that it's possible to make such comparisons.

possible world in which Thomas's actual sequence mechanism reacts to sufficient reason to keep his promise. Thus, by expanding on Fischer and Ravizza's reasons-responsive view, Coates and Swenson provide an elegant account of how moral responsibility can admit of degrees.⁹

Nelkin [2016], like Coates and Swenson, defends a scalar view of moral responsibility by arguing that the driving force behind our intuition that an individual can be more or less blameworthy or praiseworthy for performing a given action lies in the difficulty the individual faced in performing that action. Indeed, Nelkin contends that it is a desideratum of scalar views of moral responsibility to 'account for the idea that degrees of responsibility depend on (or at least seem to depend on) degrees of difficulty understood in terms of expenditure of effort and sacrifice required' [ibid.: 357]. However, Nelkin does not think that Coates and Swenson's reasons-responsiveness account can adequately explain scalar responsibility in terms of difficulty. Nelkin argues that difficulty is not entirely a feature of agents' decision mechanisms; it can also depend on environmental factors. Even if an agent's actual-sequence mechanism is maximally reasons-responsive, features of her environment can render it difficult for her to perform the morally obligatory act in a given situation, which can, intuitively, make the agent less blameworthy for not performing said act. To accommodate such cases, Nelkin offers an account of moral responsibility that relies on the notion of agents' quality of opportunity, which depends on both agents' capacities (such as their actual-sequence mechanisms) and situational features of their environment. As Nelkin argues, 'The *more* fair the opportunity the more blameworthy one is for not taking it and less praiseworthy one is for doing so' [ibid.: 371].

There is much more to say about both Coates and Swenson's and Nelkin's detailed and careful accounts, but at this juncture I would like simply to note the emphasis that both theories place on the relationship between difficulty and responsibility. Everyone who is party to the debate grants that the driving intuition behind why Marcia is less responsible than Thomas is because it's more difficult for her to keep her promise. Coates and Swenson argue that the difficulty Marcia faces can be made sense of in terms of reasons-responsiveness, while Nelkin argues that this difficulty can be explained in terms of the quality of Marcia's opportunity.

3. Quality of Reasons and Degrees of Responsibility

But difficulty is far from the only feature that is relevant to moral responsibility. The degree to which agents are praiseworthy and blameworthy is also informed by the reasons for which the agents act. An agent can have better and worse reasons to perform (or not perform) actions, and the quality of agents' reasons is directly relevant to the degree to which they are praiseworthy, blameworthy, and morally responsible. When trying to determine the degree to which an agent is morally responsible for wronging you, a natural question to ask concerns *why* the agent

⁹ In addition to scalar reasons-reactiveness, Coates and Swenson also discuss how reasons-receptivity can come in degrees and how differences in reasons-receptivity can ground differences in moral responsibility. They argue, for example, that differences in reasons-receptivity can explain why adolescents are, on the whole, less morally responsible for their behaviour than average adults are [2013: 641]. For the purposes of this paper, however, it will only be necessary to focus on Coates and Swenson's account of reasons-reactivity.

wronged you.¹⁰ Just as agents' reasons can excuse or justify morally wrong actions, they are also relevant to the degree to which agents are blameworthy. For example, if my spouse fails to go grocery shopping as they promised, I would do well to take their reasons into account when determining the degree to which they are blameworthy. If they failed to go to the store because they thought that the dog looked like she wanted some exercise (and so they took her on a walk instead), I would be inclined to think that they are less blameworthy than if they failed to go to the store because, upon reflection, they decided that grocery shopping was undignified. While neither reason is sufficient to excuse or justify breaking a promise, and both reasons may very well rest on false beliefs, they still exert different influences on how blameworthy my spouse is.

Notice that the role played by difficulty is negligible in the above case. If my spouse broke their promise because they found grocery shopping undignified, I would maintain the judgment that they were worthy of the same (high) degree of blame even if their haughty self-image made it very difficult for them to be able to keep their promise. And if they broke their promise because they believed that the dog wanted to go for a walk, I would maintain the judgment that they deserved the same (moderate) degree of blame even if this belief was relatively easy to overcome.

This intuition becomes even more salient if we imagine that my spouse fulfilled their promise, despite facing obstacles. I may judge that my spouse is somewhat praiseworthy for overcoming their desire to walk the dog and fulfilling their promise to go grocery shopping, but I don't think that my spouse would be considered at all praiseworthy if they overcame their haughty self-image to do so. And one can imagine that the more my spouse (even sincerely) dwells on how difficult it was to go grocery shopping because they deem it beneath them, the less appreciative one would feel. In many cases, agents' reasons for acting seem just as, if not more, relevant to how responsible they are as do the reasons-responsiveness of their actual-sequence mechanisms and the quality of their opportunities. While I certainly don't want to argue that concerns regarding difficulty never mitigate responsibility, it strikes me that we often inquire into agents' reasons in order to determine the degree to which they are praiseworthy, blameworthy, and morally responsible for a given action, perhaps even more frequently than we inquire as to the difficulty they faced in performing that action.¹¹

One might argue that a spouse who breaks their promise because they believe that grocery shopping is beneath them and a spouse who breaks their promise

¹⁰ In this essay, I focus on the evaluation of motivating or explanatory reasons, as opposed to normative reasons (see Alvarez [2010] for a discussion of the distinction between normative, motivating, and explanatory reasons).

¹¹ In fact, I submit that the intuition that Thomas is more blameworthy than Marcia can be maintained even if we stipulate that it was just as difficult for Thomas to keep his promise as it was for Marcia. After all, while someone like Marcia might find it more difficult to keep her promise to drive to the airport than someone like Thomas does, this needn't be the case. Imagine that Thomas's love for the movie on television makes it just as difficult for him to keep his promise as it is for Marcia to keep hers. Imagine further that Thomas's actual-sequence mechanism is exactly as reasons-reactive as Marcia's: the world in which his actual-sequence mechanism reacts to a sufficient reason to keep his promise is just as far away as the possible world in which Marcia's actual-sequence mechanism issues a different judgment. We can also imagine that the quality of Thomas's opportunity to do the right thing was exactly the same as Marcia's: both Thomas and Marcia live equally far from the airport, etc. Yet despite the fact that, on all given accounts of difficulty, Thomas and Marcia would find keeping their promises equally difficult, it still strikes me that Thomas is more blameworthy for his failure to keep his promise than Marcia is. Perhaps this is because not wanting to miss a movie on television is a worse reason to break one's promise than are the reasons to stay home that are generated by depression.

because they want to take their dog for a walk are not blameworthy to different degrees for performing the same action.¹² Rather, because the spouses act on different reasons, they perform different actions, and one is morally worse than the other. If this is the case, then agents' reasons can't really affect the degree to which agents are responsible for the same actions; they can only affect the moral qualities of the actions themselves.

Individuating actions is a tricky business, and I won't defend a particular view of action individuation here. It may very well be true that in some cases an agent's reasons can affect the kind of action that they perform. Nevertheless, this needn't preclude the possibility that in other cases an agent's reasons can affect the degree to which they are responsible for an action without thereby affecting the moral quality of that action. The view that agents can perform the same kind of action for distinct reasons is an important feature of our moral lives. It's what allows us to construct manageable and informative moral codes, for example.

Furthermore, all parties to discussions of scalar responsibility assume that it's possible to perform the same action for different reasons. Recall that Coates and Swenson motivate their reasons-responsive account with the case of Marcia and Thomas, two agents who fail to pick up their respective friends from the airport for presumably very different reasons. Yet Coates and Swenson compare the degree to which these agents are responsible for performing the very same action—breaking a promise. Nelkin also grants the intuition that Marcia is less responsible than Thomas for breaking her promise, although Nelkin offers a different explanation as to why this is so. If it turns out that two agents cannot perform the same action for different reasons, then both Coates and Swenson and Nelkin will have to rely on very different kinds of cases to motivate their views.¹³

But perhaps there are good independent reasons to favour an account of action individuation on which agents who act on different reasons necessarily perform different kinds of actions. Even if this is the case, agents' reasons can still affect the degree of blame and praise that an agent deserves, by affecting the moral quality of the actions that the agent performs. However, on this view of action individuation, the quality of agents' reasons would not directly inform the degree to which they are *morally responsible*, since deserving more blame in virtue of performing a worse action does not entail that an agent is more responsible for that wrongdoing. This would limit the role that the quality of reasons plays in discussions of moral responsibility, although the relationship between reasons, blame, and praise would still be an interesting and worthwhile topic to explore. In short, even if acting on distinct reasons entails that agents perform distinct actions, attending to the quality of agents' reasons can still play an important role in discussions of scalar blameworthiness and praiseworthiness.

¹² Thanks to Aness Webster for bringing this objection to my attention.

¹³ Additionally, a fine-grained approach to action individuation may lead to analogous worries for difficulty-based accounts of scalar responsibility. One could argue that if breaking a promise for a good reason makes the breaking of that promise a morally better action, perhaps breaking a promise when it's very difficult to avoid doing so makes that action morally better as well.

4. What Makes a Reason Better or Worse

I've submitted that the quality of agents' reasons affects the degree to which agents are praiseworthy, blameworthy, and morally responsible. While few would deny that the quality of agents' reasons is relevant to responsibility,¹⁴ the notion that the quality of our reasons can affect *the degree to which* we are responsible, praiseworthy, and blameworthy has yet to be fully explored. But what makes a reason good or bad? Reasons can be better or worse than other reasons in a variety of ways. In the following section, I'll examine two dimensions along which the quality of reasons can be measured, and I'll discuss how variance along each dimension can independently affect the degree to which we are morally responsible.

4.1. The Moral Quality of Reasons

Let's return to the case of Marcia and Thomas, with a slight modification. Imagine that Marcia, rather than being depressed, fails to keep her promise to her friend, because she decides to finish a batch of grading, which she hasn't promised to complete but which she knows her students will appreciate.¹⁵ In this case, even if it's just as difficult for Marcia to tear herself from her grading as it is for Thomas to tear himself from his couch, it still strikes me that Marcia is less blameworthy than Thomas. There is a sense in which wanting to finish a batch of grading is a *morally* better reason to break a promise than is wanting to watch a movie on television. There are many ways to conceive of the moral difference between these reasons. One way would be to focus on how Marcia's and Thomas's reasons affect their *quality of will*. Although much work has been done on understanding moral responsibility, praiseworthiness, and blameworthiness in terms of quality of will, as Nelkin argues [2016: 360], 'those who argue for [a quality of will account] do not often examine in detail how *degrees* of blameworthiness and praiseworthiness are determined and how exactly they are related to quality of will.'

An agent's quality of will can be understood as her concern or regard for others, and is determined by her reasons, in addition to her intentions, choices, character traits, etc. [Arpaly 2003; McKenna 2012]. One could argue that, while Marcia's desire to finish her batch of grading diminishes the degree of concern or regard that she has for her friend, Thomas's desire to watch a movie on television renders his quality of will far worse. After all, Marcia is prioritizing pleasing her students over the needs of her friend, which is somewhat morally objectionable, but

¹⁴ For example, several reasons-responsive accounts of moral responsibility ground reasons-responsiveness in agents' abilities to recognize and react to *good* reasons. Susan Wolf [1990: 117], in developing the Reason View, argues.

a person's status as a responsible agent rests not only on her ability to make her behaviour conform to her deepest values but also on her ability to form, assess, and revise those values on the basis of a recognition and appreciation of ... the True and the Good.

In addition, Nelkin [2011: 7] argues that 'people are responsible when they act with the ability to do the right thing for the right reasons, or a good thing for good reasons.' And, in an essay that explores moral responsibility and mental illness, Benjamin Kozuch and Michael McKenna argue: 'An agent must be able to react to a considerable range of the *good reasons* she is able to recognize by choosing and acting on the basis of such reasons' [2015: 101, emphasis added].

¹⁵ Thanks to Derk Pereboom for suggesting this modification of the Marcia and Thomas case.

Thomas is prioritizing his entertainment over his friend's needs, which is much more morally objectionable. In this way, Thomas's reason is morally worse than Marcia's, and this is why he is more morally responsible and blameworthy for breaking his promise.¹⁶

4.2. The Epistemic Quality of Reasons

Reasons can also be *epistemically* better than other reasons. For example, we can imagine that Thomas acts on an *epistemically* worse reason than Marcia does. Imagine that both Marcia and Thomas break their promises to pick up their friends from the airport because they believe that their friends' flights are delayed. However, Thomas believes that his friend's flight is delayed because he had a dream that the flight was delayed, while Marcia believes that her friend's flight is delayed because its online status was listed as such the day before it was scheduled to arrive. It strikes me that both Marcia and Thomas are morally responsible and blameworthy for breaking their promises, because they broke them for epistemically bad reasons.¹⁷ One shouldn't rely on dreams or on out-of-date online statuses to form beliefs about flight arrival times. But while both agents are morally responsible and blameworthy, they aren't equally so, because, although both reasons are epistemically bad, they're not equally bad. A day-old flight status, while unreliable, is still more reliable than a dream when it comes to justifying beliefs about the arrival times of flights.¹⁸ In this case, it seems that although the degree to which Marcia and Thomas are morally responsible for breaking their respective promises depends neither on the difficulty that they faced in breaking their promises nor on the moral quality of their reasons for doing so. Rather, our judgments of moral responsibility track the epistemic quality of their reasons: Marcia is less responsible for breaking her promise than Thomas is, because she had an epistemically better (although still bad) reason to do so.

Notice that the role that the epistemic quality of our reasons plays in determining the degree to which we are morally responsible is independent from the role that the moral quality of our reasons plays, and cannot easily be captured by quality of will accounts. While there may be cases where relying on a dream to form a

¹⁶ Nelkin [2016] criticizes quality of will views defended by Nomy Arpaly [2003] and Arpaly and Timothy Schroeder [2014], along with deep or real self views, like the account defended by David Faraci and David Shoemaker [2010], because such views cannot adequately capture the sense in which difficulty affects our scalar judgments of praise and blame. However, if the above reasoning is correct, we shouldn't expect all of our scalar judgments of moral responsibility to track only difficulty. Our judgments about moral responsibility are responsive to facts about the moral quality of agents' reasons in addition to facts about difficulty. And although quality of will views may not be able to accommodate the role that difficulty plays in moral responsibility, they can account for the effect that the moral quality of our reasons has on our scalar judgments of praise and blame. Because Nelkin, Coates and Swenson, focus solely on the phenomenon of difficulty in developing their scalar views, they will be less able to capture the relationship between the moral quality of our reasons and responsibility.

¹⁷ This is so, assuming that both Marcia and Thomas are otherwise normal moral and epistemic agents. We can also assume that, although Marcia's and Thomas's reasons are epistemically bad, they aren't *morally* bad (or at least not to a significantly different degree) and neither Marcia nor Thomas is blameworthy for their belief (or, again, at least not to a significantly different degree).

¹⁸ While I talk here in terms of reliability, one can utilize any account of the epistemic strength of reasons to measure the differences between the two reasons in question.

belief about the arrival time of a friend's flight illustrates a poorer quality of will than does relying on an out-of-date flight status, this needn't be the case. It is in principle possible for the varying epistemic qualities of two agents' reasons to affect the degree to which the agents are responsible, holding fixed all morally relevant features of the case. In this way, it is unlikely that a scalar account of moral responsibility developed to accommodate solely the moral quality of our reasons will be able also to capture the epistemic quality of our reasons. This worry could become even more vexed if we identify further dimensions along which the quality of our reasons can vary.¹⁹

4.3. Accounting for the Quality of Reasons

I have claimed that there are at least two ways of conceiving of the quality of reasons, each of which can independently affect the degree to which agents are morally responsible. One might worry that, the more dimensions along which the quality of our reasons can vary, the more difficult it will be for a single theory to adequately capture all of the ways in which the quality of our reasons affects the degree to which we are morally responsible. This may very well be true, but we cannot hope ever to develop a complete scalar account of moral responsibility if we don't first isolate and examine all of the features that determine the degree to which agents are morally responsible. For example, while Coates and Swenson and Nelkin are surely correct that difficulty affects the degree to which we are praiseworthy, blameworthy, and morally responsible, views that can account only for the role that difficulty plays in scalar responsibility are bound to be incomplete.

One may wish to argue that, despite appearances, some of these features don't actually determine the degree to which agents are morally responsible. Perhaps Coates and Swenson, or Nelkin, will deny that the quality of agents' reasons, along any dimension, should play a role in determining the degree to which agents are morally responsible. This is an in-principle tenable position. However, one cannot defend this view adequately without explaining away the appearance that the quality of reasons affects the degree to which agents are morally responsible. Furthermore, these views will likely have to compete with accounts that *can* accommodate the roles that both difficulty and the quality of our reasons play in moral responsibility, and will thus operate at a disadvantage in at least this respect. In this way, even if one wishes to ultimately deny the role that the quality of our

¹⁹ For example, one might think that the popularity of reasons affects the degree to which agents are morally responsible. In cases where all or most agents recognize and react to a set of reasons to engage in morally wrong behaviour, one could argue that this mitigates the degree to which agents are morally blameworthy for performing such actions (without thereby committing oneself to the more controversial claim that the ubiquity of a moral wrong makes that action less wrong). However, one could argue that, while we tend to judge agents as less responsible for committing moral wrongs for popular reasons, this is not because the popularity of reasons itself plays a role in determining the degree to which agents are morally responsible. After all, consensus often operates as evidence: perhaps popularity affects the degree to which we judge agents as morally responsible only because it's treated as a proxy for the moral and/or epistemic quality of a reason. It's also possible that it's more difficult (or we perceive it to be more difficult) to resist acting on a popular reason, and that this is why we judge agents who perform morally wrong actions for popular reasons to be less morally responsible. Here I won't wade too deeply into these waters, but exploring the relationship between popularity, difficulty, and the moral and epistemic quality of reasons could prove fruitful in further developing scalar accounts of responsibility.

reasons plays in determining the degree to which we are morally responsible, one cannot hope to successfully develop a scalar account of moral responsibility without addressing the role of reasons.

5. Conclusion

In this essay, I have argued that if we want to make sense of moral responsibility that comes in degrees then we must look beyond accounts that rely only on the notion of difficulty. Often, our scalar practices of attributing moral responsibility, praise, and blame to other agents are responsive not just to facts about difficulty, but also to facts about the quality of agents' reasons. If the above reasoning is correct, no account of scalar moral responsibility can be complete if it ignores the role that agents' reasons play in determining the degree to which agents are praiseworthy, blameworthy, and morally responsible.²⁰

References

- Alvarez, M. 2010. *Kinds of Reasons: An Essay on the Philosophy of Action*, Oxford: Oxford University Press.
- Arpaly, N. 2003. *Unprincipled Virtue: An Inquiry into Moral Agency*, New York: Oxford University Press.
- Arpaly, N. and T. Schroeder 2014. *In Praise of Desire*, New York: Oxford University Press.
- Capes, J. 2013. Mitigating Soft Compatibilism, *Philosophy and Phenomenological Research* 87/3: 640–63.
- Coates, D.J. forthcoming. Being More (or Less) Blameworthy, *American Philosophical Quarterly*.
- Coates, D.J. and P. Swenson 2013. Reasons-Responsiveness and Degrees of Responsibility, *Philosophical Studies* 165/2: 629–45.
- Faraci, D. and D. Shoemaker 2010. Insanity, Deep Selves, and Moral Responsibility: The Case of JoJo, *Review of Philosophy and Psychology* 1/3: 319–32.
- Fischer, J.M. 2012. *Deep Control: Essays on Free Will and Value*, New York: Oxford University Press.
- Fischer, J.M. and M. Ravizza 1998. *Responsibility and Control: A Theory of Moral Responsibility*, Cambridge: Cambridge University Press.
- Fischer, J.M., R. Kane, D. Pereboom, and M. Vargas 2007. *Four Views of Free Will*, Malden, MA: Blackwell.
- Greenspan, P.S. 1978. Behavior Control and Freedom of Action, *The Philosophical Review* 87/2: 225–40.
- Guerrero, A.A. 2017. Intellectual Difficulty and Moral Responsibility, in *Responsibility: The Epistemic Condition*, ed. P. Robichaud and J.W. Wieland, Oxford: Oxford University Press: 199–218.
- Kozuch, B. and M. McKenna 2015. Free Will, Moral Responsibility, and Mental Illness, in *Philosophy and Psychiatry: Problems, Intersections, and New Perspectives*, ed. D. Moseley and G. Gala, New York: Routledge: 89–113.
- McKenna, M. 2005. Reasons Reactivity and Incompatibilist Intuitions, *Philosophical Explorations* 8/2: 131–43.

²⁰ For helpful comments and suggestions on previous drafts of this paper, I would like to thank Michael McKenna, Shaun Nichols, Derk Pereboom, and the anonymous referees. I would also like to thank the audiences at my 2017 Pacific APA colloquium presentation, with special thanks to Aness Webster who provided comments, and at the Conference on Reassessing Responsibility, hosted at Trinity Hall, Cambridge and organized by Lubomira Radoilska. Finally, many thanks to the organizers and participants of the 2018 Philosophy Desert Workshop, especially Nicolas Cornell and Jennifer Matey for their comments.

- McKenna, M. 2012. *Conversation and Responsibility*, New York: Oxford University Press.
- Mele, A.R. 2000. Reactive Attitudes, Reactivity, and Omissions, *Philosophy and Phenomenological Research* 61/2: 447–52.
- Nelkin, D.K. 2011. *Making Sense of Freedom and Moral Responsibility*, Oxford: Oxford University Press.
- Nelkin, D.K. 2016. Difficulty and Degrees of Moral Praiseworthiness and Blameworthiness, *Noûs* 50/2: 356–78.
- Tierney, H. 2013. A Maneuver Around the Modified Manipulation Argument, *Philosophical Studies* 165/3: 753–63.
- Tierney, H. 2014. Tackling it Head On: How Best to Handle the Modified Manipulation Argument, *The Journal of Value Inquiry* 48/4: 663–75.
- Vargas, M. 2013. *Building Better Beings: A Theory of Moral Responsibility*, Oxford: Oxford University Press.
- Wolf, S. 1990. *Freedom Within Reason*, New York: Oxford University Press.