

Reasons-responsiveness and degrees of responsibility

D. Justin Coates · Philip Swenson

Published online: 9 June 2012
© Springer Science+Business Media B.V. 2012

Abstract Ordinarily, we take moral responsibility to come in degrees. Despite this commonplace, theories of moral responsibility have focused on the minimum threshold conditions under which agents are morally responsible. But this cannot account for our practices of holding agents to be more or less responsible. In this paper we remedy this omission. More specifically, we extend an account of reasons-responsiveness due to John Martin Fischer and Mark Ravizza according to which an agent is morally responsible only if she is appropriately receptive to and reactive to reasons for action. Building on this, we claim that the degree to which an agent is responsible will depend on the degree to which she is able to recognize and react to reasons. To analyze this, we appeal to relations of comparative similarity between possible worlds, arguing that the degree to which an agent is reasons-reactive depends on the nearest possible world in which given sufficient reason to do otherwise, she does so. Similarly, we argue that the degree to which an agent is reasons-receptive will depend on the intelligibility of her patterned recognition of reasons. By extending an account of reasons-responsiveness in these ways, we are able to rationalize our practice of judging people to be more or less responsible.

1

Suppose your friend Marcia promises to pick you up at the airport, but at the appointed time, she fails to show up. It would be natural to resent Marcia and just as natural to, upon seeing Marcia, to express that resentment. But suppose that as you are rebuking Marcia, she admits to you that she's been suffering from serious but

D. J. Coates (✉) · P. Swenson
Department of Philosophy, University of California, Riverside, CA, USA
e-mail: dancoates1@gmail.com

non-debilitating bouts of depression, and that she was having trouble getting off of the couch on the day she was supposed to pick you up. This admission, we think, should lead you to revise your response to Marcia. Perhaps it does not fully *excuse* her failure to pick you up, but it does seem to *mitigate* the degree to which she is responsible for that failure.

Alternatively, suppose that on a later trip, your friend Thomas similarly breaks his promise and fails to pick you up. And suppose that as you are rebuking Thomas, he simply says that although he knew that he was obligated to pick you up, he just did not feel like it at the time since one of his favorite movies was on television and he wanted to watch it. Unlike in the previous case, there is no reason to suppose that this admission should lead you to revise your response to Thomas.¹ In this case, he treated you callously, as if his relatively minor whims (e.g., seeing a movie he's already seen many times) legitimately outweighed the normative significance of his promise to you and your need at that moment. And as such, it seems clear that he is a fitting target of your resentment.

So despite suffering from depression, it seems to us that in the first case, Marcia is morally responsible for her failure to pick you up. This judgment will be borne out when we consider the conditions under which agents satisfy some minimal threshold of moral responsibility. Likewise, in the second case, Thomas is morally responsible for his failure to pick you up. And again, this is because Thomas satisfies the minimal threshold conditions on moral responsibility. However, it seems implausible that Marcia is just as responsible as Thomas. That is, they seem to possess different degrees of responsibility. Specifically, Thomas seems to deserve more blame than Marcia for his failure to pick you up. That is, in the second case Thomas is more blameworthy for his action than Marcia is for her action.

2

This suggests that moral responsibility and blameworthiness come in degrees.² But this putative fact concerning moral responsibility and blameworthiness is in need of explanation. For although many theorists supply plausible threshold conditions on responsibility, thereby outlining the minimal control agents must possess, they have typically failed to extend the account in a way that can make sense of the intuitive claim that of the agents in these two cases, Thomas is more responsible and more blameworthy for his failure to pick you up than Marcia is in the first case. This paper aims to remedy this glaring omission.

¹ In fact, if Thomas's explanation did lead you to revise your response, it would probably do so by heightening your resentment since his action was not simply motivated by forgetfulness or neglect but by callous disregard.

² Some theorists, e.g., Fischer and Ravizza (1998), hold that moral responsibility is merely a threshold concept, while maintaining that blameworthiness comes in degrees. We diverge from these theorists on this point, thinking that although an account of moral responsibility should specify the minimum threshold conditions on responsible agency, it should be consistent with thinking that among agents who satisfy those conditions, they enjoy differential degrees of responsibility. More controversially, we believe that a full account of moral responsibility must actually specify the conditions under which agents are more or less morally responsible for their actions.

As a first approximation, we think that the degree to which an agent is morally responsible depends on the degree to which she controls her action.³ But more specifically, we extend an account of control due to John Martin Fischer and Mark Ravizza.⁴ That is, we explain the conditions under which agents are more or less responsible by building upon each of the two components that Fischer and Ravizza take to be the basis of an agent's control.

First we argue that how responsible and blameworthy an agent is for her action can depend on the relevant modal properties of the actual-sequence moderately reasons-responsive mechanism that issues in her action. If those modal properties are made true by comparatively similar worlds, we argue that the agent in question enjoys a high degree of responsibility. Conversely, if those modal properties are made true by comparatively dissimilar worlds, the agent in question is less responsible. And second we argue that how responsible and blameworthy an agent is can depend on whether her sensitivity to a domain of reasons exhibits an intelligible pattern. Of course, because patterns of reasons-recognition can be more or less intelligible, agents can be more or less responsible or blameworthy in virtue of the degree to which (i) she is sensitive to the domain of reasons in question, or (ii) the degree to which her sensitivity fits an intelligible pattern.

To fully develop these claims, we begin by explaining Fischer and Ravizza's account of moderate reasons-responsiveness, which is plausibly thought of as a necessary component of the control condition on moral responsibility and blameworthiness.⁵ We then extend their account in a natural and elegant way—a way that will explain why Marcia is less responsible and less blameworthy than Thomas. In short, we explain their differential degrees of responsibility and

³ There are at least two further ways in which it is plausible to think that the degree to which an agent is morally responsible can vary. First, an agent's degree of responsibility might depend on the degree to which she satisfies the epistemic condition on moral responsibility. Second, her degree of responsibility might depend on how onerous or demanding the agent's obligations are. One plausible way of accounting for this might be by appealing to the claim (recently defended by Erin Kelley [forthcoming]) that an agent is excused when it is unreasonable to demand that she comply with moral principles. If the reasonableness of demanding that an agent act morally comes in degrees, then plausibly, an agent's responsibility will be scalar as well. Unfortunately we do not have space to explore these issues further in the present discussion.

⁴ Fischer and Ravizza's account of control is particularly well suited to the account we will develop because the components of control they identify clearly come in degrees and have the requisite modal properties our account requires. We suspect that with suitable clarification, the reasons-based accounts of moral responsibility due to Susan Wolf (1990), Jay Wallace (1994), and Dana Nelkin (2011) could adopt accounts similar to ours. But we will not explore this here. Moreover, nothing in our discussion here presupposes the truth of compatibilism (or of Fischer and Ravizza's semi-compatibilism), and so incompatibilists who accept that some form of reasons-responsiveness is a necessary condition on moral responsibility can make use of the account we offer here. Of course, the incompatibilist would need to accept reasons-responsiveness as an independent requirement for moral responsibility in order to accept our account. Favored incompatibilist components of control, such as the ability to do otherwise, do not appear to have the modal properties needed to play the role that reasons-responsiveness will play in our account. (See note 15).

⁵ We want to emphasize that even if Fischer and Ravizza's analysis of reasons-responsiveness is incorrect, something in its neighborhood is true, and our account can apply, *mutatis mutandis*, to the correct analysis of reasons-responsiveness.

blameworthiness in terms of their differing degrees of reasons-reactivity.⁶ Next we consider another case of differential degrees of responsibility and blame. Here we extend Fischer and Ravizza's account in a way that will explain why e.g., young adolescents are intuitively less responsible and less blameworthy than normal adults in terms of differing degrees of reasons-receptivity.⁷

3

We begin with a discussion of the threshold conditions on moral responsibility, i.e., those minimal conditions that an agent must satisfy if she is to be morally responsible (to any degree at all) for her actions. Like Fischer and Ravizza, we claim that an agent minimally satisfies the control condition on moral responsibility only if her action issues from an appropriately "reasons-responsive" mechanism. The thought behind this claim is that if the mechanisms that actually issue in our actions are not suitably sensitive to reasons for and against a particular course of action, and if they are not suitably responsive to those reasons, then we cannot be responsible for our actions. This is a plausible way to capture the intuitive idea that moral responsibility requires (i) that we have the capacity to recognize reasons, particularly moral reasons, and (ii) that we have the capacity to react to those reasons.⁸

But what's involved in a mechanism being "appropriately" reasons-responsive or "suitably" reasons-*receptive* and reasons-*reactive*? That is, what distinguishes an insufficiently reasons-responsive mechanism from an appropriately reasons-responsive mechanism?

Fischer and Ravizza have a sophisticated and subtle answer to these questions. According to Fischer and Ravizza, a mechanism is appropriately reasons-responsive when the mechanism in questions is "moderately reasons-responsive." More perspicuously, Fischer and Ravizza claim that:

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

⁶ Of course, as we'll see, in cases of clinical depression, agents are less reasons-receptive. Thus, it's plausible that a full explanation of Marcia's mitigated blameworthiness will require a discussion of the reasons-receptivity of those mechanisms that issue in her actions.

⁷ Again, a full account of teenagers' mitigated responsibility will require a discussion of not only their reasons-receptivity but also their reasons-reactivity.

⁸ Obviously, satisfying these conditions is not sufficient for moral responsibility, but it is necessary, and so it will be true of any responsible agent that she is appropriately reasons-receptive and appropriately reasons-reactive.

reasons (some of which are moral), and at least weakly reactive to reasons (Fischer and Ravizza 1998, pp. 243–244).

Here Fischer and Ravizza claim that a mechanism is appropriately reasons-responsive just in case it is (i) regularly reasons-receptive and (ii) weakly reasons-reactive. Of course, a bit more needs to be said to clarify Fischer and Ravizza's conception of moderate reasons-responsiveness.

First, we can say that a mechanism is reasons-receptive when it is sensitive to reasons for and against a particular course of action. Of course, *bare* reasons-receptivity is too weak to ground moral responsibility. Suppose, for example, that tickets to the Super Bowl cost \$1,000, and further suppose that (plausibly) this constitutes a reason to refrain from buying tickets to the Super Bowl.⁹ The mere fact that in deliberation Eric recognizes the \$1,000 price tag to be a reason to refrain from buying Super Bowl tickets is insufficient to show that Eric's deliberation (i.e., the actual-sequence mechanism that issues in his action) is appropriately reasons-receptive. After all, if in deliberation, Eric did not similarly recognize a \$1,001; \$1,002; \$2,000; or \$5,000 price tag to be a reason to refrain from by Super Bowl ticket then Eric seemingly fails to grasp the *significance* of the reason to refrain from buying \$1,000 tickets. As Fischer and Ravizza put it, "we are not only concerned to see that a person acting on that mechanism recognizes a sufficient reason in one instance; we also want to see that the person exhibits an appropriate *pattern* of reasons-recognition" (Fischer and Ravizza 1998, pp. 70–71).

Second, Fischer and Ravizza say that a mechanism is weakly reasons-reactive just in case given sufficient reason to do otherwise, the mechanism in question will react to that reason in at least one possible scenario. Now unlike regular reasons-receptivity, which requires that a mechanism be sensitive to reasons across a wide and patterned range of possibilities, weak reasons-reactivity is a fairly low standard. Thus, mechanisms are weakly reasons-reactive if and only if given sufficient reason to do otherwise there is at least one possible world in which the mechanism in question reacts to the sufficient reason to do otherwise.

Because this standard for reasons-reactivity is so low, some have argued that Fischer and Ravizza need to strengthen this requirement a bit. For example, Michael McKenna (2005) suggests that moderate reasons-responsiveness plausibly requires that a mechanism have "moderate reasons receptivity and *weaker* reactivity" [p. 135]. Here "weaker reactivity" requires more than just one scenario in which the agent reacts to a reason to act differently.¹⁰ However, it still does not require a pattern as broad and significant as regular reasons-receptivity. Although we find McKenna's suggestion plausible, we do not take a firm stance here on what the minimal threshold of reasons-reactivity should be. Whether one accepts Fischer and Ravizza's original proposal or McKenna's modification, reasons-reactivity will come in degrees in ways modeled by our account.

With this (brief) characterization of Fischer and Ravizza's account of moderate reasons-responsiveness in mind, we apply it to the case of Marcia.

⁹ We model this case on a case found in Fischer and Ravizza (1998).

¹⁰ Similarly, Todd and Tognazzini (2008) suggest strengthening the standards for regular reasons-receptivity.

4

Recall that although she promised to pick you up, Marcia failed to show up and you were left stranded at the airport. Upon rebuke, she confesses that she's been having a hard time lately dealing with her depression, and this softens your response. After all, her depression intuitively mitigates her responsibility and blameworthiness. However, it is plausible that despite suffering from depression, Marcia is, to some degree, responsible for her failure to pick you up. And Fischer and Ravizza's account of moderate reasons-responsiveness can help to explain this fact.

If Marcia is to be morally responsible for her failure to pick you up, the mechanism that actually issues in her action (or her omission in this case) must be regularly reasons-receptive and weakly reasons-reactive. In this case, let us suppose that the mechanism in question is Marcia's practical reasoning or deliberative mechanism. And further, let us remember that this mechanism is deeply affected by her depression.¹¹ Thus, for the purposes of our discussion, the actual-sequence mechanism that issues in Marcia's failure to come pick you up is her depressed practical reasoning. With that in mind, Fischer and Ravizza's account comes into focus.

First, it is plausible to think that Marcia's deliberations are regularly reasons-receptive. In her apology, Marcia fully admits that she had sufficient reason to come and pick you up at the airport and that she recognized the reason as such.¹² Moreover, we can suppose that her recognition of this reason is not disconnected from an understandable pattern since, as she's apologizing, she mentions that lately, she's been failing all her friends. Here, Marcia is admitting not only that she has sufficient reason to keep her promise to you, but she is also admitting that in relevantly similar cases, she has sufficient reason to keep her promises to other friends. And this reveals that her recognition of reasons is not haphazard; instead it is understandable and patterned as is required by regular reasons-receptivity.

Second, it is also plausible to think that Marcia's deliberations are weakly reasons-reactive. As we pointed out earlier, this is an especially low bar, so it should come as no surprise that Marcia clears it, even in her depressed state. To meet this

¹¹ As we will see in the remainder of the paper, this will be especially relevant to the degree of Marcia's responsibility, even though, as we argue here, it does not undermine her ability to meet the minimal threshold conditions on moral responsibility.

¹² Describing his own experiences with depression, Andrew Solomon wrote:

I ran home shaking and went to bed, but I did not sleep, and *could not get up* the following day. I wanted to call people to cancel birthday plans, but I *could not*. ... I knew that for years I had taken a shower every day. Hoping that someone else could open the bathroom door, I would, with all the force in my body, sit up; turn and put my feet on the floor, and then feel so incapacitated and frightened that I would roll over and lie face down. I would cry again, weeping because the fact that I *could not do it* seemed so idiotic to me, [Solomon (1998, pp. 46–49); emphasis from Watson (2004, p. 93)].

The fact that Solomon is so upset over his inability to do the things he needs to do to take a shower suggests that he is able to see himself as having reasons for taking a shower. After all, when we cannot do something that we see ourselves as having no reasons to do, we rarely take our inability to be "idiotic." Thus, it's plausible to suppose that Marcia would have been able to *recognize* her reason to pick you up (even if she, like Solomon, had real trouble acting on that reason).

condition, it simply has to be the case that holding fixed the mechanism that actually issues in her action, there is at least one possible world in which, given sufficient reason to do otherwise, the mechanism reacts accordingly. Here we must hold fixed the actual-sequence mechanism: Marcia's depressed practical reasoning. If her depressed practical reasoning is going to satisfy the weak reasons-reactivity requirement, it must react to a sufficient reason to do otherwise in at least one possible world. And we think it is plausible that this is true of Marcia's depressed practical reasoning.

In the actual world, Marcia's depressed practical reasoning issues in her decision to stay in her apartment by herself even though she recognizes that there is sufficient reason to do otherwise. However, suppose that Marcia received a call (from a credible source) telling her that if she left her apartment for an hour she would win \$1,000,000. Surely this is a sufficient reason to leave the apartment! And plausibly, even though she is depressed and this depression significantly affects her deliberative mechanism, Marcia would leave her apartment in this scenario. Alternatively, suppose that Marcia smells smoke and recognizes that her apartment is on fire. Again, she would have sufficient reason to do refrain from sticking around the apartment. And again, it is plausible to suppose that Marcia would get up on the basis of the very same mechanism that fails to react to a sufficient reason to do otherwise (i.e., her promise to you) in the actual world. Thus, the mechanism in question is weakly reasons-reactive even though it failed to react to sufficient reason to do otherwise in the actual-sequence.

This suggests that Marcia's practical reasoning, though significantly affected by her depression, is a moderately reasons-responsive mechanism. And because her failure to pick you up issued from this mechanism, she satisfies this condition on moral responsibility.¹³ But if Marcia satisfies the minimal threshold conditions on moral responsibility, how are we to differentiate between Marcia and Thomas, who seems to be *more* responsible, and subsequently, *more* blameworthy for his failure to pick you up?

5

For ease of exposition, let us simply assume that like Marcia, Thomas acts on a regularly reasons-receptive and a weakly reasons-reactive mechanism. Thus, let us simply assume that Thomas also acts on a moderately reasons-responsive mechanism and thereby satisfies that condition on moral responsibility. Granting these claims, we think there are still some significant differences between Marcia and Thomas that will rationalize the intuitive judgment that Marcia is less blameworthy for her broken promise than is Thomas.

Recall that Marcia and Thomas satisfy the weak reasons-reactivity condition in virtue of the fact that (holding fixed the mechanism that actually issues in their

¹³ Although this case is indeterminate with respect to the other conditions on moral responsibility, it would be possible to extend Marcia's story in a coherent way that would allow for her meet all the minimal threshold conditions on moral responsibility even though she suffers from depression.

actions) there is at least one possible world in which given sufficient reason to do otherwise, the mechanism reacts accordingly. Because both Marcia and Thomas would have left their apartments to claim a \$1,000,000 prize or to avoid burning to death, they each minimally satisfy this requirement.

But notice, the possible worlds in which someone calls offering large sums of money or apartment buildings burn down are not particularly close to the actual world. What we mean is this: there are a number of *significant* differences between those possible worlds and the actual world in which Marcia's and Thomas's practical reasoning issues in their decisions to ignore their promises to you. And these differences affect the overall *comparative similarity* of the worlds in which the phone call comes or the fire starts and the actual world in which Marcia and Thomas decide to forgo the trouble of getting off the couch to come pick you up.

By itself, this might not seem significant, but we think that ultimately, considerations of comparative similarity underwrite Marcia's and Thomas's differential degrees of responsibility and blameworthiness. Roughly, our suggestion is as follows: the greater degree of comparative similarity that obtains between the actual world and the *nearest* possible world in which the actual sequence mechanism reacts to sufficient reason to do otherwise, the greater degree of responsibility. And correspondingly, the greater the degree of blameworthiness. Likewise, if there is less comparative similarity between the actual world and the nearest possible world in which the actual sequence mechanism reacts to sufficient reason to do otherwise, the agent in question is less responsible. And correspondingly, agents whose actions issue from such mechanism will be less blameworthy for their actions.

Our contention is that in this case, the nearest world in which given sufficient reason to do otherwise, Marcia's actual-sequence mechanism reacts accordingly is *farther* or *more distant* from the actual world than is the nearest world in which given sufficient reason to do otherwise, Thomas's actual-sequence mechanism reacts accordingly. And this fact, we claim rationalizes our judgment that Marcia is less responsible and thereby less blameworthy for her action than is Thomas. But the notions of "nearness" and "farness" and of "comparative similarity" between worlds are ticklish ones, so before we can defend this claim, we turn to a brief discussion of these relations.

6

On standard analyses of counterfactuals, the truth of some counterfactual—say, if Rick Perry was elected President, Sarah would leave the country—depends on whether Sarah leaves the country in the nearest world in which Rick Perry is elected President.¹⁴ That is, whether a counterfactual is true depends on whether the consequent is true in the nearest world in which the antecedent obtains. Our thought is that the same notion of comparative similarity at stake in a counterfactual semantic is also relevant to the degree to which an agent is reasons-reactive.

¹⁴ Obviously, our discussion of counterfactuals borrows heavily from David Lewis's 1973 influential account.

Of course, this is hardly an analysis of the notion of comparative similarity, but it can help, since we can avail ourselves of the resources developed for counterfactual semantics.

Roughly (very roughly!), to say that worlds W_1 and W_2 are close in the sense of comparative similarity is to say that the facts that distinguish W_1 and W_2 are minor, trivial, or few.¹⁵ That is, suppose that the only difference between W_1 and W_2 is that in W_2 there is one more blade of grass growing in Elizabeth's front lawn than is the case in W_1 . If this is the only difference between these worlds, then they are comparatively similar and quite close indeed.

After all, suppose that in W_1 Jeremy said, "Hey Elizabeth, your lawn is perfect. If there was even one more blade of grass, I'd say the grass was too thick. But as it stands, I can not believe how good it looks!" On the standard analysis of counterfactuals sketched above, we can see that what Jeremy says is strictly speaking false. The nearest world in which the antecedent obtains is plausibly W_2 . After all, W_1 and W_2 differ only with respect to the number of blades of grass growing in Elizabeth's front lawn. But in W_2 it is plausible to think that Jeremy thinks and subsequently asserts that Elizabeth's lawn is perfectly luscious, even though it contains one more blade of grass than it does in W_1 .

Consequently, we can see that the closeness (i.e., the comparative similarity) of W_1 and W_2 is grounded in the relatively minor differences in these worlds. That is, in asserting that W_1 and W_2 are nearby possible worlds, we aren't asserting that they are "close" in a geographical sense. Rather, we are asserting that the set of coherent propositions that are true at these distinct possible worlds show little variance. So suppose that at W_1 the propositions $p_1, p_2, \dots, p_{3,974}, \dots, p_n$ are true, and that " $p_{3,974}$ " refers to the number of blades of grass in Elizabeth's front lawn at the time of Jeremy's utterance. We claim that W_2 is close to W_1 because the set of propositions p_1, p_2, \dots, p_n that are true at W_2 is, with the exception of $p_{3,974}$ (and perhaps a few related propositions that explain the truth of $p_{3,974}$ in W_1), identical to the set of propositions that are true at W_1 .

Of course, it is not just the number of propositions that two worlds share that determine their comparative similarity. For example, in ordinary contexts, the number of blades of grass growing in Elizabeth's lawn is insignificant. This explains why worlds such as W_1 and W_2 that only differ with respect to one the truth of one proposition (e.g., the number of blades of grass growing in Elizabeth's lawn is X) can be extremely close. However, if blades of grass mattered significantly in ordinary contexts, we might conclude that this difference was relevant to assessing the comparative similarity between W_1 and W_2 . For example, a world in which Dan climbs to the top of El Capitan without falling might share a great many true propositions with a world in which he slips and falls, but because of a local violation in the law of gravity is lifted to the top. But plausibly, the world in which Dan climbs to the top of El Capitan in a perfectly non-miraculous way is closer to a world in which he falls and unfortunately plummets to his death.¹⁶ After all, in

¹⁵ Admittedly, what counts as minor or trivial is contextually dependent, but given the aims of this paper we can ignore this further complication here.

¹⁶ This example comes from Alvin Plantinga (1994).

ordinary contexts, we hold fixed the laws of gravity, and being required to temporarily lift them to accommodate Dan's continued living, results in less comparative similarity overall.

This suggests that comparative similarity as a relation that obtains between worlds must be evaluated holistically; judgments of comparative similarity cannot rely on a mechanical assessment of shared true propositions. Yet despite the difficulty associated with nailing down this notion, we think that it is an intuitive one, and as we'll argue, it is relevant to the degree to which agents are morally responsible, and thereby relevant to the degrees to which agents are blameworthy.¹⁷

7

We are now in a position to explain why Marcia is less responsible and less blameworthy for ignoring her promise to you than is Thomas. A natural thought here is that Marcia is less responsible because it is more difficult (in some sense of "difficult") for her to be motivated to get off the couch to come pick you up than it is for Thomas to be motivated to get off the couch to come pick you up. But what sense of "difficult" is at stake in this explanation? It can not simply be that it is more *causally* difficult because it is consistent with the cases that the causal pressure on Marcia and Thomas is the same. So there must be another way to make sense of the intuition that because it is harder for Marcia to keep her promise, she's less responsible and less blameworthy than is Thomas.

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. W_1 is more accessible from W_2 than it is from some very distant world $W_{94,305}$ *because* there are more relevant similarities between W_1 and W_2 than there are between W_1 and $W_{94,305}$. Accordingly, we claim that the degree to which Marcia's actual-sequence

¹⁷ Despite the fact that we cannot give a complete account of comparative similarity we do want to highlight one important principle that governs judgments of comparative similarity in the context of moral responsibility. One might worry that cases involving Frankfurt-style counterfactual interveners make trouble for our suggestion that the nearness of the world where the agent reacts to a reason to do otherwise is relevant to the agents responsibility. Following a case developed by Harry Frankfurt (1969), suppose that agents A and B both commit a murder as the result of identical causal chains. In the case of A (but not B) let us further suppose that there was an evil neuroscientist prepared to causally intervene, should A show any sign that she might decide to not go through with the murder, and ensure that A commits the murder. However, as it happens, A carries out the murder without the need for any prompting from the neuroscientist.

Two things seem plausible here. First, A and B are equally blameworthy for the murders they commit and second, the world where A reacts to a reason to not commit the murder is much further away than the world where B reacts to such a reason. How can we account for these two claims? In our view cases like this show that what matters for comparative similarity in the context of moral responsibility is the similarity of the actual sequence causal chains leading to the action. Facts about portions of worlds that are causally isolated from the agents act (such as the presence of a merely counterfactual intervener) do not matter in this context. The nearest causal sequences in which A and B react to a reason to refrain from the murder are equally similar to the actual causal chain. This accounts for their equal blameworthiness.

mechanism is reasons-reactive depends on how difficult it would be for that mechanism to react to a sufficient reason to do otherwise. And we reductively analyze the notion of “difficulty” in terms of comparative similarity. On this view, it is difficult for Marcia’s actual-sequence mechanism to react to sufficient reason to do otherwise because the nearest 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.

After all, compare a world in which given sufficient reason to do otherwise, Marcia’s actual-sequence mechanism reacts accordingly to the actual world in which Marcia ignores her promise to pick you up. Plausibly, given her depression, there would have to be significant differences between these worlds—either a fire, or perhaps more optimistically, a call offering a lot of money. Because of her depression, not only does Marcia not feel like picking you at the appointed time, she would not feel like picking you up across a wide range of counterfactual scenarios. Moreover, she would not feel like doing *anything* across a wide range of counterfactual scenarios. Thus, the worlds in which she does feel like doing something and in which the actual-sequence mechanism suitably responds to sufficient reason, are likely to be distant from the actual worlds. Of course, there are such worlds—that is why the mechanism that produces her action is weakly reasons-reactive. But the nearest worlds that underwrite her weak reasons-reactivity are quite distant. And that is why she’s less reasons-reactive (and therefore less responsible and less blameworthy) than an agent not suffering from depression.

Notice, the nearest worlds in which Thomas’s actual-sequence mechanism reacts to a sufficient reason need not be distant. Very minor changes in how lazy he was feeling or the television programing would’ve been sufficient for the actual-sequence mechanism to issue in an action that accords with the weight of reasons. Thus, plausibly, there is a great deal of comparative similarity between the actual world and the nearest world in which Thomas’s actual-sequence mechanism reacts to sufficient reason. And as a result, he’s more reasons-reactive than someone like Marcia.

In short, the judgment that Thomas is more responsible and more blameworthy for failing to pick you up than is Marcia can be rationalized by appealing to how hard it would have been for them to do otherwise if they were given sufficient reason to do otherwise. And we analyzed the notion of “hardness” or “difficulty” in terms of the comparative similarity of worlds, the result being that Thomas is more responsible than Marcia because given sufficient reason to do otherwise, the nearest world in which the actual-sequence mechanism reacts accordingly is closer to the actual world. The worlds in question are close for Thomas but distant for Marcia. Thus, Thomas is more responsible and more blameworthy than is Marcia.

Now, we will argue that agent’s degree of reasons-receptivity is also relevant to her degree of responsibility. In a natural way of spelling out Marcia’s case, her depression will limit her receptivity to reasons as well as her reactivity to reasons. After all, when an agent suffers from depression, the reasons that are salient to her are restricted or altered in other ways. For example, the joy of walks on sunny days

no longer sparks her imagination in ways that lead her to see the good weather as a reason for her to take a walk. Thus, on our account there will be two ways in which Marcia differs from Thomas that reduce her degree of responsibility. Despite this, we have focused on Marcia's reactivity to reasons on order to bring out how this affects an agent's degree of responsibility. We will now go on to examine the connection between reasons-receptivity and degrees of responsibility.

8

Beginning with Fischer and Ravizza's account of moderate reasons-responsiveness as a necessary condition on moral responsibility, we argued that by extending such an account—especially by focusing on the weak reasons-reactivity requirement on moderate reasons-responsiveness—it is possible to explain why Marcia, a friend suffering from depression, is less responsible for breaking her promise to you than is Thomas, a friend “suffering” from selfishness. Comparing the nearest worlds in which it is possible to ground the actual-sequence mechanism's weak reasons-reactivity reveals that Thomas satisfies the weak reasons-reactivity requirement in virtue of a closer world than does Marcia. Thus, he is more responsible and more blameworthy.

However not all cases of differential degrees of responsibility and blameworthiness can be explained in this way. Consider Alexandra and Davis. Alexandra and Davis have each just been arrested for robbing (different) convenience stores. Let us suppose that they each satisfy the minimal threshold conditions on moral responsibility, i.e., that they are moderately reasons-responsive (and therefore, that they are regularly reasons-receptive and weakly reasons-reactive). On the face of it, it might seem that they responsible and blameworthy to the same degree. However suppose that Alexandra is 14, but Davis is 37 (and further suppose that Alexandra and Davis are healthy and normally developed for their ages). Upon learning this fact, we think it is natural to conclude that Alexandra is less responsible than is Davis. However, it is not obvious that we can explain this fact by appealing to differential degrees of reasons-reactivity. After all, had they known that there were security cameras in the convenience stores (a fact that serves as a sufficient reason to refrain from robbing a store) they each would have refrained from robbing the stores. But it is a nearby world in which they know this fact—say, a world in which they are slightly more attentive, or in which the cameras are more prominently displayed, or in which there were signs on the door announcing the presence of cameras. Thus, on our view, not only do they satisfy the condition of weak reasons-reactivity, they satisfy it in virtue of nearby possible worlds.

So appealing to differential degrees of reasons-reactivity won't rationalize the judgment that Alexandra is less responsible and less blameworthy for her crime than is Davis. In that sense, the account we develop in §§ 5, 6, and 7 is not a general account of differential degrees of responsibility. However, by extending Fischer and Ravizza's account of moderate reasons-responsiveness in a second way, we can explain why Alexandra is plausibly thought to be less responsible for her crime than is Davis. While they are equally reasons-reactive, it is doubtful that Alexandra is

reasons-receptive to the same degree that Davis is reasons-receptive. These varying degrees of reasons-receptivity, we claim, can explain why Alexandra is less responsible (and thereby less blameworthy) for robbing the store than is Davis. So if it is plausible to think that a teenager such as Alexandra is less receptive to reasons, or alternatively, is receptive to reasons in a less understandable or less “regular” fashion than Davis, then it is plausible to think that we can explain their differential degrees of responsibility and blameworthiness by appealing to this fact.

9

Unlike our extension of Fischer and Ravizza’s account of weak reasons-reactivity, our extension of their account of reasons-receptivity is quite simple. Recall that moderate reasons-responsiveness requires that the actual-sequence mechanism be regularly reasons-receptive. Of course, this is simply a minimal threshold condition on moral responsibility. And intuitively, even regularly reasons-receptive mechanisms can be more or less regularly reasons-receptive because the intelligible patterns that ground the regularity of a mechanism’s receptivity can be more or less understandable from a third-person point of view.

Higher (or lower) degrees of regular reasons-receptivity are possible in two ways. First, a mechanism could be sensitive to a wider (or narrower) range of reasons. For example, Tara is a bit of a wine snob, but deservedly so. Unlike most of us, she doesn’t just pretend to know about tannins and terroir; she’s actually sensitive to these considerations in her wine selection. So it is plausible to think that Tara is *more* reasons-receptive, at least within the domain of wine. Alternatively, compared to most of us, Jared might seem uncouth. But this is simply because he did not grow up in a home that emphasized the kind of manners and etiquette that most of us take for granted. As a result, he is not particularly sensitive to reasons of etiquette, even though he is, in general, regularly reasons-receptive. Second, a mechanism could process reasons in a more or less understandable fashion. If Larry has little disposal income, it is understandable that he sees Super Bowl tickets costing more than \$1,000 as too pricey. But it would be less understandable if he did not also see Super Bowl tickets costing more than \$500 as too pricey given his lack of discretionary funds. Of course, he would not be obviously irrational in the way that someone who saw \$1,000 tickets as too pricey but couldn’t similarly see \$1,001 tickets as too pricey is irrational. But his receptivity to reasons would be less intelligible than that of someone who, given her lack of funds, saw \$500 tickets, as well as \$1,000 tickets, as too pricey.

Plausibly, when compared to normal adults, adolescents such as Alexandra, while meeting the regular reasons-receptivity condition on responsibility, act on mechanisms that are, by and large, less sensitive to the full range of reasons and worse at processing reasons in a fully intelligible fashion. After all, because children and young teenagers are less capable of appreciating the normative significance of moral reasons *as such*, they are less principled in their application and appropriation of such reasons. This is not too surprising given the great deal of subtlety that is required to make the right decision in a wide range of moral scenarios—a level of

subtlety that only comes from life experience. Without such experience, young teenagers are unable to appreciate the significance of moral reasons to the same degree they will be able to when they reach adulthood.

These facts, we claim, can rationalize the judgment that Alexandra is less responsible and less blameworthy for her crime than is Davis. Unlike Davis, Alexandra's agency—while sufficient for satisfying the minimal threshold conditions on moral responsibility—is not fully developed. Such development comes with time and life experience. And specifically, time and life experience will afford Alexandra a greater sensitivity to reasons as well as a greater capacity for processing those reasons in an understandable way. In other words, time and life experience will transform Alexandra's already regularly reasons-receptive mechanisms into more reasons-receptive mechanisms.

As with Marcia, it is plausible to think that *both* Alexandra's reasons-receptivity and her reasons-reactivity are limited. In addition to reduced reasons-receptivity, adolescents tend to have more difficulty than adults in getting themselves to act on the reasons they do recognize. Notably, teenagers are known to have trouble with impulse control, which is plausibly understood to be a failure of reasons-reactivity. Thus, there are two reasons Alexandra is less responsible than Davis. So although we have focused on Alexandra's reasons-receptivity here in order to highlight the connection between degrees of reasons-receptivity and degrees of responsibility, we admit that a full account of why Alexandra is less responsible and blameworthy than Davis would require a discussion of reasons-reactivity.

10

Before concluding, we now wish to pause in order to consider a possible objection to the account we have developed. One might take issue with our claim that those who are less receptive to reasons are thereby less responsible for their bad behavior. For example, racists and bigots might be thought to be less reasons-receptive, since, at least in some cases, the racist or bigot seems unaware that they have genuine moral reasons reasons to treat everyone equally. However, it is also natural to think that it is precisely the racists and the bigots who are apt targets of resentment and indignation.¹⁸ Does our account (implausibly according to some) yield the result that such a person bears only a small amount of responsibility for her racist or bigoted behavior?

In order to address this case we need to introduce a further consideration that can affect an agent's degree of responsibility. We have said that an agent's degree of responsibility partially depends on the degree to which she is currently reasons-receptive. In addition we find it plausible that her degree of responsibility is often sensitive to her past degree of reasons-receptivity (and reasons-reactivity). Suppose that our racist currently has diminished reasons-receptivity with regard to reasons to

¹⁸ Thanks to an anonymous referee for bringing this case to our attention. To some extent, our statement of the worry tracks a problem that Gary Watson (1987) raises for Strawsonian theories of moral responsibility.

respect people of other races because in the past she has repeatedly decided to think hateful thoughts about people of other races. Not surprisingly, these hateful thoughts gradually inhibit her ability to recognize a certain set of moral reasons. Furthermore, suppose that when making those prior decisions she possessed a high degree of reasons-receptivity. In this case we do not think that her responsibility is substantially reduced. This is because she is herself morally responsible for her diminished reasons-receptivity.¹⁹ This suggests that the degree to which an agent is responsible is also a historical question. When an agent's past choices lead, "in the right sort of way," to current behavior, this can increase her degree of responsibility.²⁰

Now of course, not all cases of racism can be handled this way. And if our racist's lack of reasons-receptivity is explained in some other way, then we may have to grant that the racist's responsibility is mitigated to some extent. Of course, we don't have to thereby conclude that the racist or the bigot is not responsible or blameworthy. Thus, rather than a cost of our view, we take this to be the correct judgment to make in such cases. For example, it is plausible to think that a racist living in the antebellum South (call him Robert), while still morally responsible and blameworthy for his racist attitudes, is nevertheless less responsible and blameworthy for his racist attitudes than a contemporary racist is (call him David). And a plausible explanation for this would be that Robert is less reasons-receptive—i.e., he is less aware of liberal reasons for treating persons of all races as equals.²¹ Indeed, contingent facts about our environments, cultures, and upbringings can dramatically affect our reasons-receptivity dramatically affecting what we take to be reasons. Of course, we also suspect that even antebellum racists like Robert were sufficiently receptive to moral reasons, and so we suspect that they deserve (or would have deserved) our opprobrium. It is just that we think that Robert is *less* responsible than are contemporary racists like David, and thus we are willing to accept the conclusion that the prejudiced and bigoted are sometimes less responsible due to their decreased reasons-receptivity.

Of course, what we say here is not peculiar to racists and bigots. Instead, it points to an important feature of moral responsibility: that it is a historical phenomenon.²² Historical considerations—e.g., the cultural and environmental contingencies that shape our status as reasons-receptive—can affect the degree to which I am reasons-receptive. Thus, an agent's history can affect the degree to which she is morally responsible.

¹⁹ Plausibly, just as you cannot legitimately avoid an obligation by tying yourself to a chair, thereby making it impossible to fulfill the obligation, neither can you mitigate your responsibility by deciding to diminish your reasons-receptivity in the way described above.

²⁰ Unfortunately we do not have an account of what "the right sort of way" amounts to. However, it is often taken to include a foreseeability requirement. And we think something like this is on the right track.

²¹ Obviously, we are not claiming that Robert is not responsible or blameworthy for his racist attitudes—he certainly is. But it does seem as if, given his cultural context, his repugnant actions display less ill will than do those of David. And the best explanation for this difference in the quality of their wills is simply that Robert is less sensitive—i.e., less reasons-receptive—to reasons issuing from a demand of mutual regard that extends not only to caucasians, but to persons of all races.

²² Indeed, Fischer and Ravizza (1998) make this exact point.

11

In the preceding sections, we have extended Fischer and Ravizza's account of moderate reasons-responsiveness in two ways. First, we claimed (in §§ 5, 6, 7) that it is possible to explain some cases of differential degrees of responsibility and blameworthiness by appealing to features of the weakly reasons-reactive actual-sequence mechanisms. This explanation, we argue, partially explains why agents suffering from depression are less responsible and less blameworthy for their actions than are agents who do not suffer from such maladies. Second, we claimed (in §§ 8, 9, 10) that it is possible to explain other cases of differential degrees of responsibility and blameworthiness by appealing to features of the regularly reasons-receptive actual-sequence mechanisms. This explanation will partially account for why children and adolescents are less responsible and less blameworthy for their actions than are normally functioning adult agents. (Again, we do recognize that both reduced reasons-receptivity and reduced reasons-reactivity are likely to be present in both adolescents and the depressed. The full explanation of their reduced responsibility will appeal to both of these factors.)

Of course, we have only considered two cases of mitigated responsibility and blame: depression and adolescence. And obviously, the fact that our extension of Fischer and Ravizza's account of moderate reasons-responsiveness is able to accommodate these two cases is not a *decisive* reason to accept our claims concerning the connection between comparative similarity and reasons-reactivity or our attempt to account for more or less reasons-receptive mechanisms. However, we think that proof is in the pudding. That is, we think that our account has a significant explanatory and theoretical payoff for a wide range of cases of impaired or undeveloped agency—cases in which intuitively, there is mitigated responsibility and blame. We explore such cases elsewhere,²³ but in the present context, we simply offer our account of the differential degrees of responsibility and blameworthiness as a plausible rationalization of the judgments that Marcia and Alexandra deserve less blame than Thomas and Davis.

Acknowledgements For helpful comments and suggestions on earlier drafts of this paper we would like to thank John Martin Fischer, Christopher Franklin, Ben Mitchell-Yellin, Jonah Nagashima, Michael Nelson, Garrett Pendergraft, John Perry, Patrick Todd, Neal Tognazzini, and a very helpful anonymous referee.

References

- Fischer, J. M., & Ravizza, M. (1998). *Responsibility and control*. Cambridge: Cambridge University Press.
- Frankfurt, H. (1969). Alternate possibilities and moral responsibility. *The Journal of Philosophy*, 66(23), 829–839.
- Kelly, E. (forthcoming). What is an excuse? In D.J. Coates & N. Tognazzini (Eds.), *Blame: Its nature and norms*. Oxford: Oxford University Press.

²³ We apply our account to the case of psychopaths in “Mitigated Blame and Marginal Agency,” [MS], and we explore implications of our account of differential degrees of reasons-reactivity for the manipulation argument against compatibilism in “Reasons-Responsiveness and Manipulation,” [MS].

- Lewis, D. (1973). *Counterfactuals*. Boston: Wiley-Blackwell.
- McKenna, M. (2005). Reasons reactivity and incompatibilist intuitions. *Philosophical Explorations*, 8(2), 131–143.
- Nelkin, D. (2011). *Making sense of freedom and responsibility*. Oxford: Oxford University Press.
- Plantinga, A. (1974). *The nature of necessity*. Oxford: Oxford University Press.
- Solomon, A. (1998). Anatomy of melancholy. *The New Yorker*, 73, 46–58.
- Todd, P., & Tognazzini, N. (2008). A problem for guidance control. *The Philosophical Quarterly*, 58, 685–692.
- Wallace, R. J. (1994). *Responsibility and the moral sentiments*. London: Harvard University Press.
- Watson, G. (1987). Responsibility and the limits of evil: Variations on a Strawsonian theme. In F. Schoeman (Ed.), *Responsibility, character, and the emotions: New essays in moral psychology* (pp. 256–286). Cambridge: Cambridge University Press.
- Watson, G. (2004). *Volitional necessities reprinted in agency and answerability* (pp. 88–122). Oxford: Oxford University Press.
- Wolf, S. (1990). *Freedom within reason*. Oxford: Oxford University Press.