

# RESPONSE-DEPENDENCE IN MORAL RESPONSIBILITY: A GRANULARITY CHALLENGE

Shawn Tinghao Wang

## ABSTRACT

According to the response-dependence view of moral responsibility, a person is morally responsible just in case, and in virtue of the fact that, she is an appropriate target for reactive attitudes. This paper raises a new puzzle regarding response-dependence: there is a mismatch between the granularity of the reactive attitudes and of responsibility facts. Whereas the reactive attitudes are comparatively coarse-grained, responsibility facts can be quite fine-grained. This poses a challenge for response-dependence, which seeks to ground facts about responsibility in facts about the reactive attitudes. Specifically, reactive attitudes are not enough for grounding facts about degrees of moral responsibility. The response-dependence view thus requires significant revisions or supplementations.

## I. INTRODUCTION

The following biconditional is widely accepted: a person is morally responsible for *X* if and only if it is appropriate to hold reactive attitudes toward her with respect to *X*. However, when it comes to the underlying metaphysical relations, there is a deep disagreement. Response-independence theorists think a person is an appropriate target for reactive attitudes in virtue of being morally responsible; by contrast, response-dependence theorists think a person is morally responsible in virtue of being an appropriate target for reactive attitudes. Recently, the debate has received much attention. Those who are sympathetic to response-dependence suggest that it offers a simple and unified account of the conditions under which people are responsible (Shoemaker 2017; Carlsson 2017; also see Todd 2016), it characterizes responsibility in ways analogous to concepts like humor (Shoemaker 2017), it

explains the centrality of social community to moral responsibility (Watson 2014; Menges 2017), and it promises a better answer than its alternatives to the question of how people can be responsible even if metaphysical determinism is true (see, e.g., Strawson 1962; Watson 1987; Beglin 2018).<sup>1</sup> Response-independence theorists, by contrast, are resistant to granting reactive attitudes a fundamental status, suggesting that they instead gain their normative significance by tracking independent facts about people's being responsible (e.g., Fischer & Ravizza 1998; Smith 2007; Brink & Nelkin 2013; Zimmerman 2010).

However, both sides in the debate have focused almost exclusively on the explanatory relation between reactive attitudes and the binary notion of *being morally responsible*. But it is also intuitive that a person can *be morally responsible to a greater or lesser degree*. Given this, an approach like response-dependence ought to be subject to a

test of gradability: can response-dependence be applied to this scalar notion of moral responsibility? In particular, do reactive attitudes have the level of granularity to ground facts about degrees of moral responsibility? These questions have not been addressed in the literature.

This paper aims to apply such a gradability test to response-dependence in responsibility. I argue that doing so raises a new puzzle for the approach: even if response-dependence is attractive at a comparatively coarse-grained level, it turns out to be too coarse-grained to ground facts about *degrees* of moral responsibility. Though there are several available strategies one may attempt to address the worry, they all come with substantial costs and are difficult to defend. In general, the suggestion is that we should reevaluate, and properly restrict, the metaphysical role of reactive attitudes in theorizing of moral responsibility.

My discussion will focus primarily on *negative* moral responsibility, that is, moral responsibility for things that are morally wrong (or morally bad), rather than in cases of *positive* moral responsibility, that is, moral responsibility for things that are morally right (or morally good). In particular, I will remain agnostic with regard to whether there is a similar granularity challenge against the connection between reactive attitudes and positive moral responsibility.

## 2. RESPONSE-DEPENDENCE IN MORAL RESPONSIBILITY

More accurately, the *response-dependence view* can be put as follows:

(I) A person is morally responsible for something morally wrong (or morally bad) *if and only if, and in virtue of the fact that*, it is appropriate to hold some negative reactive attitude toward this person with respect to that thing.<sup>2</sup>

The *response-independence view*, by contrast, states that the reactive attitude is

appropriate if and only if, and in virtue of the fact that, the person is morally responsible. Both sides in the debate agree that the propriety condition for reactive attitudes is necessary and sufficient for ascribing responsibility; their disagreement concerns whether the order of explanation is left-to-right or right-to-left.

Response-dependence theorists disagree on what the relevant negative reactive attitudes are. The most common candidates are resentment, indignation, and guilt. But Shoemaker (2017), a defender of response-dependence, suggests that it is anger rather than resentment that figures in (I). His main worry is that “resentment” has been sometimes used in a way that implies a cognitive element of the attitude: a judgment that someone is morally responsible for a wrongdoing. But this cognitive construal of resentment almost immediately contradicts response-dependence. After all, the very idea of response-dependence is to ground facts about being responsible in facts about reactive attitudes. If reactive attitudes themselves are to be explained by their propositional contents, which further involve facts about being responsible, then the explanation from reactive attitudes would be ultimately circular—which is not a desirable feature for theories of responsibility.<sup>3</sup> Still, as Shoemaker acknowledges, we can also understand resentment in a non-cognitive manner such that it does not necessarily involve any particular cognitive element (2017, p. 494). I will proceed to use “resentment” in this latter way, so that treating resentment as the paradigm attitude in (I) does not beg the question against response-dependence. I think this ultimately captures the same emotion as intended by Shoemaker’s term “anger.”<sup>4</sup>

Another issue about (I) concerns what it means for a reactive attitude to be *appropriate*. Important for the response-dependence theorist, the reactive attitudes are appropriate not because people are morally responsible. Rather, what grounds the propriety condition

is to be found in the very facts regarding, or *internal* to the practices involving, how we hold reactive attitudes toward each other. The grounding relation can take a variety of different paths. However, if response-dependence were to have any initial appeal, it cannot characterize the propriety conditions in very local terms, for example, according to whether the wrongdoer is actually held to account. It must allow for local fallibility, that is, cases in which a person is morally responsible even though no one actually holds reactive attitudes toward her. This means that the response-dependence theorist should appeal to certain global rather than local features in our responsibility practices:

(II) It is appropriate to hold some reactive attitude toward a person for something *if and only if, and in virtue of the fact that*, doing so would accord with certain global features in our practices of holding reactive attitudes toward people.

This characterization is general enough to leave open what exactly the relevant global features are. One way to proceed is to examine the simple *dispositional* features, such that the reactive attitude is appropriate if and only if we are disposed to have the reactive attitude under some *standard condition*. The “standard condition” can be either just a condition that we normally satisfy (that is, without dispositional masks, finks, etc.) or put in more idealized terms. For example, Lewis takes it to be a condition in which a person gains “the fullest imaginative acquaintance that is humanly possible” (1989, p. 121); Railton takes it to be a state in which a person can “contemplate his present situation from a standpoint fully and vividly informed about himself and his circumstances, and entirely free of cognitive error or lapses of instrumental rationality” (1986, p. 16). Both Lewis and Railton defend response-dependence theories of value in general, rather than of responsibility in particular; however, one can easily apply what they say about the standard

or ideal condition for dispositions to cases involving moral responsibility. The other way to proceed is to examine whether the reactive attitude would be *fitting*. Shoemaker (2017) adopts this approach. Fittingness is about correct representations. Just as fear is fitting when it tracks a dangerous object, according to Shoemaker, moral anger is fitting when it tracks a property that triggers our “anger sensibilities” and belongs to “the sort of properties to which we humans are built to respond with a heated demand for acknowledgment or a tendency to retaliate” (Shoemaker, 2017, p. 510). In contrast to the simple dispositional approach, Shoemaker’s theory finds the relevant features in the dispositions of those refined and developed individuals who have a good anger sensibility (2017, p. 511).<sup>5</sup>

(I) and (II) combined together entail the following thesis:

(III) A person is morally responsible for something morally wrong (or morally bad) *if and only if, and in virtue of the fact that*, holding some negative reactive attitude toward this person with respect to that thing would accord with certain global features in our practices of holding reactive attitudes toward people.<sup>6</sup>

(III), I think, is broad enough to incorporate the various strategies that one can adopt to articulate the thesis of response-dependence.

### 3. THE GRANULARITY OF REACTIVE ATTITUDES

Response-dependence is typically presented, defended, and criticized as being concerned with a binary notion of moral responsibility, as the thesis that whether one *is morally responsible or not* is to be explained in terms of whether one is an appropriate target for negative reactive attitudes such as resentment, indignation, guilt, or anger. But it is intuitive that moral responsibility comes into degrees: one can be *more or less morally responsible* for one’s conduct. Given this, it seems that we should also ask whether

response-dependence can be applied to this scalar notion of moral responsibility. However, this connection has not been addressed in the literature.

One may think the application is just a natural extension of (III): as we can ground facts about binary responsibility in the appropriate reactive attitudes, we can ground facts about scalar responsibility in the *appropriate degrees* of reactive attitudes. However, this then raises new problems and challenges regarding the approach. In this section, I contend that reactive attitudes, such as resentment and indignation, are fairly coarse-grained in terms of their degrees; at least, this is so according to how response-dependence theorists characterize the nature of such attitudes. I will contrast this with the granularity of facts about degrees of moral responsibility in the next section.

To start with, what does it even mean to claim that emotions like resentment have degrees? On perhaps the most natural understanding, the claim just means that reactive attitudes have different degrees of *intensity*. There is no denying that emotions have varying intensities. Some differences in intensity implicate slightly different attitudes. For instance, we might distinguish between sadness (mild) and sorrow (strong). There can also be different levels of intensity within a given reactive attitude. For instance, we can feel resentment to different degrees, measured by a scale from mere annoyance or frustration to relatively stronger forms of anger, and finally to rage. However, the scale regarding degrees of emotions is usually very rough in its measure. Take happiness as an example. Suppose that Daniel feels happy about getting a monetary reward for teaching excellence. Depending on how much money the reward involves, Daniel can feel more or less happy—perhaps something like “a little bit happy” around \$5, “somewhat happy” around \$50, and “very happy” around \$500. But I take it that Daniel does not feel happier

when it is \$57 instead of \$56. I stress that this is not just because Daniel cannot intentionally adjust degrees of happiness, or because Daniel does not care about the value of 1 dollar—rather, it is because our emotional states are carved out in human psychology in a way such that they do not have the level of granularity required to register only slight differences in values. Now, there is surely a sense in which one can maintain that Daniel does feel happier about \$57 than about \$56, if by claiming so one already implies that Daniel *believes* \$57 is a little bit more valuable. However, putting aside this cognitive element that arguably may come apart with the emotion itself, degrees of emotions usually seem fairly coarse-grained.

The failure to register small differences can be applied to negative reactive attitudes as well. Imagine Heru, a somewhat mentally impaired person, who stole Ann’s passport. Now compare Heru with Henry, who stole Ann’s passport at a different time, but is just *slightly less mentally impaired* than Heru. I take it that Ann does not resent Henry more than Heru, even though Henry is slightly more competent. Again, this is not just because Ann cannot intentionally adjust degrees of resentment. Instead, it is because our emotional states simply do not have the level of granularity required to register only a slight difference in moral competence. That is just the way in which emotions are carved out in human psychology. I will refer to the general insight here as the *thesis of limited information registration*: human emotions are set up in such a way that they only track potentially responsibility-relevant factors, such as moral competence, in a fairly coarse-grained manner. More specifically, one way of measuring Heru and Henry’s relevant mental/moral competence is to look at the proportion of relevant possible worlds in which they could have acted according to moral reasons and thereby refrained from stealing Ann’s passport or committing similar wrongdoings. Suppose

this percentage is 40 percent for Heru but 41 percent for Henry. It seems unlikely Ann's emotion would change according to this slight increase in the proportion of possible worlds. I stress that the point here does not hinge on the possible-world account of moral competence. It would go through on other accounts as well, as long as moral competence is characterized in a fairly fine-grained manner. The general thought is that there are many cases in which, just like Heru and Henry, our emotions do not have the level of granularity required to differentiate between very similar but slightly different moral agents.<sup>7</sup>

Again, there is surely a sense in which we may say that Ann resents Henry more than Heru, if this already entails, or even just means, that Ann *believes* that Henry is a little bit more morally responsible than Heru for the conduct. But it is unclear whether this sort of cognitive judgment is part of the emotion of resentment. More importantly, response-dependence theorists would certainly deny that it is. As I have suggested in section 2, the relevant reactive attitudes that figure in the response-dependence approach should not involve a particular cognitive element, and surely should not involve any cognitive judgment about moral responsibility, in order to avoid circular explanation. But if we put aside this cognitive element, then the thesis of limited information registration becomes highly plausible.

To further elaborate the thesis, it is useful to consider an analogy to degrees of belief. The orthodox Bayesian view that beliefs have precise credences is frequently criticized as psychologically unrealistic (see, e.g., Mahtani 2020). It is very rare (unless perhaps in certain cases involving precise mathematical questions in particular) for one to actually have 0.64 credence in any proposition. Note that this is still consistent with thinking beliefs are gradational in some sense. Our beliefs do register information about changes in probabilities, but only do so in a fairly

coarse-grained manner. They do not, and do not have the level of granularity required to, register only slight differences in probabilities. My contention is that a similar story is true of degrees of reactive attitudes—they do not have the level of granularity required to register small differences in certain potentially normatively relevant features of moral agents, like their moral competence. It is psychologically unrealistic to posit the existence of fairly precise reactive attitudes.

From the evolutionary perspective, it is hardly surprising that information registration in emotions is limited and selective. Emotions have many important roles to play, but one of them is to trigger immediate actions when receiving certain feedback from the outside world. The evolutionary pressure is likely to associate our emotions with something that matters, but it is unlikely to make the association very accurate; coding every aspect of the external world that matters would be too costly to facilitate the efficient emotional responses that conduce to our survival.

My claim is that emotions are limited and selective in information registration, which is compatible with their registering some information in useful ways. This fits well with the current empirical evidence we have. One common measure of emotional intensity in empirical psychology is by appeal to self-report, and the measure typically adopts a coarse-grained scale, instead of, say, offering a twenty-point scale with low-to-mid-moderate, and the like. For example, Larsen et al. ask the participants “how happy are you?” using a rough scale ranging from 1 to 5, where 1, 3, and 5 are labeled as “slightly,” “moderately,” and “extremely” (2001, p. 688). Further, I think some self-report data constitute indirect evidence for the coarse-grainedness of emotional states. One potential source of evidence comes from Barrett's (2006) work on individual differences in people's ability to describe their emotional experiences in precise terms—an

ability that she calls “emotional granularity.” She distinguishes between “low granularity” and “high granularity” individuals. The low granularity individuals describe their emotional states in very general terms. One example goes like this: “I just felt bad on September 11th, really bad” (Barrett 2006, p. 38). By contrast, the high granularity individuals describe their emotions using more precise labels. One example, also in response to 9/11, goes as follows: “My first reaction was terrible sadness . . . But the second reaction was that of anger, because you can’t do anything with the sadness” (Barrett 2006, p. 38). The lesson here is not only that the low granularity individuals lend support to the coarse-grainedness of emotions. Even the high granularity individuals differentiate negative emotions using relatively general labels such as “terrible sadness” and “anger”; they do not describe their emotional states using more precise degrees. I admit this does not directly show that people’s emotional states are coarse-grained; self-report can be limited in various ways, and people might lack concepts to describe a genuine phenomenon. Still, self-report is considered a major way of measuring emotions, and, therefore, the findings in emotional granularity give some indirect support to the claim that emotions are coarse-grained. As any empirical thesis, my claim here is subject to further scrutiny; however, given the evidence we have thus far, I take it that our default thesis should still be that reactive attitudes are fairly coarse-grained.

#### 4. THE GRANULARITY OF RESPONSIBILITY FACTS

In this section, I will argue that, in contrast to reactive attitudes, facts about moral responsibility have fairly fine-grained degrees. I stress that this does not mean there are always *precise cardinal numbers* assigned to degrees of responsibility; I agree that facts about responsibility are not *that* fine-grained. Rather,

the contention is that they are still fairly fine-grained, and, most important for my purpose, more fine-grained than our reactive attitudes are able to capture.

My argument has two steps: first is to argue that there is a scalar notion of moral responsibility; second is to argue that, if we take for granted that there is a scalar notion of moral responsibility, then there are good reasons to believe moral responsibility comes in fairly fine-grained degrees. The first step, I believe, is quite straightforward since a scalar concept of responsibility is necessary for the purpose of arriving at a proper assessment of one’s degrees of blameworthiness. It is plausible that one is morally blameworthy just in case one is morally responsible for something morally wrong (or morally bad). And surely one can be morally blameworthy to different degrees, since one can deserve more or less blame. Further, it also appears one can deserve more or less blame even if we hold fixed the degree to which one’s action is morally wrong (or morally bad). But this scalarity would be mysterious if moral responsibility is simply a binary notion. In reply, one might point out that there have been important theories of moral responsibility that take it to be an all-or-nothing, threshold notion (e.g., Fischer & Ravizza 1998; Fischer 2004). However, the fact that there is a threshold concept of responsibility does *not* entail that moral responsibility does not come in degrees. Rather, the more plausible view is that there are both a threshold and a scalar concept of moral responsibility, just as in epistemology, where many agree that there are both a threshold notion of belief and a scalar notion of degrees of belief. It is thus unsurprising that even Fischer and Ravizza’s (1998) threshold framework has been later extended to applications involving degrees of responsibility (see, e.g., Coates & Swenson 2013), since the most plausible normative theory should provide both the threshold and the scalar notion.

Now I turn to the second step of the argument: if we take for granted that there is a scalar notion of moral responsibility, then there are good reasons to believe moral responsibility comes in fairly fine-grained degrees—at least more fine-grained than our reactive attitudes are able to capture. The reason for this is that the opposing view, according to which moral responsibility comes in degrees but is no more fine-grained than the gradability of our reactive attitudes, is not attractive, once we consider how such a view is related to substantial conditions of moral responsibility. While some response-dependence theorists, such as Shoemaker, may not believe that any *necessary or sufficient* conditions of responsibility can be offered without appealing to the reactive attitudes themselves, they also do not think it is irrelevant what *substantial conditions* end up being captured by the response-dependence theory. Indeed, I think one important reason why any response-dependence theory is attractive is because it seems to fit pretty well with the *plausible substantial conditions of moral responsibility*. Here is what I mean. It is generally intuitive that certain substantial conditions—again, note that they do not have to be necessary or sufficient conditions—can affect degrees of responsibility, including, for example, the reasons-responsiveness condition and the quality of will condition. The reasons-responsiveness condition specifies whether and to what degree a moral agent or her mechanism has the ability to respond to reasons;<sup>8</sup> and the quality of will condition specifies whether and to what degree a moral agent's action manifests her good or bad quality of will.<sup>9</sup> It is intuitive that such factors can affect degrees of moral responsibility, and it is also true that our emotion, at a coarse-grained level, does vary according to those conditions. After all, reactive attitudes come with *some* degrees. Thus, at a coarse-grained level, we do resent a mature, fully reasons-responsive agent more than a mentally impaired, less

than fully reasons-responsive agent, when other things are equal; similarly, we do resent a person more for an action that manifests a really bad quality of will than for an action that manifests a somewhat bad quality of will, when other things are equal. If response-dependence were to not arrive at those results even at the coarse-grained level, then, I think, the view would not be attractive at all, since it asks for too much conceptual revision. That is, the fact that response-dependence seems to capture pretty well the plausible substantial conditions of moral responsibility is a central attraction for the view. This can also be seen by Shoemaker's defense of response-dependence, where he contends that the approach explains how features like what I have referred to as the plausible substantial conditions of moral responsibility are all "blameworthy-makers" and matter for responsibility in a way consistent with the "apt variations in type and degree of anger" (2017, p. 510).

My contention is that this attraction is, unfortunately, seriously undermined. This is because a closer examination shows that response-dependence *only* gets the intuitive results about the plausible substantial conditions at the coarse-grained level, and once we examine cases about, for example, two agents with only slightly different reasons-responsiveness abilities or slightly different qualities of will, response-dependence no longer amounts to different normative judgments for those agents. My earlier case about Heru and Henry illustrates this for the reasons-responsiveness condition: though Heru is slightly less capable of responding to reasons than Henry is, this is too small a difference to be reflected in the patterns of our reactive attitudes. And this undermines the attraction of the response-dependence view not because the two agents should be morally responsible to different degrees from a response-independent perspective (which would beg the question); rather, the reason

is that they should be morally responsible to different degrees from the perspective of the plausible substantial conditions. We find the reasons-responsiveness condition intuitive not just because their effects only exist at the coarse-grained level. Instead, as illustrated in the case of Heru and Henry, our intuition supports that reasons-responsiveness should affect degrees of responsibility both at the coarse-grained and the fine-grained level. Thus, in the absence of a plausible debunking story, the apparent attraction of response-dependence is seriously undermined once it is shown that the view cannot really capture the full range of applications of our general intuition.<sup>10</sup>

The same form of argument can be made about the quality of will condition. Theorists differ on how to cash out the notion of quality of will, but it is often taken to be a sort of care or regard that one has of other people's moral interests. Care or regard may be similar to reactive attitudes in terms of their granularity. But the degree of *manifestation* of a bad quality of will can still be fairly fine-grained. To see this, note that manifestation typically implies a causal or explanatory relation. But the substandard quality of will does not have to be the full causal or explanatory story; it can just be a partial cause or partial explanation. At the intuitive level, the concepts of "more of a cause" and "more of an explanation" can be fairly fine-grained. Recent work on degrees of causation in metaphysics (for a review, see Kaiserman 2018) adds further support to this intuitive claim by laying out the various models in which the degree of causation can be fine-tuned. We can then construct cases similar to that of Heru and Henry. Imagine that Dane has a substandard quality of will toward her roommate, and this partially causes or explains her stealing her roommate's yogurt from the fridge. But this is not the full causal or explanatory story, since Dane is also under some pressure from

another person, Daniel, to steal the yogurt. We can then vary the extent of the causal or explanatory role that Dane's bad quality of will (rather than the pressure from Daniel) plays, in a way that the difference can be too small to be registered in the resentment of Dane's roommate. That is, though response-dependence initially appears attractive in virtue of being able to capture the quality of will condition, it turns out that our intuition supports that the quality of will condition should affect degrees of responsibility both at the coarse-grained and the fine-grained level. Thus, response-dependence does not really capture the full range of applications of the quality of will condition, and the initial attraction of response-dependence is again seriously undermined.

Here is another way to put the dialectic so far. Response-dependence and response-independence are *meta-normative* theories. They are committed to a particular order of explanation, but do not by themselves offer substantial conditions for being responsible. By contrast, there are *substantial* conditions like reasons-responsiveness and quality of will. A theoretical virtue for a good meta-normative theory is that it can be combined with the most plausible substantial conditions as well as their intuitive applications. Since response-dependence can only be combined with a very narrow kind of applications—applications at a coarse-grained level—of the substantial conditions, this shows that response-dependence, as a meta-normative theory, has a serious theoretical cost. This cost is obviously worrisome for a neutral party who has not yet committed to response-dependence or response-independence, but I believe that even someone already committed to response-dependence also faces a serious challenge here, given that one major attraction of response-dependence—that it appears to accurately capture the substantial conditions of responsibility—is seriously undermined.



To sum up, my contention is that there is a mismatch challenge against response-dependence:

- (1) Facts about the reactive attitudes are fairly coarse-grained.
- (2) Responsibility facts are fairly fine-grained.
- (3) Both (1) and (2) are true when the standard of granularity is held fixed.
- (4) Thus, responsibility facts cannot be fully determined by facts about the reactive attitudes.

Let me emphasize that coarse-grainedness and fine-grainedness are relative terms; central to the challenge is the comparative mismatch between the two measures of granularity, which is why (3) is a necessary step of the inference. I have defended (1) in the last section and (2) and (3) in the current section. Combined, the conclusion (4) seems to contradict the response-dependence view, which attempts to ground responsibility facts in facts about reactive attitudes.<sup>11</sup>

## 5. POSSIBLE REPLIES TO THE MISMATCH CHALLENGE

### 5.1. *The Idealization Strategy*

A response-dependence theorist may reply by appealing to the difference between the actual and the ideal or *appropriate* reactive attitudes. The suggestion is that, even if we in fact do not have fine-grained reactive attitudes, we can still use fine-grained reactive attitudes as idealizations by maintaining that they are the appropriate ones to have. And since the formulation of response-dependence uses appropriate rather than actual reactive attitudes, one might suggest, there will not be any mismatch at all.

I suspect that, however, the current strategy is not clearly available to the response-dependence theorist. Response-dependence views of responsibility aim to naturalize the propriety condition for reactive attitudes, by grounding it in the global features of our practices of holding reactive attitudes toward

people. It is then not clear how, by doing so, they can make the appropriate reactive attitudes more fine-grained than the actual ones. The specifics will of course depend on how one understands the propriety condition. Remember two common approaches are the simple dispositional approach and the fitting approach. The former route finds the appropriate reactive attitudes in our dispositions under standard conditions. The standard conditions are sometimes understood as just the *normal* conditions, those that we usually satisfy when there are no dispositional masks, finks, and the like. But this does not help with the granularity issue. Putting the standard conditions in more idealized terms does not help either, at least according to the classic accounts. Remember that Lewis takes the condition to be one in which we have “full imaginative acquaintance” with the subject, and Railton takes it to be one in which we are fully informed of the circumstances and fully free of cognitive errors. But neither more acquaintance nor being more informed can help make our reactive attitudes much more fine-grained than they already are. Again, both Lewis and Railton aim to defend response-dependence for value in general; however, it is reasonable to expect some response-dependence theorists of responsibility to appeal to very similar ways of understanding the standard condition as well.

The fitting route finds the appropriate reactive attitude in the fitting responses. Remember that the notion of fittingness is used to rule out those responses of people with bad “anger sensibilities” (Shoemaker 2017). According to Shoemaker, a good moral sensibility is analogous to the wine sensibility of a sommelier, in contrast to that of a diner (2017, p. 511). One may then suggest that appealing to this refined sort of sensibilities can help avoid the mismatch challenge. However, I am inclined to think that the challenge can also be applied to those with good anger sensibilities.

I tend to think that, in the earlier example, Ann does not feel angrier toward Henry than toward Heru, *even if* Ann has the good anger sensibilities. This is because the limitation of information registration is not due to the poor sense of any individual, but due to a general limitation in human psychology. Again, there is surely a sense in which we can develop fairly sophisticated, refined capacities to track moral responsibility facts, if we include among these capacities the cognitive capacity of making judgments about degrees of moral responsibility. However, if response-dependence theorists were to appeal to such a cognitive capacity to understand good anger sensibilities, then the alleged fundamental status of reactive attitudes would again be undermined. But if they were to exclude this cognitive capacity, then the development of anger sensibilities would be greatly limited. It does not seem that we could refine the granularity of anger or resentment that much without employing our capacity to learn and apply what to *believe* about degrees of moral responsibility.

It thus seems that the two common routes of understanding the propriety condition are not helpful for formulating a convincing reply to the mismatch challenge. As a result, for those who want to pursue the idealization strategy in response to the challenge, they would need a new way to explicate the notion of idealization in enough detail and provide independent motivations for favoring it over its alternatives. It is unclear to me how one can do so, but I believe that this will lead to much valuable inquiry.

### 5.2. *The Restriction Strategy*

On a different kind of reply, one may suggest that response-dependence should be restricted to *facts about whether one is morally responsible*, which is compatible with response-independent explanations for *facts about degrees of moral responsibility*. Indeed, this may even be a reasonable

reinterpretation of the actual view of many proponents of response-dependence. The conclusion of my mismatch challenge then, one may suggest, turns out to be consistent with a response-dependence approach about moral responsibility.

I think this is a promising reply, but I worry that it significantly compromises the theoretical payoff of endorsing response-dependence. For example, one motivation for response-dependence is its simplicity; but this would be undermined if we need to supplement the account using response-independent explanations anyway. It is also unclear the extent to which this hybrid picture will help answer the question of how people can be responsible if metaphysical determinism is true. After all, if only facts about whether one is morally responsible is response-dependent, then the approach provides no explanation for why it cannot be the case that everyone is morally responsible to *an extremely small degree* if determinism is true. To say the least, defenders of the current strategy need to explain why a hybrid picture is well-motivated, and, in particular, why it does not already concede too much and still preserves many of the theoretical advantages of the pure response-dependence approach.

## 6. CONCLUDING REMARKS

I have raised a granularity challenge against the response-dependence view of moral responsibility. Response-independence theorists, by contrast, do not face the same challenge since they are free to use factors external to our practices of holding each other responsible in grounding the propriety conditions for reactive attitudes. Thus, the granularity challenge lends support to the response-independence theories. However, it is worth noting that the mismatch in granularity also has important methodological upshots even for the response-independence theorist—for example, the mismatch suggests that we will not be able to use fine degrees of

emotions as *epistemological* guides to facts about degrees of responsibility. In general, the upshot is that the mismatch challenge should lead us to reevaluate and properly restrict the theoretical role of reactive attitudes.

It remains an open question whether a similar granularity puzzle is as challenging in domains other than moral responsibility. For example, are our pro-attitudes too coarse-grained to be the metaphysical basis of values? Are the psychological features of respect too coarse-grained to be the metaphysical basis for dignity? Are our legal practices too coarse-grained to be the metaphysical basis for legal culpability? These questions are

not to be solved here; but note there is also no need to demand the same kind of answer to this cluster of questions—for example, it might turn out that dignity is a fairly coarse-grained notion and can be fully determined by the psychological features of respect, but value is too fine-grained to be fully determined by our pro-attitudes. The bottom line is that there is a real granularity puzzle for response-dependence in moral responsibility, and the plausibility of response-dependence approaches in other domains should be subject to a similar granularity test.

*University of California, San Diego*

#### NOTES

For valuable comments and feedback on earlier versions of this paper, I want to thank many people including Craig Agule, William Albuquerque, Santiago Amaya, David Brink, Rosalind Chaplin, Kathleen Connelly, Cory Davia, Joseph Martinez, Michael McKenna, Dana Kay Nelkin, Ahmed Siddiqi, David Shoemaker, Marshall Thompson, Manuel Vargas, Daniel Weltman, Monique Wonderly, the participants at the 2019 Free Will, Moral Responsibility, and Agency Graduate Conference at Florida State University, the audience at Lingnan University, and the participants at the Agency and Responsibility Group at UC San Diego. Special thanks to the two anonymous referees for *American Philosophical Quarterly* for their helpful comments.

1. Note that it is not always clear the extent to which the above authors endorse response-dependence. However, they all clearly identify the view as occupying an important conceptual space.
2. Some versions of response-dependence might prefer replacing “appropriate” with “fitting,” “reasonable,” “apt,” and the like. I use “appropriate” as a broad normative notion that incorporates all those formulations.
3. That is, unless one endorses a mutual dependence view according to which facts about being responsible and facts about reactive attitudes are mutually dependent. McKenna (2012) suggests such a view. I will put this view aside and instead focus on response-dependence and response-independence, both of which take the order to dependence to be one-way rather than mutual.
4. Carlsson (2017) argues that guilt, instead of resentment or indignation, is the reactive attitude that figures in (I). Though I will focus more on resentment, I believe that similar issues about granularity will arise for the emotion of guilt as well.
5. Another influential account that is often referred to as being response-dependent is Wallace’s (1994) fairness view, according to which reactive attitudes are appropriate just in case they are *fair*. However, as Todd (2016) has pointed out, there are important differences between Wallace’s view and other theorists who endorse or are sympathetic to response-dependence. In particular, Todd makes the interesting observation that Wallace never identifies traditional accounts of responsibility (e.g., libertarianism) as taking the wrong side in the order of explanation, in contrast to other response-dependence theorists. I’m inclined to think Wallace’s view is open to both response-dependent and response-independent

interpretations, depending on how we understand the notion of fairness. Wallace is therefore the target of the current paper only insofar as his notion of fairness can be explained in a purely response-dependent way.

6. The inference relies on the widely accepted assumption that grounding is a transitive relation.
7. There are parallel cases where two agents have the same capacities, but one's action is slightly more wrong than the other's. Imagine that Fei stole \$100 from Ann, whereas Fara stole \$101 from Ann. Though Fara's action is slightly more wrong than Fei's, and Fara is thereby more blameworthy, Ann's resentment does not seem to have the level of granularity required to track this difference. To the extent that response-dependence theorists about moral responsibility also endorse response-dependence about blameworthiness, cases like this can be used to construct a structurally similar granularity challenge for them as well.
8. See Fischer & Ravizza (1998), Fischer (2004), Wolf (1990), Nelkin (2011), Vargas (2013), McKenna (2016).
9. See Scanlon (2008), Smith (2008), Talbert (2008), Hieronymi (2004), McKenna (2012), Shoemaker (2015).
10. See Coates & Swenson (2013) and Nelkin (2016) for alternative, but also fairly fine-grained, characterizations of degrees of responsibility in a reasons-responsiveness framework.
11. There is a different perspective to look at the mismatch in degrees of reactive attitudes, by appealing to the acceptable stability and variation for appropriate reactive attitudes. It seems the propriety condition for reactive attitudes should allow for a reasonable degree of instability. It is typically not inappropriate for a person to feel a certain degree of resentment toward a wrongdoer at one time, but at a later time feel a bit less resentment toward the same wrongdoer. This can be true even when the person neither forgives the wrongdoer nor receives any further information. If both the initial, heated-up emotion and the later, cooled-down emotion are appropriate, how can the response-dependence theorist tell a story of grounding degrees of responsibility using degrees of emotions? More work needs to be done to tell a convincing story here. Bykvist (2009) appeals to similar concerns to argue against fitting attitude theories of value. He contends that one difficulty consists in the fact that "how strongly we should react emotionally seems also to depend on temporal matters" (Bykvist 2009, p. 16).

## REFERENCES

- Barrett, Lisa F. 2006. "Valence is a Basic Building Block of Emotional Life," *Journal of Research in Personality*, vol. 40, no. 1, pp. 35–55.
- Beglin, David. 2018. "Responsibility, Libertarians, and the 'Facts as We Know Them': A Concern-Based Construal of Strawson's Reversal," *Ethics*, vol. 128, no. 3, pp. 612–625.
- Brink, David O., & Nelkin, Dana K. 2013. "Fairness and the Architecture of Responsibility," in *Oxford Studies in Agency and Responsibility*, Vol. 1, ed. David Shoemaker (New York: Oxford University Press), pp. 284–313.
- Bykvist, Krister. 2009. "No Good Fit: Why the Fitting Attitude Analysis of Value Fails," *Mind*, vol. 118, no. 469, pp. 1–30.
- Carlsson, Andreas B. 2017. "Blameworthiness as Deserved Guilt," *Journal of Ethics*, vol. 21, no. 1, pp. 89–115.
- Coates, D. Justin., & Swenson, Philip. 2013. "Reasons-Responsiveness and Degrees of Responsibility," *Philosophical Studies*, vol. 165, no. 2, pp. 629–645.
- Fischer, John M. 2004. "Responsibility and Manipulation." *The Journal of Ethics*, vol. 8, no. 2, pp. 145–177.

- Fischer, John M., & Ravizza, Mark. 1998. *Responsibility and Control: A Theory of Moral Responsibility* (Cambridge: Cambridge University Press).
- Hieronymi, Pamela. 2004. "The Force and Fairness of Blame," *Philosophical Topics*, vol. 18, pp. 115–148.
- Kaiserman, Alex. 2018. "'More of a Cause': Recent Work on Degrees of Causation and Responsibility," *Philosophy Compass*, vol. 13, e12498.
- Larsen, Jeff T., McGraw, A. Peter., & Cacioppo, John T. 2001. "Can People Feel Happy and Sad at the Same Time?" *Journal of Personality and Social Psychology*, vol. 81, no. 4, pp. 684–696.
- Lewis, David. 1989. "Dispositional Theories of Value," *Proceedings of the Aristotelian Society, Supplementary Volumes*, vol. 63, pp. 113–137.
- Mahtani, Anna. 2020. "The Dispositional Account of Credence," *Philosophical Studies*, vol. 177, pp. 727–745.
- McKenna, Michael S. 2012. *Conversation and Responsibility*. New York: Oxford University Press.
- . 2016. "Reasons-Responsiveness Theories of Freedom," in *The Routledge Companion to Free Will*, ed. Kevin Timpe, Meghan Griffith, & Neil Levy (London: Routledge), pp. 49–62.
- Menges, Leonhard. 2017. "Grounding Responsibility in Appropriate Blame," *American Philosophical Quarterly*, vol. 54, no. 1, pp. 15–24.
- Nelkin, Dana K. 2011. *Making Sense of Freedom and Moral Responsibility* (New York: Oxford University Press).
- . 2016. "Difficulty and Degrees of Moral Praiseworthiness and Blameworthiness," *Nous*, vol. 50, no. 2, pp. 356–378.
- Railton, Peter. 1986. "Facts and Values," *Philosophical Topics*, vol. 14, no. 2, pp. 5–31.
- Scanlon, Theodore M. 2008. *Moral Dimensions: Permissibility, Meaning, Blame* (Cambridge: Harvard University Press).
- Shoemaker, David. 2015. *Responsibility from the Margins*. Oxford University Press.
- . 2017. "Response-Dependent Responsibility; Or, a Funny Thing Happened on the Way to Blame," *Philosophical Review*, vol. 126, no. 4, pp. 481–527.
- Smith, Angela M. 2007. "On Being Responsible and Holding Responsible," *The Journal of Ethics*, vol. 11, no. 4, pp. 465–484.
- . 2008. "Control, Responsibility, and Moral Assessment," *Philosophical Studies*, vol. 138, pp. 367–392.
- Strawson, Peter. 1962. "Freedom and Resentment," *Proceedings of the British Academy*, vol. 48, pp. 1–25.
- Talbert, Matthew. 2008. "Blame and Responsiveness to Moral Reasons: Are Psychopaths Blameworthy?" *Pacific Philosophical Quarterly*, vol. 89, pp. 516–535.
- Todd, Patrick. 2016. "Strawson, Moral Responsibility, and the 'Order of Explanation': An Intervention," *Ethics*, vol. 127, no. 1, pp. 208–240.
- Vargas, Manuel. 2013. *Building Better Beings: A Theory of Moral Responsibility* (New York: Oxford University Press).
- Wallace, R. Jay. 1994. *Responsibility and the Moral Sentiments* (Cambridge, MA: Harvard University Press).
- Watson, Gary. 1987. "Responsibility and the Limits of Evil: Variations on a Strawsonian Theme," in *Responsibility, Character, and the Emotions: New Essays in Moral Psychology*, ed. Ferdinand Schoeman (Cambridge: Cambridge University Press), pp. 256–286.
- . 2014. "Peter Strawson on Responsibility and Sociality," in *Oxford Studies in Agency and Responsibility, Vol. 2*, ed. David Shoemaker & Neal Tognazzini (New York: Oxford University Press), pp. 15–32.
- Wolf, Susan. 1990. *Freedom within Reason* (New York: Oxford University Press).
- Zimmerman, Michael J. 2010. "Responsibility, Reaction, and Value," *Journal of Ethics*, vol. 14, no. 2, pp. 103–115.