

Two Ways of Socializing Moral Responsibility

Circumstantialism versus Scaffolded-Responsiveness

JULES HOLROYD

1 Introduction

This chapter is concerned with the implications of two important developments in the recent literature on moral responsibility: first, there has been a sustained defense of the view that moral responsibility is constituted by our social practices, rather than any metaphysically deep notion antecedent to our social relations (the *social thesis*). Second, it has been asked whether and why we might want such practices (the *justification thesis*). The answers to these questions then guide the articulation and justification of the standards for responsible agency and the norms internal to the practice that govern when individuals are liable to the moral responses of praise or blame.

My main aim here is to evaluate two competing views of morally responsible agency, each of which has been developed to sit within a view of responsibility that endorses the *social thesis* and the *justification thesis*. These two views are Vargas's *circumstantialism* (2013, 2015)—the view that responsible agency is a function of the agent and her circumstances, and so highly context sensitive—and McGeer's *scaffolded-responsiveness* view (2015; McGeer and Pettit 2015), according to which our responsible agency is constituted by our capacity for responsiveness both to reasons themselves, and to the expectations of our audience (whose sensitivity may be more developed than our own). I develop and defend a version of the *scaffolded-responsiveness* view. I then make some suggestive remarks about how such a view might be located within a broader conception of our practices of responsibility: that such a view coheres with a picture of our practices in which moral responsibility is implicated in a range of moral

responses beyond blame or praise. Finally, I suggest that these conclusions about moral responsibility and the range of moral response may be further bolstered by approaching these issues via what has been called an “ameliorative” (cf. Haslanger 2000) rather than revisionist analysis. Revisionist analyses seek to revise our existing concept of responsibility, and so remain anchored in our extant concept. Instead, an ameliorative analysis starts by asking what we want the concept of responsibility for and what concept will serve those purposes, with no assumption that the answers we give will yield an analysis that closely tracks our existing understanding of moral responsibility.

2 The Social Thesis

If we want to know what moral responsibility is and when we are warranted in deploying the concept, how should we conduct our inquiry? One approach would be to analyze our intuitions about the concept, revealing what conditions have to be in place in order to be warranted in deploying the concept in our interactions with others. What does the world have to be like (what metaphysical conditions must obtain) and what do we have to be like as agents, in order to stand in certain relations to our actions such that we are morally responsible for them?

Revisionists argue that the concept of moral responsibility must be revised if it is to deliver a coherent set of conditions for responsibility. This idea is supported by pointing to seemingly inconsistent intuitions about responsibility (for example, that ultimate control is, and is not, necessary for responsibility; that determinism is, and is not, necessary for responsibility). Some of these intuitions, then, must be forsaken, and a revised concept adopted. How might one go about this revisionary project? A revisionist approach considers (1) what work the concept of moral responsibility does for us in our social practices, (2) whether these practices are justifiable, and (3) whether a defensible notion of responsibility is embedded within those practices, or can be developed from our understanding of them. Revisionism about responsibility has been most fully developed and defended in the recent work of Manuel Vargas (2013, 2015). In Vargas’s revisionism, we end up with a revised concept that rejects incompatibilist intuitions, according to which responsibility is dependent upon some metaphysically deep feature of agents and their actions (say, the need or possibility of ultimate control). Rather, when we look at our social practices, we find a concept that (1) has a role in keeping track of deserved praise and blame. Vargas claims that (2) these practices can be justified if the practices as a whole serve to cultivate a certain valuable kind of agency, namely, agency that is sensitive to moral considerations. And the notion of responsibility that is

embedded within these practices is (3) a circumstantialist view of responsible agency—more details of which to follow. Whilst the core of the social thesis is that morally responsible agency is a matter of standing in certain social relations (rather than certain metaphysical conditions obtaining), there are various revisionist alternatives consistent with this position. Some such alternatives may hold little appeal: consider, for example, the options of nihilism—the view that moral responsibility is impossible for us—or eliminativism—the prescription that the notion of moral responsibility should be eliminated from our normative discourse. Each of these latter options rejects some of our intuitions about moral responsibility. But there would be costs to endorsing either view—in particular, giving up the important benefits that can be gained from our practices of holding each other responsible (cf. Vargas 2013, 107). If an adequate justification of the social practices of holding certain agents responsible can be delivered, and if the conception of responsibility derived from this is a defensible one, then we have little reason to pursue either of these bleaker alternatives.¹ However, even if we seek to reject these two bleak options, there may be other—still revisionist—alternatives to Vargas’s circumstantialism about moral responsibility, such as the scaffolded-responsiveness view defended here.

3 The Justification Thesis

The *social thesis* has been defended before, both by consequentialists, and famously by P. F. Strawson (1962). Both views see responsibility as constituted by a set of social practices: but each give different accounts of what justifies those practices, and thereby of what responsibility is. On a ‘Strawsonian psychologism,’ we hold that we cannot but engage in these practices—our psychology is simply geared up to feeling the reactive attitudes (resentment or gratitude) in response to the (poor or good) quality of will we find manifested in others’ actions. To be responsible is to be liable to be held responsible, on this view. But this justification takes a strong view on our psychological limits, and supposes implausibly that these patterns of affective response are unchangeable: a fixed point in human relations. Moreover, even if we accept such constraints, this would explain, rather than justify, our practices (cf. Vargas 2013, 160).²

However, the Strawsonian option has seemed all the more appealing when set alongside the competing claim, from consequentialist thinking, that what justifies the practice is the good outcome of praise and blame. On this view, praise and blame have a certain moral influence, encouraging people to do good and discouraging or deterring them from bad action. You simply are morally responsible if you are susceptible to be influenced in this way. You are blameworthy if blaming you effects good outcomes (see, e.g., Smart 1973). This sort of

view seems coldly manipulative, and it circumvents the agential capacities that we might hope to be engaged in moral communications (see Vargas [2013, 166] for detailed discussion of each of these views).

Both of these views endorse the *social thesis*, and see responsibility as essentially a social rather than metaphysical fact: an agent is responsible and liable to blame given the constraints of our psychological dispositions to view her in that way, or given the efficacy of viewing her thus. But neither of these views provides adequate justification to shore up these social practices.

An alternative option has recently been developed by Vargas (2013) and “heartily endorsed” by McGeer (2015, 2637). Vargas’s important revisionary move is to develop a revised moral influence view, which points to the good effects of the responsibility system for fostering a certain sort of agency, namely, agency that is sensitive to moral considerations. This is the ‘agency cultivation model.’ On this view,

when we hold one another responsible, we participate in a system of practices, attitudes and judgments that support a special kind of self-governance, one whereby we recognize and suitably respond to moral considerations. So, roughly, moralized praise and blame are justified by their effects, that is, how they develop and sustain a valuable form of agency. (Vargas 2013, 2; cf. McGeer 2015, 2637)

The practice as a whole—the norms regulating judgments about and expressions of deserved praise or blame—can be justified if it promotes this valuable sort of agency.³ Insofar as the practice has this function—supporting a certain kind of self-governance, cultivating and sustaining agency that is sensitive to moral considerations—it looks like we have good reason to regard these practices as justified.

4 Morally Responsible Agency

If responsibility is constituted by a social practice, and that social practice can be justified, we still lack an answer to our question of what responsible agency is: when is an individual a responsible agent? In other words, *who* is a participant in those practices? We should hope the account of responsible agency has independent plausibility—that it coheres with our intuitions, or the contours of the practices with which we are familiar, and that it resonates with why we care about responsibility. We have additional resources we can draw on in evaluating an account of responsible agency, as provided by the *social thesis* and *justification thesis*. The account of responsible agency should be one that can function

in the social practice of holding each other responsible. Further, the notion of responsible agency should be such that, when deployed, it coheres with and even promotes the overall aims of the practice, namely, the cultivation of morally sensitive agency (cf. Vargas 2013, 220). We already have a head start in this task, since Vargas and McGeer both endorse this revisionist approach, and have each developed accounts of morally responsible agency, located within this framework. However, they each arrive at radically different conceptions of morally responsible agency. Next I articulate the competing accounts; then I go on to evaluate them.

4.1 Circumstantialism

Circumstantialism about moral responsibility builds in a familiar claim about the relevance of the agent's reasons-responsiveness to her responsibility. This form of morally responsible agency is understood in terms of a (context-specific) capacity to recognize moral considerations, and to govern ourselves in light of our recognition of these. As with standard reasons-responsive accounts, the capacity involves responsiveness in sufficiently many (counterfactually specified) occasions, in a sufficiently systematic way. But unlike the more familiar versions of reasons-responsiveness views (e.g., Fischer and Ravizza 2000), Vargas rejects atomism and monism about responsibility. That is, he rejects the claim that the relevant capacity is a nonrelational property of agents (Vargas 2013, 204); and he rejects the claim that there is one psychological structure or mental capacity that is important for or constitutive of being suitably responsive to reasons.

Instead, we get a picture whereby the relevant capacities are relational: “a function of whether the agent (with the relevant features in the considered context of action) stands in a particular relationship to the normative practice” (Vargas 2015, 2622); and whereby the capacities are plural: there are “multiple agential structures of combinations of powers that constitute the control or freedom required for moral responsibility” (Vargas 2013, 205). There are good reasons for adopting this sort of relational and plural view of the relevant capacities: literature from empirical psychology has demonstrated that an individual's ability to respond to reasons can be highly dependent on the particular circumstances, and influenced by, for example, the primes provided by the environmental context. One example that Vargas uses to illustrate “the startlingly localized and context-specific potential for degrading our rational capacities” (2013, 207) is that of stereotype threat: the phenomena whereby the activation of a stereotype provokes a stress response that can hinder an individual's performance—desensitize them to reasons, we might say. In a context in which stereotype threat is in operation, an agent's capacities are degraded by those circumstances and the stereotypes activated in them, which impair an individual's ability to respond

to the reasons she otherwise would.⁴ Even from the armchair, Vargas suggests, it would appear that various different capacities are involved in different aspects of moral agency (2013, 208). For example, the capacities involved in seeing that someone is upset, or in arguing that one has a particular duty, or in recognizing a right, may each differ significantly. Perceptual, emotional, and cognitive reasoning capacities may each be demanded, and different situations may demand more or less from one or other of these capacities. Moreover, situational features may impair each capacity in variegated and unexpected ways.

What we end up with is a view whereby highly context-sensitive sets of capacities are the constituents of an agent's reasons-responsiveness. But the crucial question for this kind of view is what degree of reasons-responsiveness an agent must be capable of in order to be a responsible agent. We already know that for the circumstantialist this will not be an all-or-nothing judgment: for Vargas, an agent might be a responsible agent in one context, and not in another. What does it take in any one context, though? For any context, there is no antecedently given answer to how sensitive to reasons an agent must be. Rather, that standard is given by the forward-looking aims of the practice, to cultivate the agent's moral sensitivity. Accordingly, it depends on "whether the relevant specification of the capacity would best support agents in the actual world in the recognition and suitable response to moral considerations" (Vargas 2015, 2622). The relevant standard—the degree and kind of responsiveness—constitutive of responsible agency in a context is set by a standard, which is "actually co-optimal or better for fostering agency that recognizes and suitably responds to moral considerations in the actual world, in ordinary contexts of action" (Vargas 2013, 219).⁵ Vargas notes that a secondary aim in setting this standard is that it should serve to expand the contexts in which agents demonstrate such sensitivity (Vargas 2013, 220). So, a particular contextual feature (such as stereotype threat) may impair an agent's sensitivity to reasons; but whether it does so in a way that undermines an individual's responsibility is not something that can be answered without recourse to the forward-looking aims of the practice. Would a system of norms that supports holding the agent responsible in such contexts, given such sensitivities, better serve the aims of cultivating sensitive agency?⁶ This delivers the answer to whether an agent, in a context, is responsible. We will consider an example of how responsibility is determined in more detail shortly.

There are two crucial points to emphasize about this circumstantialist view: first, the circumstantialist claim is one about morally responsible agency, rather than liability to blame. This means that our responsible agency is more "patchy" on this view than on others. For example, if an environment blocks our sensitivity in some way, it is not simply that we are responsible agents who are excused: rather, our responsible agency itself is absent. One implication of this is that much attention is needed to the "moral ecology" in which we function as

moral agents (Vargas 2013, 243–249; see also 2017) and the environmental and social relations that might best cultivate such sensitivity. Secondly, it is important to note that Vargas is keen to avoid specifications of the relevant standards for sensitivity in a way that would render almost no one responsible. For example, we shouldn't require sensitivity across *almost all* possible similar contexts, articulated in fine-grained terms, as almost no one would meet that standard, and “we would be required to forgo the real benefits of responsibility” (Vargas 2013, 219). This wish to avoid general skepticism is motivated in part by Vargas's principle of conservatism (that we should not too radically revise our conception) and by the idea that we should seek to preserve the benefits attendant on the practices of responsibility.

4.2 Scaffolded-Responsiveness

Let us now contrast this circumstantialist view with what I am calling the *scaffolded-responsiveness view* of morally responsible agency, developed by McGeer (2015), and McGeer and Pettit (2015) across a series of papers. As mentioned above, both McGeer and Vargas articulate their respective accounts of responsible agency within a view that sees responsibility as constituted by social practices, and that maintains that these practices are justified by a forward-looking aim of fostering morally sensitive agency. However, the conceptions of responsibility yielded are in many respects radically divergent. Like Vargas, McGeer (2015) endorses a broadly reasons-responsive view of responsible agency, whereby an agent must have the capacity to register and respond to certain reasons to qualify as a responsible agent. However, whilst both Vargas and McGeer reject atomism, they do so in different ways. As set out above, Vargas construes the capacities as standing in relation to a normative practice in a context. Meanwhile, on the *scaffolded-responsiveness* view of reasons-responsiveness, this capacity is not simply a matter of being able directly, here and now, to register and act on certain reasons—McGeer and Pettit (2015) draw attention to another way in which we might be responsive to reasons: namely, indirectly. On this view, one is reasons-responsive also if one is able to adjust or sensitize to the reasons that there may be. Just as an immune system has the capacity both to defend against diseases it is already sensitized to, it can also become reactive to new diseases. Similarly, we may already be sensitive to some reasons, but also we may have the capacity to become sensitized to new reasons (McGeer and Pettit 2015, 163). In particular, this capacity to become sensitized operates via our sensitivity to our “audience”—those whose (anticipated) judgments about our conduct matter to us (168). Accordingly, “as a reason-responsive agent you can be expected, not only to be moved by the reasons you confront in making a choice, but also by the audience, actual or prospective, that you confront” (170).

The audience should provide an environment in which “we are continually exhorted by one another to exercise that capacity; to think about our actions; to justify them to one another; to work on our weaknesses . . . to nurture our strengths” (McGeer 2015, 2647). The expectation of one’s audience provides a “situational force”—one that we may care about both in itself and instrumentally. Accordingly, the responsibility-relevant capacities are relational in that their scope is dependent on one’s audience, and this may expand the range of reasons to which an individual is sensitized: both those she is able to detect, and those she is sensitized to via her (actual or prospective) audience.⁷ In this way, others “scaffold” our responsiveness to reasons, expanding the range of reasons to which we may be properly said to be responsive.

Note that this articulation of audience involves a certain optimism: for some, the audience may not in fact provide this sort of scaffolding. Thus, one underspecified aspect of the account is the interpretation of one’s *prospective* audience: does this include only the audience one realistically expects to face? If so, there may be cases in which one’s actual audience is uncritical and unsupportive. On this interpretation, an uncritical audience closes off opportunities for enhancing sensitivity. Alternatively, the conception of ‘prospective audience’ may extend to audiences one can conceive of, and who populate the moral community, but whom in fact one is unlikely to face in one’s actual experience.⁸ For example, under one idealization, one’s audience will include the sorts of subgroups with new moral insights, who problematize unconsidered moral issues and “force attentiveness in agents who are already, for the most part, morally reflective and sensitive” (Isaacs 1997, 681). On this interpretation, even if one’s sensitivity is restricted by one’s actual audience, it is enhanced if one could be sensitized to the reasons pointed to by this possible prospective audience.

On either way of developing the view, the scaffolded-responsiveness picture rejects atomism and articulates a relational standard for the sensitivity required for responsible agency. Note the contrast with circumstantialism, which indexes the scope of the agent’s capacities to include her circumstances. This has the effect of limiting the scope of an agent’s responsibility: where circumstances impair or degrade an agent’s capacities, she may not be responsible in that context. In contrast, on the *scaffolded* view, making the constituent sensitivity relational to *others’* sensitivity has the function of potentially expanding her capacities and the range of considerations to which an agent might be—directly or indirectly—sensitive. Crucially, even if an agent’s capacities are not in fact expanded, the insight that to be a responsible agent is to have the potential to be scaffolded by our (actual or prospective, realistic or idealized) audience captures an important relational dimension of responsible agency.⁹ This, it is argued, coheres with the forward-looking aims of the practice, since engaging with individuals as responsible agents—engaging in moral response such as expressions of blame—can

serve to bolster or reinforce those capacities, exhorting the agent to exercise it and aim for greater (direct) attentiveness to reasons. As McGeer and Pettit (2015) put it, addressing each other as morally responsible agents can expand the scope of reasons to which we are responsive: we can *capacitate* each other.¹⁰

We have, then, two considerably different pictures of responsible agency that we may plug into our framework for understanding morally responsible agency as grounded in social practice. Is one or the other of these views more defensible? In the next section, I outline some considerations that speak in favor of the scaffolded-responsiveness view.

5 How Should We Think about Responsible Agency?

How might we decide which account of responsible agency to endorse?¹¹ First, we can consider which notion might best serve the justificatory aim of the practice. Second, we can consider which notion has independent plausibility. Third, we can look to our existing practices and ask how they might inform our thinking about which notion of responsibility we should endorse. To this end, we can start our evaluation with some examples of cases that have already generated much critical reflection; this helps us to see both how *circumstantialism* and *scaffolded-responsiveness* would deal with these cases, and the extent to which the considerations articulated above are met.

Instructive cases to focus on are the cases that fall under what Vargas describes as “new, unusual and particularly challenging” (2013, 183) contexts of action, since the two accounts (as we shall see) deal radically differently with such cases. Furthermore, we have a wealth of critical reflection on just such cases, since they have been the focus of much philosophical attention. These sorts of cases include those in which a moral issue is not conceived as such, so as to make difficult or inaccessible the moral insight required for sensitivity to the moral reasons at issue. One such case, discussed by Cheshire Calhoun (1989), concerns the use of gendered pronouns in a way that perpetuates sexist assumptions, and the difficulty of recognizing this in a cultural context in which language use of this sort is not widely considered a “moral issue”:¹²

C1: Bert, a male philosopher writing in the second half of the twentieth century, always uses the male pronoun “he” to refer to supposedly universal individuals, whilst writing his philosophical papers. He does so in a cultural context in which this is the norm, and in which the idea that the use of gendered language, in this way, is problematic is extremely marginal—maintained only by a few themselves marginalized feminist

AQ: 10: “I offer an alternative way of thinking about moral response, in section 5.6.” Please clarify if the cross reference to section should be section 6 rather than 5.6. (p. 14)

thinkers. It is certainly not a moral viewpoint with which Bert is familiar, and indeed he has not even considered this aspect of his language use as a “moral issue.”

Another sort of case, recently much discussed, can also be thought of as a special case of moral ignorance: namely, cases of discriminatory behavior that results from implicit bias. In these cases, we can suppose, our concern is that agents are largely ignorant of the reasons that they have to avoid implicitly biased, discriminatory behavior: namely, the fact that they likely harbor implicit associations, which influence their judgments and behaviors.

C2: Alfred is evaluating the CVs of predominantly white male and female philosophy graduates, with a view to drawing up a shortlist of candidates for a job. Alfred harbors implicit biases that influence his judgments of the CVs of female graduates, and those graduates who are not racialized white: although showing qualities comparable with those on the CVs of the white male graduates, he sees the latter as “shining” with a quality that the other CVs don’t seem to have. He draws up an all-white, all-male shortlist for the job. Alfred is not aware of the role of implicit bias in his judgment, nor indeed of facts about the risk of bias, and is confident that he has picked “the best” candidates.

As with the case of sexist language, for many people—Alfred included—it may not even register that there is a “moral issue” to which they are being insensitive. Yet the individual fails to be appropriately responsive to reasons that there are—to adopt procedures that may insulate decisions from bias, for example—and thus the question of their responsibility for this failure is activated. The general question of the blameworthiness of individuals in cases of moral ignorance—and indeed, in cases of implicit bias—is one that has been widely addressed, so I do not consider that here.¹³ Rather, what I want to ask is what determinations are made on either the circumstantialist or scaffolded-responsiveness view of responsible agency. This helps us to build a plausibility argument for one or other of the accounts.

5.1 Circumstantialism about Moral Ignorance

On the circumstantialist approach, whether individuals such as Bert or Alfred are morally responsible will depend upon whether they stand in the right relation to the normative practice. That is, it depends on whether the agent is sufficiently sensitive; but our determinations of what “sufficient sensitivity” involves look to the degree of sensitivity requisite according to the forward-looking aims

of the practice. In other words, the degree of sensitivity that an agent must have is specified according to whether holding agents responsible, in such contexts, serves to improve sensitivity to the relevant moral considerations—in these cases, the moral reasons that there are to use nonexclusionary language or to attend to the possibility of bias. Note that this is not a question of whether any particular token of blame or instance of holding responsible will be effective in cultivating sensitivity. Rather, the question concerns whether a system that specifies *this* standard (whatever that turns out to be) of sensitivity required for responsibility in such contexts will have these agency-improving effects.

Of course, this is an empirical question. It is not at all clear to me from the arm-chair what the answer to this question might be. However, Vargas's application of circumstantialism to "new and unusual" cases indicates the conclusion that agents are *not* responsible agents in such cases, that the requisite sensitivity in that context is not yet present, such that moral response (notably, blame) for failure to conform with reasons would not cultivate greater sensitivity. In relation to these sorts of cases, Vargas indicates that the agent's capacities are not supported by the context—the environment blocks the agent's sensitivity (2013, 183)—such that we reach the conclusion that the individual is not a responsible agent. And, recall that on the circumstantialist account responsible agency is patchy—the agent is not simply excused from blame; rather, they simply lack the status of morally responsible agent; hence they are not a candidate for liability to blame, given the circumstances. This is because on the circumstantialist view it is responsible agency, rather than any excusing condition, that does the work here.¹⁴ (This strikes me as problematic, but I hold off further evaluation until section 5.3.)

We get a clearer statement of this conclusion when Vargas takes circumstantialism in application specifically to the case of implicit bias (see Vargas, 2017). There may be some reasons to think that holding responsible may have good, sensitizing, effects: it may provide helpful moral feedback and promulgate narratives of self-governance that help agents to live up to them (25). However, all things considered, Vargas argues that the present context is not one in which agents are responsible for implicitly biased actions: blaming may provoke hostility and may prevent individuals from "buying in" to norms that prohibit implicit bias (26). Further, too few individuals are equipped with the relevant knowledge about the phenomenon of bias, or the possibility of mitigating it, to suppose that our current context is sufficiently supportive of sensitive agency. As such, Vargas suggests that the forward-looking aims of the practice are not met by holding people responsible: we do not yet have the moral ecology that supports such determinations (but may yet, in the not too distant future) (31).¹⁵ So, in "new, unusual or particularly challenging contexts" (2013, 183) such as the ones I have described, circumstantialism delivers the verdict that the individuals do *not* meet the standards for responsible agency.

5.2 Scaffolded-Responsiveness about Moral Ignorance

Before we evaluate the circumstantialist response, let us consider what the alternative view of responsible agency delivers in relation to this case. The scaffolded-responsiveness view is again concerned with whether the agent has the requisite capacities to register and respond to the reasons to avoid (in this example) gendered language. On this view, the capacity for sensitivity includes not only an individual's capacities here and now to register the moral reasons in play, but also an agent's capacities given the ability to adjust or sensitise to the reasons there are via sensitivity to an "audience." An individual meets these conditions for responsible agency—meets the standard of sensitivity—either if she is presently sufficiently sensitive to the reasons, or if her present capacities are such that she can be further capacitated by the prospective moral address of others—address thus furthers the forward-looking aims of the practice of improving sensitive agency. Recall that much depends on how audience is specified within the scaffolded-responsiveness view. If one is scaffolded by one's actual audience, then in instances in which one's audience is uncritical and provides no new considerations or challenges, one's sensitivity will not be enhanced by its presence. But, as we saw, an interpretation of the view that incorporates some idealization—that one's prospective audience does include those with moral insight or who present moral challenges—will provide greater opportunity for capacitation.

One example of the sort of capacitating interaction at issue may be the sorts of address discussed by feminist philosophers' engagement with such cases. (McGeer and Pettit [2015] do not suggest this; I am hereby elaborating their account in what I take to be defensible ways.) In dealing with cases of moral ignorance, a range of responses have been considered: Miranda Fricker (2010) delineates the notion of moral disappointment to articulate the contours of the appropriate moral response where we might have *hoped* for greater moral insight, even whilst cultural norms make that insight hard to achieve.¹⁶ Michelle Moody-Adams suggests there is scope for a kind of "forgiving moralizing" (1994, 303) in cases of moral ignorance, whereby we continue to maintain that an individual's moral agency has been implicated in a wrong, but also recognize *the great effort needed* to adopt an appropriately critical stance toward heretofore unquestioned aspects of the cultural context. Cheshire Calhoun has argued for the importance of reproach in such cases, in order to convey the obligation to change behavior (e.g., avoid sexist language), to *motivate* individuals to reflect and do better, and to *reinforce* a view of each other as agents capable of registering and reacting to the relevant reasons (i.e., to be "self-legislating," as Calhoun puts it [1989, 404]).

I do not here want to take a stance on which of these responses is the appropriate one in particular: the important point is that all of these authors focus on these kinds of moral responses to individuals as appropriate despite the agent's

failure of direct responsiveness to the moral reasons. Thinking about cases of moral ignorance helpfully provides substance, I think, to the claim from McGeer and Pettit (2015) that our sensitivity is a function not just of our present responsiveness to reasons, but also of our sensitivity to audiences and their exhortations to adjust and expand the range of our direct attentiveness to reasons. Let us briefly note—this is a point to which we later return (in section 6 below)—that none of these authors focus on blame or blameworthiness per se as the appropriate reaction to the moral agents involved. The main point to tease out for now is that these forms of moral address are construed as appropriately directed toward the agent's failure of sensitivity. Moreover, the responses are construed as capacitating and enabling the agent to adjust so as to be directly sensitized to the considerations at issue. Thus, the scaffolded-responsiveness view yields the determination that agents in the sorts of moral ignorance cases I have described may yet be responsible agents. What matters is that they have the capacity to be sensitive to reasons indirectly, via a (perhaps idealized) prospective or actual audience (as well as the capacity to be directly sensitive to certain reasons themselves).

5.3 Which View?

The two views under consideration appear to yield substantially different conclusions about how to deal with cases of moral ignorance. Is one approach more defensible? In this section, I use three parameters in order to evaluate the competing views. First, I consider whether either view better serves the aim of the practice (of cultivating moral-considerations-sensitive agency). Second, I ask whether each view has independent plausibility. Third, I consider the extent to which each view coheres with our extant practices of holding each other responsible.

The Aims of the Practice

Each view of responsible agency is positioned within the *social thesis* and the *justification thesis*, whereby responsible agency arises out of our social practices, which have the purpose of cultivating reasons-sensitive agency. Each view positions itself as able to serve that aim. Yet the views, as I understand them in relation to our cases of moral ignorance, deliver radically different pictures of responsible agency and its implications in these cases. Do we have reason to believe that one or other of these views would better serve the aim of cultivating reasons-sensitive agency?

One of the key insights from the feminist literature on moral ignorance is that moral responses to failures of sensitivity are important for moral change.

As Calhoun puts it, “[A]n excusing response to moral ignorance precludes the social growth of moral knowledge” (1989, 406). And whilst Vargas raises worries about holding individuals responsible in such new and unusual cases—that this may provoke hostility or backlash (2017, 25)—some empirical studies suggest otherwise. In the case of implicit bias, moral confrontations have been found to reduce the later expression of bias (Czopp, Monteith, and Mark 2006). And, in a study specifically on the effects of blaming responses, whilst blame did not reduce implicit bias significantly more than neutral social interactions, nor did it increase bias; and, most significantly, blamed individuals more strongly expressed intentions to change future behaviors to avoid bias (Scaife et al., ms.). So, holding each other responsible may well increase sensitivity to moral reasons. Moreover, there is the danger, to which Vargas is alert, of promulgating hopelessness or complacency via narratives of lack of control. Indeed, in one empirical study, when people were told about the pervasive nature of implicit bias, in the absence of a strong moral norm against this pattern also being communicated, more implicit bias was subsequently displayed (Duiguid and Thomas-Hunt 2015). The circumstantialist view exempts individuals from moral responsibility if contexts do not support sensitivity to the relevant reasons. But in fact, this recent evidence supports Calhoun’s contention that in cases of culturally supported moral ignorance, morally engaging individuals may better cultivate moral sensitivity, and so foster the aims of the practice.

Independent Plausibility

Even whilst offering a revisionary account of moral responsibility, we should still hope that the account offered remains recognizably a concept of moral responsibility, one that resonates with the notion—or notions—we deploy in our social practice. It is worth noting that much of the philosophical discussion of moral ignorance has focused on *what kind of* moral response toward the morally responsible agent is appropriate, rather than on the question of whether the individual is a responsible agent to whom moral responses are appropriate. The defeasible assumption here, then, is that the agents *are* morally responsible, such that some form of moral response is appropriate; the issue is what mode of response is most effective in increasing sensitivity to the norms at issue. The warrant for such moral responses depends upon holding the individual responsible for her failure of sensitivity. These practices provide support for a view that yields the verdict that individuals are responsible for such failures. The scaffolded-responsiveness view generates this verdict.

One might deny that our practices provide support for the scaffolded-responsiveness view, by suggesting that all the observations I have made show

is that the agent must be in general sensitive to such considerations, such that there is a point in engaging her in moral response. The cases I have described do not show that the individual was sensitive to the relevant considerations *at the time at which they behaved problematically*—and so do not support the idea that the agent was at that time morally responsible. There are two responses to this worry. First, the moral responses I outlined above are not properly understood as directed at the agent in recognition of her general sensitivity to moral reasons. They are responses of disappointment, forgiveness, or reproach directed toward the agent *for their failure* to behave as they were expected to. They are ways of holding the individual responsible *for that failure*. These responses are quite different from that of merely pointing out reasons in the hope that the individual will notice them in future. But this latter response seems to be the only one available to the circumstantialist, who, as we have seen, is committed to holding that individuals in such cases are not responsible agents.

The second response returns to the adjudication that the circumstantialist *should* make. If it turns out that holding agents morally responsible in such cases does in fact foster morally sensitive agency—as I suggested above—then the circumstantialist *should* seek an account that also delivers the verdict that the agent targeted by these responses is morally responsible. But the circumstantialist view developed by Vargas does not seem to present us with those resources, since emphasis is placed not on the capacitating relational aspects of responsibility, but on the ways circumstances may impair or erode our capacities. These considerations suggest that it is the scaffolded-responsiveness view that has greater independent plausibility.

Coherence with Practice

There is another dimension of our practices that an independently plausible account of moral responsibility should seek to respect. This feature—and indeed, one of the points of grounding an account of moral responsibility within such social practices—is that, for the most part, we do interact with each other on the assumption that we are morally responsible agents. In the Strawsonian phrase, we occupy “the participant standpoint” (Strawson 1962) in our dealings with each other much of the time. Accordingly, an account that is independently plausible would avoid the determination of widespread skepticism about responsible agency—that is, it would avoid the implication that we lack responsible agency much of the time. Indeed, part of the revisionist project is to save the notion of responsibility from the skepticism oft generated by incompatibilism. This is due to the recognition that responsible agency is valuable, and being able to interact with each other as such is an important—and beneficial—part of human life. Hence Vargas’s aspiration that we should avoid setting the standards

for responsibility such that almost no one is responsible, as we miss out on the benefits of these social practices.

But the circumstantialist view is in fact ill placed to avoid such skepticism and shore up such value. On the one hand, for all we know, we may frequently be in situations of moral ignorance. For much of the time, there may be reasons to which we are not (yet) sensitized. On the other hand, for all we know, there may be many ways in which situational and contextual factors impinge on our agency. If such moral ignorance or situational influences are—as the circumstantialist view prescribes—precisely the sorts of impairments that render us (not excused, but) *not responsible agents*, then for all we know, we lack moral responsibility in relation to some or other moral consideration much of the time. McGeer also presses this challenge when she argues that, for Vargas, “the fact that [an agent] is *pro tanto* blameworthy for failing to respond to [moral] considerations M in [context] C does not mean that she is blameworthy *tout court*” (2015, 2645). This is because there may be a range of other considerations in play in C to which the agent is not sensitive. Accordingly, “circumstantialism actually raises the bar for (*tout court*) blameworthiness” (2015, 2645). But the point holds not just with respect to blameworthiness, but also with respect to responsible agency itself. For the range of considerations in relation to which sensitivity is lacking, the agent may not be a responsible agent if the moral ecology is not yet in place. This is a direct function of having the concept of moral responsibility do the work of exempting agents from responsibility in such circumstances, rather than the notion of excuse, which upholds the agent’s status as morally responsible whilst mitigating blame (see Vargas 2013, 183). This means, as Vargas acknowledges, this account of responsibility is “more targeted” than those views that endorse a generalized capacity for responding to reasons (Vargas 2013, 228). As a result, this targeted approach leaves circumstantialism open to skepticism, insofar as it is eminently likely that there are almost always some moral considerations in a context to which we lack sensitivity. Accordingly, it may never be the case that we are responsible agents *tout court*. This is to countenance a form of skepticism about morally responsible agency that does not cohere with the plausible assumption that we can adopt the “participant standpoint” in our daily interactions; that we can engage with each other as responsible agents—albeit imperfect, often defective responsible agents, with much to learn—much of the time. The scaffolded-responsiveness view—which identifies the ways in which our sensitivity may be incomplete but open to adjustment via sensitivity to others’ viewpoints—seems better placed to capture this aspect of our responsibility practice.

Of course, a proponent of circumstantialism may respond that she too can capture this feature of our practice: any specification of the standards of responsibility that generated such skepticism about responsible agency would indeed

be at odds with the aims of the practice, so the specification should be adjusted so as to permit the judgment that we are responsible agents on more occasions than the skeptical determinations supply. However, once that move is made, it is not clear what remains *circumstantial* about the circumstantialist view. Setting the standard for responsibility in a way that insulates responsibility from pervasive threats from situational context or culturally pervasive ignorance removes the supposed sensitivity to the circumstances of responsibility. I think this is a good move, but it is not one that renders responsible agency *circumstantial*.

The considerations raised in this section suggest that the circumstantialist view does not cohere with valuable aspects of our practice, and they also reveal tensions within the competing aims of the account: to capture the context sensitivity of our capacities, whilst also setting the standard for the relevant capacities guided by the forward-looking aim of the practice. The circumstantialist restricts the scope of responsible agency, and must maintain that there is some discontinuity in the kind of agency implicated in cases of moral ignorance and that displayed in cases where the agent is in fact (according to the circumstantialist) morally responsible. Meanwhile, the scaffolded-responsiveness view is better placed in terms of the independent plausibility of a notion that delivers determinations of responsibility that are consistent with the value of adopting the participant standpoint in our everyday interactions.

Moreover, whilst the scaffolded-responsiveness view is a somewhat revisionary characterization of reasons-responsiveness as understood within the philosophical literature, the view in fact better coheres with the social practices we engage in (cf. McGeer and Pettit 2015, 174). I take it that the examples from the literature on moral ignorance bear out this claim. Insofar as some sort of moral response is engaged (reproach,¹⁷ forgiveness, moral challenge, disappointment), this is because the agent has the capacities, in the scaffolded sense, to be sensitive to these considerations. This consideration, then, provides further support for the scaffolded-responsiveness view, rather than the circumstantialist account of responsible agency.

However, one crucial point on which we should remark is that our attention to these social practices engages a host of moral responses that are not coextensive with blame. How can we maintain that responsible agency is engaged in such cases, if the moral responses involved are not blame? After all, isn't liability to blame simply what it is to be a responsible agent, such that if we are not concerned here with blame, we cannot then be concerned with responsible agency? Both proponents of circumstantialism (Vargas 2013) and scaffolded-responsiveness (McGeer 2015; McGeer and Pettit 2015) seem to endorse what we can refer to as the *liability assumption*: that to be a responsible agent is to be liable to praise or blameworthiness. We see this in the claim that what it is "to blame someone [is] to hold them responsible in a negative sense" (McGeer

and Pettit 2015, 161). Similarly, Vargas proceeds on the assumption that being morally responsible is simply a matter of the agent having the relevant capacities and deploying them in an action that is praiseworthy or blameworthy (2013, 213). Later, it is explicitly held that “we should reserve the phrase ‘is morally responsible’ for cases in which moral praise and blame can arise” (309). If *this* is how we are thinking about moral responsibility—as tightly bound with praise or blameworthiness—then the fact that we find moral responses such as disappointment, moralizing forgiveness, reproach, and moral confrontation in cases of moral ignorance, does nothing to support the claim that there is morally responsible agency present. If this is right, then no support is garnered by the scaffolded-responsiveness view, and the circumstantialist seems better placed to capture this absence of responsible agency. However, in the next section, I point toward strong reasons for rejecting the characterization of moral responsibility as liability to praise or blame.

6 Beyond Blame

The story so far: I have presented some examples of culturally pervasive moral ignorance—about language use, about implicit bias—and considered how circumstantialism and scaffolded-responsiveness views of moral responsibility deal with them. We have arrived at the following possible positions regarding the responsibility of the agents involved:

1. We could maintain that agents in these cases are not responsible (circumstantialism), so not blameworthy (I raised various concerns about this approach);
2. We could maintain that agents in these cases are responsible (scaffolded-responsiveness) and so are blameworthy (given the liability assumption);
3. We could maintain that agents are responsible (scaffolded-responsiveness) but may not be blameworthy (reject the liability assumption).

Given the considerations we have aired so far, I have suggested there are reasons to construe moral responsibility on the model of scaffolded-responsiveness, such that agents in moral ignorance may be morally responsible. I suggest this move is strengthened if we reject the liability assumption. Whilst a full defense of this claim is beyond the scope of this chapter, I suggest that such a move coheres with a view of responsibility grounded in our social practices.

In articulating the regulatory norms of the practice, Vargas focuses on the norms governing praise and blame: indeed, the concept of responsibility itself is set up as serving the purpose of keeping track of deserved praise and blame (2013, 250). Importing a broadly Strawsonian picture within the social

practices, the regulatory norms hold that blame is deserved if a suitably sensitive agent demonstrates a deficit of good will; and praise is deserved when good will is manifested (see Vargas 2013, 250). But is this really an adequate characterization of the use to which we put the notion of moral responsibility? Certainly it is part of the use of the concept. This is a particularly important aspect of the concept from the point of view of metaphysical conceptions of moral responsibility: we must be the kind of agents—ultimately in control, free from deterministic forces—who *deserve* either punitive blame or the rewards of praise (cf. Strawson 1994).

But once we endorse the social thesis and focus on our social practices, the exclusive attention on praise and blame is unwarranted. Rather, my proposal is that whilst deserved praise and blameworthiness are among the regulatory norms of this practice, there is also a range of other moral judgments and responses that do important regulatory work—those implicated where moral failures warrant responses other than blame, for example. These are illustrated by our cases of moral ignorance above. Indeed, it is notable that when P. F. Strawson first introduced the reactive attitudes as constitutive of the social practice of holding responsible, his focus was not limited to praise and blame, but extended to a range of attitudes including resentment, gratitude, indignation, pride, and love. At least, then, it looks like a mistake to preclude, via conceptual fiat, the importance of other kinds of moral response to moral failures. Further, there are good reasons to suppose that these kinds of moral response may be at least as well placed as praise or blame to serve the purpose of cultivating sensitive agency. We should hope for a characterization of responsible agency and moral response that respects these broader possibilities.

How should we make room for these other kinds of moral response? One option would be to see moral responses such as disappointment, reproach, indignation, and so on simply as part of, or subsumed within, a model of blame and the regulatory norms governing its attribution and expression. With this move, the various moral responses are already captured by the liability assumption: that an agent is responsible when she is liable to blame (or praise) broadly construed. But this is obfuscatory on two counts. First, characterizations of blame are often importantly distinct from these other moral responses; supposing that all of them are a species of blame obscures important differences between them. For example, blame may aim to make an individual feel guilty; indignation may not. Blame may aim to cause suffering; disappointment may not (cf. Bennett 2002). Second, this move obfuscates the notion of liability, since there may be different conditions of warrant for these different kinds of moral response. For example, the warrant for blame may require that the agent meets epistemic conditions (she should have known the language was exclusionary); the warrant for disappointment may not (cf. Fricker 2010).

Another option, then, is to see these other modes of moral response as distinct from blame, and filling out our conception of the kinds of moral responses with which we engage each other as morally responsible agents. These other modes of moral response sit alongside blame as ways of holding each other to account (on tokened occasions) and improving our sensitivity to moral considerations (via our engagement in the practice). There are independent reasons for making this move: reflection on our practices does indeed reveal a broader range of moral communication than is captured by simply blame and praise. Some philosophical literature has addressed these modes of moral response, analyzing the structure and role of responses such as indignation, anger, contempt, and disappointment (e.g., Srinivasan 2017; Mason 2003; Fricker 2010; Westlund, chapter 10, this volume). Yet these responses and analyses are ignored by a conception of responsibility that treats responsible agency as constituted by liability to blame (or praise), narrowly construed.

Such a broadening of our moral responses is particularly appealing in combination with a scaffolded-responsiveness view of responsible agency. This view permits that agents may be morally responsible for failures of responsiveness in virtue of their capacity to be sensitized. Where sensitivity is indirect (via prospective audience), liability to a wider range of moral responses than blame may be particularly apt. There is clearly much more to be said here in articulating a fuller picture of these modes of moral response and the conditions for their warrant. My main aim in this section has been to motivate a framework that makes room for such work.

7 Beyond Revisionism about Responsibility

I have been engaged with the following set of theses about our practices of holding each other morally responsible:

- The social thesis
- The justification thesis
- Scaffolded-responsiveness about responsible agency
- A rejection of the liability assumption and an expanded repertoire of moral response.

This set of theses together stand in need of further defense. I want to close by offering some suggestive remarks about the methodological approach we might take whilst developing these theses. My proposal is that they can better be defended if we frame the task of doing so as an explicitly ameliorative analysis

of responsibility, rather than as merely a revisionary project. What distinguishes these two approaches, and why does it matter which we adopt?

Revisionism, as we have seen, starts with “our” concept of moral responsibility, but acknowledges that we may have to revise some aspects of thinking about the concept and embrace certain new implications. But the method of inquiry is essentially conceptual analysis, with a commitment to rationalizing that concept as far as possible. We proceed by testing intuitions about the limit of our concept of responsibility (when someone is or is not responsible, and why), and try to get an analysis that best matches it whilst being both consistent and capable of doing the justificatory work needed. But such a project is essentially bound by a “principle of conservation,” whereby “even though we are entertaining a revisionist proposal, there is still reason to limit the scope of revisions” (Vargas 2013, 103), so as not to depart too far from “our” concept. Such an analysis is therefore beholden to what Kelly McCormick (2013) calls the reference-anchoring problem: the problem of showing that the revised concept of responsibility yielded remains sufficiently close to the folk psychological concept, in order to avoid the charge of changing the subject. Moreover, revisionists face what McCormick calls the normativity-anchoring problem, which challenges the revisionist to explain why we should care about the new concept. Revisionists are skeptical about the force of folk intuitions per se, jettisoning some (incompatibilist ones) in formulating a revisionary account. Then, when we ask the justificatory question of why we should care about the revised concept—in this case, of responsibility-related practices—we find an appeal to the value of a certain form of agency. But the revisionist cannot appeal to intuition in support of this value, for, as we have seen, the revisionist is skeptical about folk intuition. So either revisionists must explain why some intuitions count for more than others, or—absent further argument—we lack adequate justificatory grounds for the revised concept.

I want to propose that an ameliorative analysis of responsibility is preferable and in particular is on a better footing in relation to these problems.¹⁸ Ameliorative analyses do not try to unpack and articulate our concept. Instead, this sort of inquiry is normatively motivated: we start by asking what the legitimate purposes are for which we want and use the concept, and then, having articulated those purposes, we identify which concept we ought to use given those purposes. For example, an ameliorative analysis of moral responsibility could start by asking, do we want the concept to maximize utility? Or to track desert? Or to improve morally sensitive agency? Do we want the concept to be inclusive or elitist? Or do we want the concept at all? These questions are foregrounded in an ameliorative analysis, and the concept that can serve the legitimate purposes is then formulated. As such, there is no need for a principle of conservation, since the aim is not to give an analysis of “our” concept, but to

set out an account of what the concept should be and should be used for. For example, we might identify the legitimate purpose to which we put the notion of moral responsibility as that of cultivating a certain kind of agency, via the means of moralized response.¹⁹ We then ask what concept of responsibility we ought to deploy for this purpose. Such an approach avoids the reference-anchoring problem, since the ameliorative approach asks what our concept *ought to be*, and so it maintains no pretense that the analysis will yield a concept close to the extant folk concept. It is not an attempt at conceptual analysis of “our” concept: an ameliorative approach is explicit that it may change the subject. To the extent that the ameliorative analysis has to deal with questions about how far the analysis departs from our intuitive concept, this is a practical question about whether people could and will take up the newly formulated notion of responsibility, and not simply a worry about “changing the subject.”

Further, conceiving of the project as an ameliorative analysis can also address the normativity-anchoring problem. Such an analysis need not reject (some, or all) intuitions as useful in theorizing. Indeed, proponents of this analysis can maintain that intuition—for example, about the value of certain aspects of our practice—provides important data for conceptual analysis. But the ameliorative project is simply different in kind: it attempts to ask what concept we ought to use for our legitimate purposes (rather than ask what our concept is). Accordingly, adopting this sort of analysis brings no methodological commitment to general skepticism about intuition, and we can rely on intuition as helpful and methodologically respectable in guiding our understanding of what we want the concept to do for us, or in explaining the values that underpin the justification of the practice. Moreover, the ameliorative project needs to show that the concept argued for serves legitimate purposes, and it would have to argue for those purposes as legitimate. In this context, say, we would have to argue for the purpose of cultivating moral-considerations-sensitive agency as legitimate.²⁰ Arguing that this aim is legitimate is in principle consistent with conceptual pluralism; we might have different concepts for different legitimate purposes, subject to feasibility constraints.²¹

These considerations provide some reason for pursuing an ameliorative analysis, rather than a merely revisionary approach. However, there is a further strong reason for adopting this approach: doing so leads us to explicitly reflect on the question of what we want our concept of responsibility for and what work the concept *ought to be* doing for us. If our practices of responsibility are social practices, and thus contingent on the sustaining, critiquing, and reforming of certain social arrangements, it is right that we should carefully scrutinize and ask whether those social practices can be altered or improved. That is, once the social thesis is endorsed, and we move away from a view on which responsibility is simply a matter of certain metaphysical conditions being met, the question of how we want our responsibility practices to be is opened up.

I have suggested that if we want the concept to be doing work for us in practices that aim to cultivate moral-considerations-sensitive agency, then we should endorse the scaffolded-responsiveness account of moral responsibility. This view can serve those purposes by expanding the range of reasons to which we are (directly and indirectly) responsive at any time. This enables us to see morally responsible agency as implicated in a range of cases—such as those of moral ignorance—in which moral responses other than blame are engaged in order to further sensitize agents to the reasons in play. Some of the considerations I raised in defending the scaffolded-responsiveness view indeed drew on features of our practices as we find them. But this view may be best defended, I have suggested, by pursuing an ameliorative analysis. Such an account is better suited to engage both the role of the social practice of holding responsible, and of responsible agency, in contexts in which we often find ourselves, namely those in which our moral understanding is impoverished, but—we hope—improving. And surely this is what we want these social practices for.²²

Notes

1. Note that revisionists face certain challenges: first, they must show that the account remains an analysis of *moral responsibility*. That is, they must show that the revisions made have not ended up changing the subject (McCormick [2013, 5] calls this the reference-anchoring problem). Second, they must show why the account delivered remains one we should still care about: since it does not capture all of our intuitions, does it still capture other important ones—or are there other things that can shore up the value of the practice we are left with (what McCormick [2013, 13] calls the normativity-anchoring problem)? Addressing these two concerns will be aided by facing a further challenge: namely, articulating what the practical implications are of adopting revisionism. For example, Vargas notes that if revisionism requires us to give up incompatibilist intuitions (intuitions that favor alternative possibilities, or require ultimate control, say), “there may be isolated pockets of our practice that . . . cannot continue as before” (2015, 2623). However, Vargas does not articulate what those parts that we may have to jettison would be. Yet my sense is that it will be difficult to fully meet the first two challenges before we have a good grasp on exactly how our responsibility practices themselves (rather than our concept) may change. I address these methodological issues in the final section.
2. We might worry about the fairness of Vargas’s objections to Strawsonian psychologism, critiquing the practice as a whole (our dispositions, as humans, to engage in certain kinds of practice) by using a case that focuses on an individual’s psychological dispositions (the example of jealous Dave, who has certain fixed psychological dispositions to feel jealous of his partner’s other friendships, but comes to question the justifiability of the expression of these feelings [Vargas 2013, 162]). Just because individuals’ dispositions are explanatory rather than justificatory, this does not mean that the set of dispositions as a feature of agents engaged in a practice has no justificatory power. Sommers (2013) has also raised this concern.
3. There are important differences between Vargas’s model, which focuses on the effects of the system as a whole (i.e., a system of norms that regulate praise and blame), and McGeer’s model, which focuses rather on the cultivation of sensitivity in particular tokens of moral response (such as blame or praise).
4. Vargas wants to resist framing this in terms of the agent’s capacities remaining the same, but exercising them becoming more difficult—since this just pushes the issue back a step: our

abilities to exercise those capacities are not cross-situationally stable, even if the capacities themselves are (2013, 208). I actually think that “pushing the issue” back a step matters a great deal, for reasons articulated in section 5—namely, it allows us to maintain the claim that the individual is a responsible agent, albeit an agent who did not exercise her capacity to recognize reasons in that context.

5. Vargas spells out the details of how sensitive agents must be to qualify as responsible in counterfactual terms, which concern the frequency (“a suitable proportion of worlds” [2013, 214]) with which one would respond to moral considerations in relevantly similar contexts. Details are needed about this frequency (what proportion of possible worlds?) and about the description of contexts that are relevantly similar (coarse- or fine-grained descriptions of morally relevant considerations and contexts?). These matters are settled by looking to what proportion of possible worlds, described at what degree of fine- or coarse-grainedness up for discussion, is at least as good as any other in fostering moral-considerations-sensitive agency (2013, 219).
6. Of course, this is a difficult question for us to answer. Vargas relies on an ideal observer to determine this. It is not clear to me how this part of the picture integrates with an account whereby responsibility is constituted by our social practices, but I set this worry aside for another occasion.
7. McGeer and Pettit (2015) develop this model to explain how we should understand blame: as a sanction that enforces a specific or general injunction against inattentiveness to reasons, and that serves as a post hoc exhortation to be attentive. I am not persuaded by this aspect of their account, but it is not entailed by the scaffolded-responsiveness view, so we can regard these parts of the account as separable.
8. A related issue concerns the remarks about one’s audience being “authorized” as a source of moral guidance or challenge (McGeer and Pettit 2015, 172), and whether this includes those one in fact views as authorized, or those one *should* authorize qua members of the moral community. The issue here and in the text is ultimately how much idealization is built into the conception of audience in the account.
9. I am grateful to Marina Oshana for helping me to pinpoint this dimension of difference between the two accounts. For an alternative way of formulating the relational aspects of responsibility, see Westlund, chapter 10, this volume.
10. On their view, the sanctioning component of the blame serves to fulfil the promised threat that accompanies the exhortation “attend to reasons or else” (McGeer and Pettit 2015, 183–185). I am not persuaded by this part of their view, but it seems to me inessential to the scaffolded view of responsible agency. I offer an alternative way of thinking about moral response, in section 5.6. Note also that this feature of McGeer and Pettit’s view applies to particular tokens of blame, rather than the practices as a whole, which is a point at which their justification departs from Vargas’s.
11. Or indeed, any notion of responsible agency—there may be others that fare better than either of the two available. I consider this the start of our inquiry, which may end up with a different account yet.
12. Other prominent examples in the literature on moral ignorance include Michele Moody-Adams’s (1994) discussion of the alleged difficulty of recognizing the wrong of slavery in cultures in which slave ownership was common. Another, from Miranda Fricker’s (2010) work, is that of a schoolteacher whose cultural context makes it difficult to grasp the moral insight that corporal punishment is wrong. I focus on Calhoun’s case because it is (I assume) more proximal to the cultural context in which readers are located. This is not to say there are not many other analogous cases to which similar thoughts may apply (exploitation of cheap labor, climate pollution, the food industry, and so on and so on).
13. See Isaacs 1997; Calhoun 1989; Moody-Adams 1994; Fricker 2010; Saul 2013; Holroyd 2012; Brownstein 2016; and Washington and Kelly 2016.
14. Vargas writes, “[I]f the threat to a responsibility ascription operates via some threat to the normal capacity to respond to moral considerations (as seems more likely in cases of new, unusual, or particularly challenging contexts of action), then the issue is how these concerns are accommodated internal to some account of the capacities required for being subject

- to the responsibility norms. In other words, we look to our theory of responsible agency” (2013, 183).
15. Note that it seems not at all clear that we should say similar things in the more standard case of moral ignorance: that holding responsible for gendered language is inappropriate since it provokes hostility, prevents buy-in to the norms against gendered language, and that the knowledge is not in place to permit sufficient sensitivity. This, for me, provides additional reason for doubting the circumstantialist view.
 16. Though note that Fricker is concerned with cases in which we are at a historical distance from the wrongdoers, such that it may no longer be possible to capacitate them.
 17. Note, however, Calhoun’s detachment of the justification of reproach (responsibility) and its point (moral change).
 18. For a statement of ameliorative analyses, see Haslanger 2000. For discussion of the contours of such an inquiry see Burgess and Plunkett, 2013a, 2013b.
 19. Compare a consequentialist ameliorative analysis, which might specify the legitimate purpose as maximizing utility and specify a concept of responsibility that serves this purpose: blame-worthy if susceptible to moral influence.
 20. Compare Haslanger’s (2000) ameliorative analysis of gender, which argues that ending gender inequality is the legitimate aim.
 21. See Saul 2006; Jenkins 2016.
 22. I am grateful for extremely useful feedback from the editors of this volume, and for constructive discussion with Federico Picinali, Yonatan Shemmer, Manuel Vargas, and audiences at Macquarie University and the University of Sheffield.

References

- Bennett, C. 2002. “The Varieties of Retributive Experience.” *The Philosophical Quarterly* 52 (207): 145–163.
- Burgess, A., and D. Plunkett. 2013a. “Conceptual Ethics I.” *Philosophy Compass* 8 (12): 1091–1101.
- Burgess, A., and D. Plunkett. 2013b. “Conceptual Ethics II.” *Philosophy Compass* 8 (12): 1102–1110.
- Brownstein, M. 2016. “Attributionism and Moral Responsibility for Implicit Bias.” *Review of Philosophy and Psychology* 7 (4): 765–786.
- Calhoun, C. 1989. “Responsibility and Reproach.” *Ethics* 99 (2): 389–406.
- Czopp, A. M., M. J. Monteith, and A. Y. Mark. 2006. “Standing Up for a Change: Reducing Bias through Interpersonal Confrontation.” *Journal of Personality and Social Psychology* 90 (5): 784–803.
- Duguid, M. M., and M. C. Thomas-Hunt. 2015. “Condoning Stereotyping? How Awareness of Stereotyping Prevalence Impacts Expression of Stereotypes.” *Journal of Applied Psychology* 100 (2): 343–359.
- Fischer, J. M., and M. Ravizza. 2000. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Fricker, M. 2010. “I—The Relativism of Blame and Williams’s Relativism of Distance.” In *Aristotelian Society Supplementary Volume* 84 (1): 151–177.
- Jenkins, K. 2016. “Amelioration and Inclusion: Gender Identity and the Concept of Woman.” *Ethics* 126 (2): 394–421.
- Holroyd, J. 2012. “Responsibility for Implicit Bias.” *Journal of Social Philosophy* 43 (3): 274–306.
- Haslanger, S. 2000. “Gender and Race: (What) Are They? (What) Do We Want Them to Be?” *Noûs* 34 (1): 31–55.
- Isaacs, T. 1997. “Cultural Context and Moral Responsibility.” *Ethics* 107 (4): 670–684.
- Mason, M. 2003. “Contempt as a Moral Attitude.” *Ethics* 113 (2): 234–272.
- McCormick, K. 2013. “Anchoring a Revisionist Account of Moral Responsibility.” *Journal of Ethics & Social Philosophy* 7 (3): 1–19.

- McGeer, V. 2015. "Building a Better Theory of Responsibility." *Philosophical Studies* 172 (10): 2635–2649.
- McGeer, V., and P. Pettit 2015. "The Hard Problem of Responsibility." In *Oxford Studies in Agency and Responsibility*, Vol. 3, edited by D. Shoemaker, 160–188. Oxford: Oxford University Press.
- Moody-Adams, M. M. 1994. "Culture, Responsibility, and Affected Ignorance." *Ethics* 104 (2): 291–309.
- Saul, J. 2006. "Philosophical Analysis and Social Kinds: Gender and Race." In *Aristotelian Society Supplementary Volume* 80 (1): 119–143.
- Saul, J. 2013. "Implicit Bias, Stereotype Threat, and Women in Philosophy." In *Women in Philosophy: What Needs to Change?*, edited by K. Hutchison and F. Jenkins, 39–60. New York: Oxford University Press.
- Scaife, R., T. Stafford, A. Bunge, and J. Holroyd. "The Effects of Moral Interactions on Implicit Bias." Unpublished manuscript.
- Smart, J. J. C. 1973. "An Outline of a Utilitarian System of Ethics." In *Utilitarianism: For and Against*, edited by J. J. C. Smart and B. Williams, 3–76. Cambridge: Cambridge University Press.
- Sommers, T. 2013. Review of *Building Better Beings: A Theory of Moral Responsibility*, by Manuel Vargas. *Notre Dame Philosophical Reviews*. <https://ndpr.nd.edu/news/42827-building-better-beings-a-theory-of-moral-responsibility/>.
- Srinivasan, A. 2017. "The Aptness of Anger." *The Journal of Political Philosophy*. doi: 10.1111/jopp.12130.
- Strawson, G. 1994. "The Impossibility of Moral Responsibility." *Philosophical Studies* 75 (1): 5–24.
- Strawson, P. F. 1962. "Freedom and Resentment." *Proceedings of the British Academy* 48: 1–25.
- Vargas, M. 2013. *Building Better Beings: A Theory of Moral Responsibility*. Oxford: Oxford University Press.
- Vargas, M. 2015. "Précis of *Building Better Beings: A Theory of Moral Responsibility*." *Philosophical Studies* 172 (10): 2621–2623.
- Vargas, M. 2017. "Implicit Bias, Responsibility and Moral Ecology." In *Oxford Studies in Agency and Responsibility*, Vol. 4, edited by David Shoemaker, 219–247. Oxford: Oxford University Press.
- Washington, N., and D. Kelly. 2016. "Who Is Responsible for This?" In *Implicit Bias and Philosophy*, Vol. 2, edited by M. Brownstein and J. Saul, 11–36. New York: Oxford University Press.