

Power, Social Inequities, and the Conversational Theory of Moral Responsibility

MICHAEL MCKENNA

We should think of the many different kinds of relationships which we can have with other people—as sharers of a common interest; as members of the same family; as colleagues; as friends; as lovers; as chance parties to an enormous range of transactions and encounters. Then we should think, in each of these connections in turn, and in others, of the kind of importance we attach to the attitudes and intentions towards us of those who stand in these relationships to us, and of the kinds of *reactive* attitudes and feelings to which we ourselves are prone. In general, we demand some degree of goodwill or regard on the part of those who stand in these relationships to us, though the forms we require it to take vary widely in different connections. The range and intensity of our *reactive* attitudes towards goodwill, its absence or its opposite vary no less widely.

—P. F. Strawson, “Freedom and Resentment”

1 Strawsonian Theories of Responsibility and the (Dubious?) Demand for Good Will

How should we understand the social and relational dimensions of moral responsibility?¹ Understood in one way, this question has an obvious answer: sociality is essential to moral responsibility’s nature. So it is for P. F. Strawson (1962) and numerous others embracing a broadly Strawsonian approach to theorizing about moral responsibility.² On this approach, moral responsibility is essentially interpersonal because being responsible is conceptually connected to holding responsible, which in turn is understood in terms of social practices. Hence, responsibility turns out to be a deeply social phenomenon. This is

intimated in the well-known passage from Strawson's seminal paper "Freedom and Resentment" quoted above.

On Strawson's own proposal, *being* morally responsible in the sense of being either praiseworthy or blameworthy is most fundamentally about the quality of will with which an agent acts—whether it is from good or ill will. How is this thesis conceptually linked with considerations of sociality? Here is how: standards for a competent agent's acting from a reasonable quality of will are understood by reference to the expectations of the moral community positioned to hold responsible. As indicated in the preceding quotation, the community makes the moral demand that agents show a sufficient degree of good will. Deviations from these demands then render fitting reactive responses such as praise or blame in the form of outwardly manifested emotions like gratitude or resentment. In this way, the fittingness of praising and blaming in social practice signals the scope of praiseworthiness and blameworthiness. One involves surpassing the standards set down by the basic demand for good will; the other involves falling below them.

In this essay, I will be concerned with the requirements of sociality in light of the Strawsonian enterprise.³ More precisely, like other contributors to this volume, I will focus upon a strikingly neglected topic within the literature on free will and moral responsibility: the social and relational dimensions of moral responsibility in light of both social inequities and asymmetrical relations of power. How do, and, perhaps more importantly, how *should* these factors influence our moral responsibility practices and judgments? By drawing upon my Strawson-inspired conversational theory of moral responsibility, I intend to expose a rather unseemly dimension of our moral responsibility practices. To foreshadow what will come, I call attention to Strawson's remark, included in the opening quotation above, that in general we demand some degree of good will or regard on the part of those with whom we are variously related. This widely shared assumption is often heralded as one of the deepest and most celebrated insights in Strawson's paper (e.g., Watson 2014). Of course, in a certain respect it is undoubtedly true, given that it is qualified with 'in general.' Nevertheless, it masks much that is a source of worry, and so perhaps should not be celebrated without caution. After all, not all *do* demand a reasonable degree of good will from pertinent others. For some, there is little point in doing so because there is little reason to expect that their demands will even register at all among those whose authority is liable to have great sway over them. For many others, the forms good will can be expected to take will be settled by standards they have no part in shaping.

In what follows, I will focus primarily upon disparities regarding the forms that good and ill will can be expected to take. By drawing upon my proposed



Strawson-inspired conversational theory of moral responsibility, I will argue that many agents who are morally responsible for their conduct, even when they do act from a reasonable degree of good will, operate in contexts that are morally unfair to them as agents. Or, at any rate, there is something morally suspect about the social conditions facilitating exercises of their agency when they act in ways that are morally praiseworthy (and also morally blameworthy). This is because, as the conversational theory reveals, quality of will is to be identified and explained by a community of interpreters who take some kinds of actions as indicative of good will and other kinds as indicative of lack of good will. Since some in this community are socially empowered, in contrast with others who are socially disempowered, the conditions for what signals good and ill will are liable to arise from potentially unjust social circumstances.

2 The Conditions for Moral Responsibility

I turn to the conditions for moral responsibility. Considering these will help assess the influences on moral responsibility of asymmetric relations of power and various forms of social inequities. I will assume in what follows that we can safely set aside worries about free will and moral responsibility skepticism. Perhaps, contrary to my own position, skeptics like Derk Pereboom (2014) and Galen Strawson (1986) are correct that because no one has free will, no one is morally responsible (in an important sense). I leave that as an open question. But the issue before us is about potential effects on responsibility due to *special* conditions of our sociality. Hence, it will be most useful to assume free will and moral responsibility realism and then proceed by raising questions about unique social and relational conditions in which these can be attenuated, compromised, or extinguished.⁴

Some Strawsonians hope to exhaust the conditions for moral responsibility by focusing upon quality of will alone, or instead quality of will and a general capacity for engaging in adult interpersonal relationships.⁵ My own view is that this attempts to do too much with too few resources. I prefer a mixed theory. Moral responsibility as I understand it requires at least two conditions that are neutral between Strawsonian interpersonal theories and other competitor proposals: a *control condition* and an *epistemic condition*.⁶ These conditions apply both to the conditions for being a morally responsible agent—which concerns the *status* of some but not all persons—and also to the conditions in which a person who *is* a morally responsible agent is morally responsible *for* something, such as an action, an omission, or the consequences of one of these things.

Attending to the control and the epistemic conditions gives rise to several interesting philosophical issues regarding the influence of significant social



inequalities and asymmetric relations of power. The control condition (which can plausibly be understood as a free will condition) invites worries about coercion and duress, especially in contexts of interactions between members from disparate groups whose social status and power are asymmetrically distributed. A further and more subtle possibility here concerns domination in the sense Pettit (2001) intends it. In such cases, a person's options might be constrained in light of the mere prospects that others are easily able to exercise domination over that person with no repercussions. Hence the person self-regulates in ways that are liable to diminish her freedom.⁷ These factors can sometimes serve as grounds for mitigating or instead excusing otherwise blameworthy behavior on the part of those who are disadvantaged by the pertinent disparities. Naturally, they can also give rise to *greater* degrees of freedom for those advantaged by these conditions, and so to a potentially greater range of conduct for which one ought to be regarded as accountable.

In the case of the epistemic condition, the social and relational variables at issue raise worries about disproportionate access to education, complex sources of information, or for that matter the opportunity to attend to matters of moral and political concern with sufficient care—people working three jobs to make ends meet, for instance, can perhaps be excused for failing to stay abreast of efforts to properly regulate the banking industry. Hence, some appeals to ignorance or limited understanding as grounds for mitigation or excuse are far more compelling than others.⁸ Those massively disenfranchised and without the power to guide the aim of various inquiries will be far more likely to *merit* some form of mitigation or excuse in certain contexts, although it should be noted that, ironically, they are often the same group of people far less likely actually to *be excused*.⁹ Likewise, of course, those with *greater* access to epistemic resources are liable to bear a greater degree of moral responsibility for their conduct since the charge “You should have known better” will apply more liberally.

The considerations mentioned in the two preceding paragraphs provide ample resources for exploring the social and relational dimensions of moral responsibility in light of extreme disparities in relations of power and significant social inequities. What insights might we be able to draw specifically from my conversational theory of moral responsibility? Again, my goal will be to attend to the Strawsonian demand for a reasonable degree of good will.

Some Strawsonians apparently think that a quality of will condition on moral responsibility is encompassed by the control and the epistemic condition already identified.¹⁰ When, for instance, an agent knowingly and freely does morally wrong, this alone, some might contend, is sufficient for her having a morally objectionable quality of will and so being morally blameworthy. My own view (McKenna 2012), which I grant is controversial, is that quality of will is a *further* condition over and above the control and epistemic conditions. In



the case of blameworthiness, what is involved is either an ill will or an insufficiently good will.¹¹ Nothing much turns on settling this issue here. Either way, on a Strawsonian theory, when an agent is morally responsible for some action, she *does* satisfy a quality of will condition—either because the control and epistemic conditions already ensure it, or because on a view such as mine a distinct quality of will condition is also satisfied. Moreover, for Strawsonians, it is quality of will that is most salient in settling questions of an agent's responsibility. And quality of will, as explained above, is linked to the moral community since it is the community who, by making demands, sets the bar for what will count as sufficient or insufficient quality of will.

Quality of will, as I understand it (McKenna 2012), is a matter of the value of an agent's regard for others.¹² When we hold morally responsible, what we are reacting to, or what we are prepared to react to, is an agent's quality of will as manifested in her conduct. Eventually, I will further develop the quality of will condition in terms of my conversational theory of moral responsibility. However, before doing so, I turn in the next section to an especially illuminating application of the epistemic condition in evaluating cases of putative blameworthiness. Are those ignorant of largely unrecognized moral wrongs due to conditions of oppression blameworthy for participating in and benefitting from those conditions? Is their ignorance exculpating? This will help shed a light on the particular issue I wish to consider.

3 Excusing Morally Ignorant Oppressors?

In her superbly argued “Responsibility and Reproach,” Cheshire Calhoun (1989) confronts a challenge for feminists faced with questions about how to respond to social oppression when carried out or at least perpetuated by seemingly innocent participants to the oppressive practices. Is reproach warranted in such cases? If so, it seems that these apparently innocent participants are not innocent after all but instead blameworthy. If not, it seems that feminists are not entitled to “use moral reproach as a tool for effecting social change” (1989, 389). Each horn comes at quite a cost.

Consider the first horn—that the seemingly innocent are blameworthy after all. As Calhoun rightly notes, many individuals who are, as she puts it, “morally unflawed” (1989, 389) unknowingly commit wrongdoing through participating in oppressive social practices that strongly disfavor women.¹³ Yet they do so by means that appear to arise from nonculpable ignorance. How so? In contexts in which some moral knowledge is not widely shared, which Calhoun calls “abnormal moral contexts” (396), it is not just the unsavory, like pimps and other misogynists, who are prone to harm women by means of various oppressive



practices, but many well-meaning men (and women too, of course). Calhoun mentions parenthetically “male bias in psychological and other theories, the design of female fashion, the use of ‘he’ neutrally, [and] heterosexual marriage” (397). Even if one wished to dispute any of these as significant moral wrongs, it is beyond question that various shared social practices regarded widely by many as morally neutral do after all contribute to morally wrong forms of oppression that undermine women’s autonomy. To think that reproach *is* warranted even for those Calhoun refers to as “morally unflawed” is, it seems, to think that these moral agents are after all culpable for engaging in what they take to be perfectly innocent behavior. But how could that be? We cannot expect even the most virtuous among us to have access to all the best information about morality. Presumably, we are all vulnerable to moral blind spots at the limits of our current social and cultural settings.

Now consider the other horn—that in these contexts the seemingly innocent are indeed as they seem to be, innocent. If so, it appears that feminists are not entitled to reproach—by which I take it Calhoun means moralized blame. The cost of thinking that reproach *is not* warranted because these parties are not blameworthy means that feminists must let too much pass. The oppression identified is, after all, deeply harmful and far-reaching. Failing to reproach, Calhoun argues, comes dangerously close to endorsement and so to participating in the modes of oppression feminists correctly see as deep forms of moral wrongdoing.

Calhoun offers an elegant solution to this dilemma. She distinguishes between reasons to blame in light of blameworthiness, and reasons to blame as a means of effecting social change. A person could be entitled to blame those who are not blameworthy, say in the context of excused wrongdoing, given *other* powerful moral considerations. A significant need to achieve social reform is such a reason, according to Calhoun, and, moreover, failing to reproach is highly likely to signal endorsement or sanctioning of the moral order (1989, 400–405). (Of course, it should be clear that the reasoning at issue here applies to other cases as well, like that at issue in claims of white privilege or class privilege.)

Some might be inclined to downplay Calhoun’s proposal as perhaps correct but no more than a pedestrian application of a simple point: blameworthiness merely provides a *pro tanto* reason to blame. All blameworthiness ever establishes is that a wrongdoer is deserving of something, blame—and all desert ever provides is a *pro tanto* reason for some sort of response. But there can be reasons other than desert for treating a person a certain way, and these reasons can override the *pro tanto* reasons. Or these other reasons can instead do duty to justify a treatment like blame when the desert-based reasons are altogether absent. For instance, young children who are not yet morally responsible agents are not yet blameworthy for their conduct. Hence, they do not deserve blame—granting the near truism that only the blameworthy deserve blame. But we might have

good moral reason to blame them as a means of training them up into becoming morally responsible agents who will eventually deserve blame (and praise). So, too, it can be argued, we can have reasons to blame those who are not blameworthy for participating in oppressive practices because it is in the service of an overriding social good. While this might be true, there is nothing especially philosophically illuminating going on here.

To downplay Calhoun's proposal in this way would be to misunderstand the striking philosophical point she brings into relief. Her argument, as I understand it, is not *simply* about overriding reasons to ignore considerations about what a nonculpable wrongdoer does not deserve (blame). It is, rather, about the *special social setting* giving rise to those overriding reasons and how that social setting shapes not only the moral landscape but our moral responsibility practices as modes of responding to that moral landscape. (It is this point in particular that I wish to draw upon in my own proposal as I develop it below.) As Calhoun puts it, the sort of moral wrongdoing at issue in these kinds of cases occurs at the social rather than at the individual level (1989, 394). This gives rise to a special class of moral reasons that can then be “turned back on” our responsibility practices in the hope of refashioning them.¹⁴

Think about it this way. The architecture of our responsibility practices is, so to speak, built up out of a stock of accessible bits of moral knowledge. Good moral reasoners, being fallible creatures and so lacking moral omniscience, rely upon this stock as a resource for engaging with others in ways that can give rise to blameworthy as well as praiseworthy behavior. But if moral knowledge, like all knowledge, depends upon background social conditions rendering pertinent truths accessible or instead inaccessible, then some forms of moral knowledge will sometimes be, so to speak, outside the reasonable scope of even a well-meaning person's epistemic radar.¹⁵ Of course, perhaps a few elites at the fringes of moral knowledge might be better positioned, but their understanding cannot be expected to be accessible to most others.¹⁶ As such, the moral landscape will be affected insofar as even morally virtuous agents will be liable to participate unwittingly in wrongdoing.¹⁷ And, moreover, because of this, our practices of holding morally responsible will be likewise affected insofar as this sort of epistemic limitation will do two things. First, it will provide legitimate grounds to claim nonculpable ignorance as a basis for excusing conduct that contributes to oppression. Second, it will limit the moral community's resources and so impede the moral community from deploying the very machinery of our blaming practices that could be used to correct these sorts of wrongs.

These special features of our social setting and the pressure they place on our own practices of holding morally responsible provide reasons for justifying blame even when directed at those whose wrongdoing, due to epistemic considerations, is excusable—that is, even at those who do not deserve blame. Doing

so allows for the possibility of moral reform. This, I take it, is precisely an application of what Manuel Vargas (2013) has in mind in writing of “building better beings.” We modulate our responsibility practices with the aim of encouraging people to be more alive to moral reasons. By blaming those wrongdoers who are not blameworthy in these contexts, we thereby allow for the possibility of making accessible this sort of moral knowledge. *We bring it to light.* This in turn can help refashion our responsibility practices in such a way that those failing to act properly in light of this moral knowledge would no longer be able to claim nonculpable ignorance. Then they could after all be regarded as blameworthy and so not excused. There is, in this sense, a kind “feedback” loop that helps us bootstrap our moral responsibility practices to be better positioned to hold accountable those who unwittingly do engage in wrongdoing that contributes to oppression.

4 The Role of the Conversational Theory and the Demand for Good Will

In the next section, I draw upon Calhoun’s proposal to help cast a critical light of a different sort on our moral responsibility practices. In this section, I first set out only briefly my conversational theory as a natural way of extending a Strawsonian account of moral responsibility. As explained above (section 1), on my view, quality of will as manifested in an agent’s conduct is a matter of the value of the regard or concern an agent displays for others. The quality of this regard is in turn evaluated in terms of a moral community’s implicit demand that co-members display a sufficient degree of good will toward others. Perceived departures are then taken to be grounds for a fitting response in the form of a reactive attitude of some sort. In the case of blameworthiness, this is a matter of responding with resentment or indignation. Overt manifestations of these reactive emotions, when directed at a blamed party, have, as Gary Watson (1987) has noted, not only an expressive but a communicative role; they serve to communicate our moral demands and expectations.

According to my conversational theory, and drawing upon Watson’s proposal, I have argued (McKenna 2012) that moral responsibility is not only essentially interpersonal and communicative, but that it also has a *conversational* dimension. This is crucial. On the conversational theory, an agent’s actions—those that are candidates for blameworthiness or praiseworthiness—are potential bearers of meaning, where meaning is a function of the quality of an agent’s will. This meaning is analogous to the meaning a competent speaker conveys when she engages in conversation. I call this *agent meaning*. Like speaker meaning, agent meaning can be affected by the interpretive framework whereby others



interpret the agent. In the case of agent meaning, a moral community assigns saliences to types of actions, and they do so in light of expectations about the cooperative constraints of something analogous to a conversational transaction. Hence, the interpretive community expects that the agents whose actions they interpret will have uptake with respect to the interpretations the community employs. They therefore expect that agents will be able to modulate their behavior in light of those interpretations. For instance, a type of action liable to be interpreted as manifesting poor quality of will gives reason for competent moral agents to avoid signaling such quality of will by acting in this type of way. By emerging conventions, types of conduct come to indicate appropriate or instead inappropriate quality of will. A competent moral agent thus acts in a social context wherein these interpretive pressures are liable to affect her judgments about what does and does not signal good or ill will and so what might or might not be communicated to others who stand prepared to hold her to account by way of praising and blaming practices.

The same applies to the practices associated with various modes of blaming and praising. Consider overt manifestations of reactive emotions like resentment or indignation directed at a blamed party. These modes of expressing emotions involve modulating one's behavior against the backdrop of otherwise expected social interactions, interactions unfolding under the assumption that the basic demand for good will has been met. A blamed person is, for instance, excluded from an otherwise routine lunch outing, or is greeted more coolly in simple exchanges, or directly corrected for some misdeed with angry rather than kind treatment. All of these are modes of interacting by way of altering what would otherwise be regarded as treatment of one who *does* show reasonable good will. In short, just as the actions for which we are accountable have meaning, so too the means of blaming or praising through outward manifestations of reactive emotions have meaning. And in each case, this sort of meaning is shaped in part by a background set of loose interpretive conventions against which instances are evaluated.

One more point about the preceding sketch of the conversational theory. Grice (1957) distinguished speaker from sentence meaning. When we interpret a speaker in actual conversation, our interpretive goal is to understand what the speaker meant by what she said, not just what the sentences mean that the speaker used to convey what she meant. Often these go together, but they can come apart. We can, for instance, use the sentence ‘That was a good meal’ to mean the meal was terrible by making use of sarcastic cues and the like and relying upon our audience to cooperatively receive our expression as we intended it. So, while conventions regarding sentence-types help shape what we mean to say, they do so in ways consistent with variation that allows for particularizing our own intended meanings and departing from the conventional meanings



assigned to sentence-types. An analogous point applies to interpreting the actions of morally responsible agents, and it *also* applies to interpreting the praising and blaming responses of those holding responsible when they react to those blamed (or praised). A shove, for instance, might characteristically show ill will, but in some contexts could be taken to express solidarity or playful aggression (say on a basketball court). Why is this important for what is to come? Because our ability to function in an interpretive space where others are equipped to see us as acting with good will, and our ability to communicate to them our reactive assessment of them when they do not, rely upon an expectation that we share enough interpretive resources to facilitate successful communication. And one kind of resource is just social authority—being positioned to have one's interpretive scheme do the work in settling meaning.¹⁸ This resource, as it happens, is usually not evenly distributed.

Given this very simplified description of the Strawsonian project and my proposed conversational theory as an extension of it, return now to the passage from Strawson with which I started. I repeat it here for ease of reference:

We should think of the many different kinds of relationships which we can have with other people—as sharers of a common interest; as members of the same family; as colleagues; as friends; as lovers; as chance parties to an enormous range of transactions and encounters. Then we should think, in each of these connections in turn, and in others, of the kind of importance we attach to the attitudes and intentions towards us of those who stand in these relationships to us, and of the kinds of *reactive* attitudes and feelings to which we ourselves are prone. In general, we demand some degree of goodwill or regard on the part of those who stand in these relationships to us, though the forms we require it to take vary widely in different connections. The range and intensity of our *reactive* attitudes towards goodwill, its absence or its opposite vary no less widely. P. F. Strawson “Freedom and Resentment”

Note Strawson's observation about the variability and particularity of our relationships. He focuses on how much we care about how others regard us in their interactions with us (their quality of will) and our demand for good will on their behalf. This extends to our reactions to those who fail to meet the demand when we hold them responsible via a reactive attitude. Note also his observation that the forms we require the demand for good will to take vary widely (compare with the preceding point about the particularity of speaker meaning in relation to generic sentence meaning). Now let us inject into this picture a fact about nonideal human communities. Those who make the demands for a reasonable degree of good will and those prepared to react to departures from those



demands are not all equally empowered. If we add familiar facts about significant social inequalities and asymmetric relations of power, we bring to light a set of considerations that are not necessarily so flattering to a Strawsonian understanding of our responsibility practices.

5 Something Insidious Rooted in Our Responsibility Practices?

So now, looking through a critical lens, focus yet again on Strawson's remark that in general we demand some degree of good will or regard on the part of those who stand in varying relationships with us. As it happens, it is a contingent albeit inescapable fact that some who express a demand for good will are taken to have an authority that others lack. As I noted above (section 1), there are some who take their own social status to be so limited that they do not demand that they are shown a reasonable degree of good will—at least in relation to members of comparatively advantaged groups. In some cases, they do not see any point in even making the demand. Others do, but they have no part in settling the interpretive standards regarding what plausibly counts as expressions of good will.¹⁹ In this way they are, in a certain sense, outsiders to those positioned to exercise a greater degree of moral authority.

Consider, for instance, etiquette or manners. While standards of etiquette or manners are usually not directly regarded as relevant to morality (however, see Buss 1999), drawing upon these conventions in how one comports oneself can often be a *vehicle* for showing or instead failing to show deference or respect for others, which is a matter of morality. In this way, especially as understood through the lens of the conversational theory, superficial conventions of etiquette can function as a way of manifesting morally significant behavior. However, the social conventions giving rise to what counts as polite behavior—such as pausing to hold a door open for someone, refraining from interrupting another who is speaking in certain social contexts, how one behaves while dining in certain settings rather than others, the cutlery one uses, when one begins to eat, and so on—are all established by social contexts wherein those empowered set the expectations. Departures from expected behavior can be taken to show a lack of respect or concern for the feelings of others—such as one's hosts. The upshot is that those who act in such contexts have available to them resources for displaying good or ill will, but those resources are themselves structured and constrained by certain groups empowered to set expectations and “police” departures from expected behavior. Of course, if this were limited just to matters of etiquette, there would be little interest here. But my contention is that these interpretive pressures on the (conversational) context of action are ubiquitous.



They pervade nearly every aspect of our social lives. Those marginalized by existing power-structures—cultural, social, and economic—live out their lives shouldering the burdens of acting in a context in which much of the interpretive framework signaling what counts as constituting good or ill will is settled by others whose social lives are in some way alien, inaccessible, or unwelcoming to them.

Return to Calhoun's insight. Special social contexts can provide reasons to blame some who are not in fact blameworthy in order to achieve social change. Her focus was on social structures involving massive disparities in social advantage that supported forms of wrongdoing: contributing to oppressive social conditions through forms of moral ignorance that were exculpating. These social structures helped support the *perpetuation* of the oppressive conditions by deploying the very responsibility practices that, were they exercised by better-informed moral agents, would have instead functioned as a tonic in correcting those oppressive conditions. Calhoun's proposal shows us how to alter our resolution and look upon our responsibility practices critically so as to determine whether their design is in certain respects deleterious. This is what I intend to do in this section.

There is, however, a significant difference between the cases I wish to focus upon and the ones Calhoun was interested in. Calhoun focused on cases of moral *wrongdoing*. She argued that in these special sorts of contexts moral ignorance really was excusing. (She then argued that reproach could be warranted anyway—a conclusion I agree with.) My interest is different. I am interested primarily, albeit not exclusively, in cases where well-intentioned moral agents, those Calhoun would describe as morally unflawed (I would prefer a different term), do *not* engage in moral wrongdoing at all. Instead, they act well. They do right. In doing so, moreover, they act in ways that *do* show adequate moral regard for others. These agents *do* meet reasonable demands for good will. Indeed, they might even be regarded as praiseworthy. My claim is that in many of these cases, those who act well often do so in contexts in which the resources for interpreting their quality of will are framed by the interests of others—sometimes, even often, quite innocently so. (An example is forthcoming in the next paragraph.) Yet these framings have power insofar as a comparatively socially disadvantaged yet fully competent moral agent, in acting from good will, will be responsive to that interpretive framework. The standards the empowered set for signaling compliance with the demand for good will are shaped by those whose status is in some manner or other dominating. And this can happen, as in the sorts of cases Calhoun has in mind, even when the parties in positions of domination are all well-meaning and innocent of any wrongdoing. In such cases, the parties involved need not be individually culpable for any particular conditions resulting in these asymmetric social conditions.

To return for just a moment to a superficial case of etiquette, imagine a young provincial boy from an impoverished family off to college at some very elite school. Through his talent and good fortune, he comes to find himself at “high table” dining with a sophisticated class of people, all well-meaning and welcoming to him. He tries his best to display grace and gratitude, somehow managing to comport himself well in dining and conversing with these elite, even managing to be able to talk a bit about high art, like a recent opera production he was lucky enough to attend. Suppose, all going smoothly, his hosts later think well of him, and he too about himself. Here is the thing. He showed good will (not necessarily *morally* good will, but good will), and were he in various ways to have departed from these forms of etiquette, he could have been *accurately* regarded as an ass—as having shown disrespect for his hosts or other company. Nevertheless, the conventions deployed to discern his good intent, the cues he was aware of as potential signals of disapproval (even if only possible and never actually on display), all arise from a world where the persons setting those interpretive conditions are the ones empowered to do so, while he comes from a world as an outsider. Were he later to feel some pride for being a bit of a success, he might also, not unreasonably, feel degraded and in some way burdened by his disadvantaged place. It is not that there is a type of oppression of which he is a victim, although maybe there is that too. It is rather that his own exercises of responsible agency are shaped by conditions that still are liable to leave him feeling alienated.²⁰

My contention is that *this is everywhere*. Our lives are rife with circumstances like this. Academics reading this will naturally think of the familiar dynamics of what counts as appropriate behavior in conferences or colloquia settings, how to maneuver in a graduate seminar, or for that matter, when and how to show deference in responding to a referee. In philosophy settings, there are modes of argumentation, styles of asking questions, or knowing when to shut up, which are simply settled by those empowered. Even when one is at one’s best and is regarded as showing the best quality of will in his or her dealings with others, the means of showing that are by conventions shaped by an elite few.²¹

What more is there to say about this observation? Think about the social conditions informing Calhoun’s proposal. Special social circumstances—what Calhoun called “abnormal moral contexts”—set a baseline for forms of conduct regarded as acceptable in a way that shields us from accurate moral knowledge. In the cases Calhoun had in mind, certain sorts of moral wrongs, such as those involving gender disparities, were not recognized by most in the moral community, and so objectionable moral behavior was not seen as such. Here, something similar is going on: a baseline for deploying an interpretive scheme regarding what counts as signaling good and ill will is set in place by a set of background cultural practices. Participants to the practices, both those who are comparatively advantaged and those who are not, cannot be expected to have

antecedently fashioned these by reference to some morally ideal standard. They come to be taken as given background conditions of the cultural milieu in which agents get trained up into the moral community and learn to function. But as it happens, these interpretive schemes, and the conventional meanings they assign to patterns of action, have baked-in forms of bias that serve as the basis for even well-intentioned people to engage with each other. Moreover, since on the conversational theory the interpretive enterprise involves efforts to understand the particularized meaning of an agent's actions, moral agents rely upon the interpretive community being inclined to interpret them well. Those disadvantaged due to significant asymmetries in relations of power then risk alienation by defying or departing from these conventions, both in attempting to act with good will and in reacting to others when holding them accountable. Hence, our moral responsibility practices—our actual practices as they normally function—are in a sense morally tainted, or at least they are morally dubious. At the very least, they need to be assessed from a critical distance. Perhaps an example will help.

Consider a more serious matter than one simply about etiquette, one that does after all have moral import. In a comical exchange in the opening pages of E. M. Forster's *Howards End* ([1910] 1986, chapter 2), Meg discloses to Aunt Juley her sister Helen's secret—that Helen and young Paul Wilcox are in love. This is taken by Aunt Juley and Meg for apparently different reasons to be a crisis of the first order. Why? The reader is left to infer that two young people being in love is nearly tantamount to an engagement. It seems someone who has Helen's interests at heart is needed to go size up the situation by visiting Helen at the Wilcoxes' estate, where she is staying. Meg and Aunt Juley debate who is suited for this task and how to proceed, with Meg delighted for the news and Aunt Juley wary. The relative social positions of the two families are a central factor in these considerations. (We are led to believe that young Paul Wilcox comes from a wealthy family.) In explaining herself with youthful exuberance, Meg remarks that if Helen had fallen in love with a shop assistant or penniless clerk, it would not matter. Nothing other than Helen's being in love counts—although perhaps a very long engagement might be needed. To this, pleading to be the diplomat, Aunt Juley then remarks to Meg,

Now, just imagine if you say anything of that sort to the Wilcoxes. I understand it, but most good people would think you mad. Imagine how disconcerting for Helen! What is wanted is a person who will go slowly, slowly in this business, and see how things are and where they are likely to lead. (Forster [1910] 1986, 10)

Of course, as the reader soon learns (only pages later), Helen's relationship with Paul Wilcox quickly falls apart. Nevertheless, this lighthearted opening

exchange is very revealing. While Meg and Aunt Juley both have something morally important at stake—the well-being of Helen—they are alive to what would and should count as proper decorum in assessing the situation. Why is this relevant to the current topic? As the reader learns, Meg and Aunt Juley's situation is influenced by asymmetric relations of power shaped by wealth, gender, and social status.²² They are simply not suitably positioned to exercise much power over the Wilcoxes, given their place in the social life of that time. Revealingly, it never occurs to either that one might just go as relative moral and social equals and speak plainly with the Wilcoxes about the prospects for this young couple. Instead, to show good will and act well (alive to “what most good people would think”), what is needed, as Aunt Juley remarks, is the discretion of one who can signal to these people what is expected of proper folks. Hence Aunt Juley promises just to visit and discretely look about, making no mention of engagements or anything of the sort.

This might seem a curious example for me to pick from *Howards End*, given the famous scene, featured prominently in Angela Smith's (2008) work on moral responsibility, when Meg confronts her husband Henry (Mr. Wilcox, Sr.) for blindness to his own moral hypocrisy in being so unforgiving of Helen for her adultery:

“Not any more of this!” she cried. “You shall see the connection if it kills you, Henry! You have had a mistress—I forgave you. My sister has a lover—you drive her from the house. Do you see the connection? Stupid, hypocritical, cruel—oh, contemptible!—a man who insults his wife when she's alive and cants with her memory when she's dead. A man who ruins a woman for his pleasure, and casts her off to ruin other men. And gives bad financial advice, and then says he is not responsible. These, man, are you. You can't recognize them because you cannot connect.” (Forster [1910] 1986, 243–244)

A critic at this point might object that I have gotten it wrong: Meg's performance in this famous scene places on full display that, disproportionately disempowered or not, Meg has equal moral footing in the moral responsibility game for holding to account those she blames—in this case, her ass of a husband. But as I see it, this deepens my point rather than cuts against it. Why are we to regard Meg as the heroine in *Howards End*? I say it is in part because she is able to act well as a moral agent—as a morally responsible agent—holding to account others who would tarnish her sister or judge her cruelly. But her acting so well is to be regarded by us as heroic partly because she takes her role as a moral agent by operating within a social context in which she and her family are socially disadvantaged. That is, her moral agency is “against the odds,” and yet she is able

to hold others to account despite her marginalized social position. The earlier lighthearted story gives us a window into the world into which Meg, Helen, and Aunt Juley enter, and we see them socially disadvantaged in the circumstances in which they are initially to engage the Wilcoxes. Apparently, as they see their own circumstances, they are expected to operate within the conventions taken to bear on what counts as showing good will, and so on.²³

6 Conclusion

In works like *Beyond Good and Evil* ([1886] 1966) and *Genealogy of Morals* ([1887] 1967), Nietzsche famously argued that our moral responsibility practices were benighted. They concealed ugly facts about our nature and our true motivations. Although I have come at it by very different means, I too have attempted in this essay to cast a critical eye on these practices. I have no interest in drawing the sorts of conclusions Nietzsche wished to draw. But I have been at pains to scrutinize an element of Strawson's project that, to the best of my knowledge, no one has ever even considered looking upon critically—the demand for good will that according to Strawsonians serves as the foundation for our moral responsibility practices. Drawing upon my conversational theory of moral responsibility, and attending especially to the interpretive dimension to our perceptions of what counts as good and ill will, I have argued that our moral responsibility practices are tarnished or in some way benighted.²⁴ They have baked into them problematic moral assumptions placing pressure on those who are comparatively disadvantaged given significant asymmetries of power in the social relations between disparate groups.

When I first encountered this dimension of the Strawsonian enterprise, and when it occurred to me that my conversational theory took on wholesale these problematic elements of our responsibility practices, I took it to be a damning criticism of the Strawsonian enterprise and of my conversational theory. But upon reflection, I do not think it is. Indeed, I now take it to be an advantage of the conversational theory that it helps to bring these facts about our responsibility practices into clear focus. Bear in mind that the point of a Strawsonian theory is to explain our moral responsibility practices and not necessarily to endorse them. So, as a descriptive resource, I think it just helps in getting something correct. As a diagnostic resource, I also think it is useful; we are pointed in the direction of what needs correcting. But what about as a prescriptive resource? Here, I must say, the Strawsonian program and my conversational theory are completely silent.

In the opening section of this essay, I asked how both social inequities and asymmetrical relations of power affect our responsibility practices and

judgments. Drawing upon my conversational theory, I hope I have been successful in the preceding discussion in helping to answer that question. But I also asked, noting this to be the more important question, how *should* social inequities and asymmetrical relations of power affect our responsibility practices and judgments? So far as I can tell, to this normative question, Strawsonians have little to offer, nor do I from the resources of my conversational theory. Perhaps we might learn from Calhoun's proposal for how we ought to respond to those in abnormal moral contexts acting from some forms of nonculpable moral ignorance. With an eye to reform, she argues that we ought to treat these parties as if they are culpable. Maybe something similar is called for here in thinking about revising our responsibility practices in ways that more equitably reorient the standards for what counts as signaling good and ill will.²⁵

Notes

For helpful advice and comments on this chapter, I would like to thank Keith Lehrer, Kate Manne, Elinor Mason, David Shoemaker, Jason Turner, and Manuel Vargas. I would also like to thank Katrina Hutchison, Catriona Mackenzie, and Marina Oshana, for kindly inviting me to contribute to this volume, *Social Dimensions of Moral Responsibility*, and also for a set of detailed critical comments.

1. In this essay, I will focus only upon what many now call the *accountability* sense of moral responsibility wherein one who is responsible is liable to be held to account for her blameworthy behavior. Gary Watson (1996) and then David Shoemaker (2011, 2015b) have carefully identified other senses of responsibility other than the accountability sense. I make no claims about those other senses in this essay.
2. Gary Watson (2014, 17) offers the clearest expression of the thesis that sociality is at the heart of Strawson's theory. For those adopting a Strawsonian approach, see, for example, Bennett 1980; Darwall 2006; McKenna 2012; Oshana 1997, 2004; Russell 1992, 2004; Scanlon 2008; Shoemaker 2015b; Vargas 2013; Wallace 1994; and Watson 1987. Some, such as Oshana and Scanlon, might reject the Strawsonian project for other reasons (such as emphasizing the importance of the reactive attitudes), but not, I assume, the tight connection with an interpersonal constraint.
3. Strawsonian interpersonal approaches can be contrasted with intrapersonal approaches whereby the conditions for being morally responsible are accounted for just by attending to facts about the agent who is responsible and so without essential reference to considerations of holding responsible (e.g., see Glover 1970; Haji 1998; Zimmerman 1988). The most familiar version of this alternative strategy is a ledger theory. On a ledger theory, an agent's moral responsibility, including her praiseworthy and blameworthy conduct, can be understood on analogy with entries in a ledger of the moral record of an agent's conduct. The ledger will include, for instance, entries registering culpable violations of one's moral obligations, compliance with one's obligations, supererogatory conduct, and so on. What responses of holding responsible by praising or blaming are rendered apt, and from whom, are settled on different grounds. Thus, the social and relational dimensions of the practices of holding morally responsible are not essential to *being* morally responsible. The two are only contingently related. This is consistent with those contingent relations being deeply embedded in our practices and informed by important norms regarding social relations between those who are responsible and those who hold responsible.
4. For similar reasons, we can also remain neutral between those free will and moral responsibility realists who are compatibilists and those who are libertarians. The metaphysics of free

agency is not what is directly at issue here. For a survey of realist views of a compatibilist variety, see McKenna and Coates (2015), and of a libertarian variety, see Clarke and Capes (2017). Another option is to be a free will and moral responsibility realist but also an agnostic about the compatibilist/libertarian divide (e.g., see Mele 1995, 2006).

5. See Bennett (1980), Russell (1992, 2004), and more recently Shoemaker (2015b).
6. This is a widely shared thesis. For example, see Fischer and Ravizza 1998; Haji 1998; McKenna 2013; Nelkin 2011; Sartorio 2016; and Wolf 1990.
7. I am grateful to Dave Shoemaker for pointing this out. Also, as Katrina Hutchison has pointed out in her comments on this essay, there has been a considerable amount of work on this topic on connection with the topic of relational autonomy. See, for example, Baier 1985; Christman 2004; Friedman 1997; Meyers 1989; and Oshana 2006.
8. The following is a comment on this chapter written by Catriona Mackenzie, for which I am grateful: Another set of issues relevant to the epistemic condition and to problematic practices of holding responsible relates to what Miranda Fricker calls “epistemic injustice.” To suffer epistemic injustice is to be subject to prejudicial exclusion from knowledge practices on the basis of one’s social group membership. Epistemic injustice encompasses testimonial injustices, such as discounting the credibility of a person’s testimony due to identity prejudice, and hermeneutical injustice, which is “the injustice of having some significant area of one’s social experience obscured from collective understanding owing to hermeneutical marginalization” (Fricker 2007, 158). For another who has also developed the notion of epistemic injustice, see Charles Mills (2007). Thanks to Elinor Mason for the latter reference. For discussion of the relevance of Fricker’s work on epistemic injustice for moral responsibility ascriptions, see the chapters in this volume by Catriona Mackenzie and Elinor Mason.
9. Thanks again here as well to Catriona Mackenzie for the insightful observation.
10. The textual evidence is not decisive, but there is some reason to think that Strawson himself (1962) endorses this view, since when he gives his list of pleas that would defeat judgments of responsibility (and acting from objectionable quality of will), they naturally parse into control and epistemic categories. Also, Fischer and Ravizza (1998) endorse a Strawsonian account of responsibility, and seem only to identify a control and an epistemic condition on responsibility.
11. Marina Oshana helpfully advises me to clarify this point. It just comes to the following. In my view, a person might do morally wrong, might do so knowingly and freely, and yet not be blameworthy because she does not act from a morally objectionable quality of will. Suppose, for instance, she acts in the context of a moral dilemma and so cannot avoid wrongdoing. In that case, if she harbors no poor regard for anyone, and if she shows sufficient regard for all involved, then she might do wrong, but she is not blameworthy. Cases like this, I maintain, show that quality of will is a further condition on blameworthiness for wrongdoing beyond a control and an epistemic condition. Again, little turns on it in the present context.
12. This is slightly truncated. My full view also includes not just regard or concern for others but also for, as I have elsewhere put it, salient moral considerations. I include this to capture cases like harm to the environment or to nonhuman animals. David Shoemaker (2015a) has convinced me that there are problems with this formulation, however, but there is no need to work them out here for the purposes of this chapter.
13. I worry that the expression ‘morally unflawed’ is misleading, since it suggests that what is at issue is a background characterological consideration and that what blameworthiness tracks is that. But Calhoun’s point is clear enough regardless. She has in mind, I take it, the idea that those whose quality of will is morally acceptable and even laudable might still engage in certain sorts of wrongdoing by participating in oppressive practices ignorant of the wrong that they are doing.
14. Calhoun never explicitly states this last point, but as I read her, it is clearly an intended implication.
15. See note 8. Also, see the essays in this volume by Jules Holroyd, Neil Levy and Elinor Mason.
16. This is how Calhoun understood the feminist community at the time she penned her essay (1989, 397–398).
17. This is how I think of many of the issues regarding the moral status of nonhuman animals.
18. In her chapter in this volume, Catriona Mackenzie drawing on Fricker (2007), notes a finer distinction that bears on this issue. Testimonial injustice concerns prejudicial failures of

uptake on the part of hearers to have uptake due to a speaker's perceived lack of authority and credibility. Hermeneutical injustice concerns a prejudicial failure of collective interpretive resources to interpret a speaker's experiences. As will become clear, while I am interested in both, it is the latter that is more directly relevant to my thesis.

19. See Andrea Westlund's contribution to this volume, as she discusses this problem.
20. As Marina Oshana has thoughtfully noted, there are two issues at play here. One is about the liability to feel alienated. Another is about a deeper threat: the evaluations of morally responsible agency in the light of which one exercises one's own agency are often shaped outside the reach of one's own influence.
21. Here is another example, suggested to me by Elinor Mason, who permitted me to quote her from our personal correspondence. Mason writes,

[A]s we welcome our new students this week . . . I am deluged with inadvertently rude emails. It is a constantly frustrating aspect of my job that, as a woman, it is very hard to give critical feedback to students about anything other than their work. So, for example, the terrible emails they write, (which are, of course examples of them getting the etiquette rules wrong: sometimes through social disadvantage, but sometimes the opposite—some of our students are so over-privileged that they think everyone else must be the servant class). I feel like it is part of my job to tell my personal tutees how to write to their professor. But I am always met with instant hostility. So I can't do my job as well as I should be able to do, because my good will in correcting them is perceived as bitchiness or something. If a man did it, it would be perceived as avuncular good advice. So in Fricker's terminology, I am suffering a sort of epistemic injustice. But (and I take it this is part of your point), the particular way in which I am not taken seriously affects my *responsibility*, because I cannot undertake certain things trusting that my will is going to be read correctly. I cannot do my job as well as a man.

The special irony of this case is that when I received her comment, I was having the very same dealings with one of my students, and my experience is pretty much exactly as Mason predicted it would be. Indeed, in all of my years of teaching (over twenty-five now), I have *never once* had my similar corrections be treated in an unwelcome manner.

22. As it happens, we learn that the Wilcoxes apparently have considerable wealth. Moreover, we learn in a prior letter from Helen to Meg, that Mr. Wilcox (Sr.) had found occasion to say "the most horrid things about women's suffrage," but "so nicely" (Forster [1910] 1986, 506), leaving Helen to feel terribly ashamed for saying she believed in equality.
23. Thanks to Katrina Hutchison for suggesting this way of expressing my point.
24. To be clear, I take it as a problem of *any* Strawsonian approach if the demand for good will is uncritically accepted as a foundation for our moral responsibility practices. Focusing on the conversational theory as a way of bringing this to light is helpful, albeit not required to make the point. Why helpful? The theory explicitly attends to a community's interpretive resources as a strategy for discerning the "conversational" significance of an agent's actions (as a way of identifying the quality of her will). This makes it easy to reflect upon how those interpretive resources can be asymmetrically shaped.
25. This project comes down squarely within the boundaries of the approach Manuel Vargas (2013) advocates. He argues that we should evaluate our existing moral responsibility practices with an eye to improving the way agents best respond to moral reasons.

References

- Baier, Annette. 1985. *Postures of the Mind. Essays on Mind and Morals*. Minneapolis: University of Minnesota Press.
- Bennett, Jonathan. 1980. "Accountability." In *Philosophical Subjects: Essays Presented to P.F. Strawson*, edited by Zak van Straaten, 14–47. Oxford: Clarendon Press.

- Buss, Sarah. 1999. "Appearing Respectful: The Moral Significance of Manners." *Ethics* 109 (4): 795–826.
- Calhoun, Cheshire. 1989. "Responsibility and Reproach." *Ethics* 99 (2): 389–406.
- Christman, John. 2004. "Relational Autonomy, Liberal Individualism, and the Social Constitution of Selves." *Philosophical Studies* 117 (1): 143–164.
- Clarke, Randolph, and Capes, Justin. "Incompatibilist (Nondeterministic) Theories of Free Will." *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition), Edward N. Zalta (ed.), forthcoming URL = <<https://plato.stanford.edu/archives/spr2017/entries/incompatibilism-theories/>>.
- Darwall, Stephen. 2006. *The Second-Person Standpoint: Morality, Respect, Accountability*. Cambridge, MA: Harvard University Press.
- Fischer, John Martin, and Mark Ravizza. 1998. *Responsibility and Control: An Essay on Moral Responsibility*. Cambridge: Cambridge University Press.
- Forster, Edward M. (1910) 1986. *Howards End*. New York: Penguin Books, Inc.
- Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Friedman, Marilyn. 1997. "Autonomy and Social Relationships: Rethinking the Feminist Critique." In *Feminists Rethink the Self*, edited by D. T. Meyers, 40–61. Boulder, CO: Westview.
- Glover, Jonathan. 1970. *Responsibility*. New York: Humanities Press.
- Grice, Paul H. 1957. "Meaning." *Philosophical Review* 66 (3): 377–388.
- Haji, Ishtiyaque. 1998. *Moral Appraisability*. New York: Oxford University Press.
- McKenna, Michael. 2012. *Conversation and Responsibility*. New York: Oxford University Press.
- McKenna, Michael. 2013. "Reasons-Responsiveness, Agents, and Mechanisms." In *Oxford Studies in Agency and Responsibility*, Vol. 1, edited by David Shoemaker, 151–184. New York: Oxford University Press.
- McKenna, Michael, and Coates, D. Justin. "Compatibilism." *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2016/entries/compatibilism/>>.
- Mele, Alfred. 1995. *Autonomous Agents*. New York: Oxford University Press.
- Mele, Alfred. 2006. *Free Will and Luck*. New York: Oxford University Press.
- Meyers, Diana Teitjens. 1989. *Self, Society and Personal Choice*. New York: Columbia University Press.
- Mills, Charles. 2007. "White Ignorance." In *Race and Epistemologies of Ignorance*, edited by Sharon Sullivan and Nancy Tuana, 11–38. New York: State University of New York Press.
- Nelkin, Dana. 2011. *Making Sense of Freedom and Responsibility*. Oxford: Oxford University Press.
- Nietzsche, Friedrich. (1886) 1966. *Beyond Good and Evil*. Translated by Walter Kaufmann. New York: Random House.
- Nietzsche, Friedrich. (1887) 1967. *On the Genealogy of Morals*. Translated by Walter Kaufman. New York: Random House.
- Oshana, Marina. 1997. "Ascriptions of Responsibility." *American Philosophical Quarterly* 34 (1): 71–83.
- Oshana, Marina. 2004. "Moral Accountability." *Philosophical Topics* 32 (1–2): 255–274.
- Oshana, Marina. 2006. *Personal Autonomy in Society*. Aldershot: Ashgate Publishing.
- Pereboom, Derk. 2014. *Free Will, Agency, and Meaning in Life*. New York: Oxford University Press.
- Pettit, Philip. 2001. *A Theory of Freedom: From the Psychology to the Politics of Agency*. New York: Oxford University Press.
- Russell, Paul. 1992. "Strawson's Way of Naturalizing Responsibility." *Ethics* 102 (2): 287–302.
- Russell, Paul. 2004. "Responsibility and the Condition of Moral Sense." *Philosophical Topics* 32 (1–2): 287–305.
- Sartorio, Carolina. 2016. *Causation and Free Will*. New York: Oxford University Press.
- Scanlon, Thomas M. 2008. *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, MA: Belknap Harvard Press.

- Shoemaker, David. 2011. "Attributability, Answerability, and Accountability: Toward a Wider Theory of Moral Responsibility." *Ethics* 121 (3): 602–632.
- Shoemaker, David. 2015a. "McKenna's Quality of Will." *Criminal Law and Philosophy* 9 (4): 695–708.
- Shoemaker, David. 2015b. *Responsibility from the Margins*. Oxford: Oxford University Press.
- Smith, Angela. 2008. "Control, Responsibility, and Moral Assessment." *Philosophical Studies* 138: 367–392.
- Strawson, Galen. 1986. *Freedom and Belief*. Oxford: Clarendon Press.
- Strawson, Peter F. 1962. "Freedom and Resentment." *Proceedings of the British Academy* 48: 187–211.
- Vargas, Manuel. 2013. *Building Better Beings*. Oxford: 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*, edited by Ferdinand Schoeman, 256–286. Cambridge: Cambridge University Press.
- Watson, Gary. 1996. "Two Faces of Responsibility." *Philosophical Topics* 24 (2): 227–248.
- Watson, Gary. 2014. "Peter Strawson on Responsibility and Sociality." In *Oxford Studies in Agency and Responsibility*, Vol. 2, "Freedom and Resentment" at 50, edited by David Shoemaker and Neal A. Tognazzini, 15–32. Oxford: Oxford University Press.
- Wolf, Susan. 1990. *Freedom within Reason*. Oxford: Oxford University Press.
- Zimmerman, Michael. 1988. *An Essay on Moral Responsibility*. Totowa, NJ: Rowman and Littlefield.