

Corporate Moral Credit

Grant J. Rozeboom

Saint Mary's College of California USA

When do companies deserve moral credit for doing what is right? This question concerns the positive side of corporate moral responsibility, the negative side of which is the more commonly discussed issue of when companies are blameworthy for doing what is wrong. I offer a broadly functionalist account of how companies can act from morally creditworthy motives, which defuses the following Strawsonian challenge to the claim that they can: morally creditworthy motivation involves being guided by attitudes of “goodwill” for others, and these attitudes involve affect and/or phenomenal consciousness, which corporate agents cannot maintain. In response, I show that what matters about being guided by attitudes of goodwill is being directly concerned for others in one’s practical deliberation. Companies can achieve this direct concern through their decision-making procedures without affect or phenomenal consciousness. I also explore how a company’s moral creditworthiness, or lack thereof, should shape stakeholders’ relationship with it.

Key Words: corporate moral responsibility, moral worth, Strawson, moral credit, corporate motivation

In response to the Black Lives Matter movement and the most recent spate of publicized police violence against Black people in the United States, many companies enacted a variety of antiracist policies and programs. For instance, Salesforce made a public commitment to double the number of Black employees in leadership positions, to require implicit bias and equity training of all employees within one year, and to spend US\$100 million on Black-owned businesses and suppliers, among other measures (Zaveri, 2020). Taken at face value, there is something morally good about this corporate response to ongoing problems of systemic racism. Companies like Salesforce purport to be guided by a sense of justice and empathy for the victims of racism. But it is natural to take a more cynical view, which suspects that these companies are ultimately only focused on protecting themselves from financial losses due to public criticism of, and employee dissatisfaction with, their prior inaction toward racial injustice.

In raising these concerns, we are wondering, in part, whether companies’ actions are morally creditworthy.¹ Assuming they did the right thing, do they deserve credit

¹ I say “creditworthy” rather than the more common “praiseworthy” or “moral worth.” Moral praiseworthiness is broader than moral credit—some actions deserve praise that are not creditworthy, perhaps because they involved extraordinary sacrifice but were done from the wrong motives—and moral worth (at least in its original Kantian form) is narrower than moral credit, because it attaches only to actions that are performed in

for this? This question about corporate moral creditworthiness concerns the positive side of corporate moral responsibility, the negative side of which is the more commonly discussed question of whether and when companies are blameworthy for acting wrongly. In addition to knowing when companies are blameworthy—deserving of blame, anger, and ostracism—for doing what is wrong, we also want to know when they deserve credit—appreciation, solidarity, trust, and gratitude—for doing what is right. And because, as I explain later, the central component of deserving credit for doing what is right is acting from the right kind of motives, it is not enough merely to know that a company has acted morally rightly and is not blameworthy, for it might have done what is right and avoided blameworthiness and still failed to act from the right kinds of motives.²

Why does this matter? Why not just focus on whether companies do what is morally right and set aside questions about whether they deserve credit for doing what is right? It can seem out of touch and morally fetishistic to pore over the purity of a company's motives.³ But a concern with morally creditworthy corporate motivation can allow for mixed motives; we need not demand that a company not care about its profitability or attaining competitive advantage to deem its motives morally creditworthy (di Norcia & Tigner, 2000).⁴ And it is not just a moralistic concern with keeping corporate motives pure that can ground a deep concern with the moral quality of corporate motives. As Scanlon (2008) argues, developing ideas first advanced by Strawson (2008), our participation in relationships with others is regulated by our view of their motives (or “quality of will”). When others exhibit deficient motives, this “impairs” our relationship with them and calls for adjustments in our “reactive attitudes” of trust, appreciation, and gratitude toward them. Something similar seems true of our relationship to companies, as I argue later on. We depend on and are vulnerable to them, and we should regulate our attitudes toward them according to our view of their motives and whether their quality of will supports or impairs our relationship with them.⁴ For instance, how trusting should Salesforce's stakeholders be of the company in light of its antiracist efforts? To what extent should they feel solidarity with Salesforce in the fight against racism? How

the face of opposing inclinations. However, a potential infelicity in using “morally creditworthy” is that this term naturally applies to two different corporate phenomena: attributing a company, rather than some of its individual members, with a right action and (what I am focusing on here) giving a company credit for doing a right action from the right kinds of motives. In this article, I use “morally creditworthy” and “creditworthy” to refer only to the second phenomenon. Many thanks to Amy Sepinwall for encouraging me to be clearer about this.

² It also is not enough to know whether a corporation is a moral agent, because it is not clear what kind of exercise of moral agency is sufficient for acting from the right kinds of motives. Obviously, merely deciding to the right thing (as an exercise of free agency) is insufficient, and even responding to moral reasons is also insufficient, as I argue subsequently.

³ See, for example, Andrew Stark's (1993) well-known excoriation of business ethicists fussing over corporate motives.

⁴ An assumption here, then, is that a company's long-term stakeholders do not merely rely on it to act well, as I rely on my computer to function well, but also sensibly form bonds of trust and solidarity, which involve expectations of goodwill (see, e.g., Pettit, 1995) and not merely expectations of good behavior. On this aspect of corporate moral responsibility, see Tollefsen (2008: 10).

appreciative should they be of Salesforce's actions? These questions cannot be answered without assessing the moral quality of Salesforce's motives—without determining whether Salesforce deserves credit for doing what is right.⁵

My aim here is to provide a broadly functionalist account of how companies (and, by extension, similarly complex group agents) can act from morally creditworthy motives, in light of a serious challenge to the claim that they can. I assume from the outset that we must show how companies can act from the same general kinds of motives that, in individual agents, are morally creditworthy. This would be straightforward to show if, as many think, morally creditworthy motivation consists in responding to moral reasons (Arpaly & Schroeder, 2013; Markovits, 2010). But if we take seriously a Strawsonian view of the relational context of deserving moral credit and, specifically, the interpersonal expectations conveyed by the “reactive attitudes,” it seems clear that morally creditworthy motivation requires more. In addition to responding to moral reasons, agents must be guided by goodwill toward others (Rozeboom, 2017). If this is right, then it is difficult to see how companies can act from morally creditworthy motives, given how attitudes of goodwill involve affect and/or phenomenal states of consciousness.

Here is how I will proceed. In [section 1](#), I provide a brief overview of the debate about the nature of morally creditworthy motivation in individual agents, and I defend the Strawsonian view that such motivation involves goodwill toward others (which goes beyond responding to moral reasons). This will make clearer the challenge we face in accounting for morally creditworthy corporate motivation. In [section 2](#), I argue that the reason why morally creditworthy motivation involves goodwill toward others is that attitudes of goodwill entail being directly concerned for others. These attitudes ensure that those involved with our actions figure non-accidentally into our deliberations. I also explain how the practical functioning of two attitudes of goodwill—respect and care—realizes such direct concern. In [section 3](#), I then argue that, if a company can realize the practical functioning of these attitudes, then it can be sufficiently guided by attitudes of goodwill to have morally creditworthy motives. I describe how the decision-making procedures (and associated processes) of companies can realize the practical functioning of respect and care while being agnostic about the extent to which actual companies do so. In [section 4](#), I close by considering two objections that call into question the plausibility of my conception of morally creditworthy corporate motivation. I also explore further what is at stake in determining whether companies deserve credit for doing what is right. How should external stakeholders respond to companies that act rightly but do not deserve credit for doing so? Does whether we should give what Kant (2008: 14) calls mere “praise and encouragement” to companies, as opposed to more unqualified forms of trust and admiration, depend on whether their motives are creditworthy?

Note that I am working from within the broad stream of scholarly work that takes seriously the idea that corporate agents can be morally responsible for what they do

⁵This is one way, then, of thinking about the overarching moral function of giving moral credit, which I will develop throughout this article. It functions to commend and reinforce the quality of will required by the moral standards inherent in our relationships. It is not merely a device for signaling praise for good motives.

(Bratman, 2017; French, 1984; Hess, 2018; List & Pettit, 2011; Silver, 2005). It would detract from my aims here to explain why we should take this idea seriously (although my arguments will provide indirect support for it). Instead, my aim is to show those who are amenable to the thought that corporate agents can be morally blameworthy but may be reluctant to ascribe credit to them (because they lack affective and/or phenomenal states) that corporate agents can, in fact, be morally creditworthy.

1. THE CHALLENGE FOR MORALLY CREDITWORTHY CORPORATE MOTIVATION

Return to the Salesforce example. Why is it natural to greet news of Salesforce's actions with some skepticism, and how exactly does this skepticism involve withholding moral credit from Salesforce? A natural suspicion, again, is that Salesforce is merely holding its finger to the wind, doing what is needed to shield itself from the financial costs of appearing complacent in the face of current social justice movements, which its economically powerful stakeholders care about. But even if true, this suspicion does not undermine the (putative) fact that Salesforce is doing what it morally ought to do—that it acts rightly in increasing the representation of Black employees in leadership positions, in requiring implicit bias and equity training, and so on. What does our suspicion entail, then? It adds a caveat to our commendation of Salesforce for doing what is right: “Yes, it did what it should have done, but it wasn't doing it for the right reasons.” This caveat about Salesforce's motives preempts the attitudes of trust, gratitude, and appreciation that we extend to agents when we give them credit for doing what is right (more on this in [section 4](#)).

This is a familiar phenomenon, not just in dealing with corporate agents, but also, and perhaps even more so, in dealing with one another. A normally withdrawn neighbor unexpectedly drops by for some chitchat and then “just happens” to mention that they'll be moving some heavy old cabinets out of their basement next week, and oh, now that you mention it, it would be great to have an extra pair of hands. Although it is morally right to engage in friendly conversation with one's neighbors, they were not doing this from the right kinds of motives and thus do not deserve moral credit. This suggests that the central requirement for deserving moral credit concerns one's motives.⁶ One must do what is right from the right kinds of motives. And if we understand what the right kinds of motives are in individual human agents, my starting assumption, again, is that this will shed light on what these motives must be in corporate agents.

The Kantian idea of “acting from duty” initially shaped many theorists' understanding of what the right kinds of motives are, at least at the level of individual

⁶ But this is likely not the only requirement. Other theorists plausibly propose that the subjective difficulty of doing what is right (Smith, 1991) and the strength of countervailing considerations (Massoud, 2016) also impact one's level of moral creditworthiness. For companies, this might suggest that the convenience of a right action diminishes its moral creditworthiness, for example, a restaurant donates leftovers to a food pantry when it would be just as costly to dispose of them.

human agents (e.g., Herman, 1993), but there is now some consensus that the motive of duty is neither necessary nor sufficient for morally creditworthy motivation (Arpaly & Schroeder, 2013; Markovits, 2010; but see Johnson King, 2020; Sliwa, 2016). Agents can act from duty without deserving moral credit, because they might have a deformed sense of duty, and they can deserve moral credit without acting from duty, because they might be sufficiently guided by a sense of what morally matters without cognizing their action as being morally required. In place of the motive of duty, moral theorists have converged on the general idea that creditworthy motivation consists in responding to moral reasons—the reasons why one’s action is right (Arpaly & Schroeder, 2013; Markovits, 2010).

If this was all that morally creditworthy motivation required, then it would not be difficult to explain how companies can have morally creditworthy motives. The existing literature on corporate agency shows how companies can respond to moral reasons, such as reasons of economic fairness, at least as some of the reasons behind their actions (see, e.g., Hess & Björnsson, 2017; List & Pettit, 2011; Tollefsen, 2015). For all we need is a model of how companies can act on the basis of different kinds of reasons—a model of rational, intentional corporate agency. And there are many such models, which, although they are subject to some dispute, plausibly underwrite the idea that companies can act on the basis of moral reasons (although perhaps not by regarding them as moral reasons, but this is not required for moral creditworthiness on the generally accepted moral reasons–based view). (In addition to the formerly cited literature, see French, 1984, 1996.)

But once we examine more closely the relational context within which moral credit is earned and given, it is clear that responding to moral reasons is not enough for full moral creditworthiness. It helps to start here with Strawson’s (2008) influential account of moral responsibility in “Freedom and Resentment.” One central idea in this essay is that being morally responsible for some action—either blameworthy for wrongdoing or creditworthy for right-doing—consists in being an apt target of “reactive attitudes,” such as resentment and gratitude. Determining the conditions for being blameworthy or creditworthy is a matter of determining what it takes to violate or satisfy the internal standards of the reactive attitudes—the standards that, when thought to be violated or satisfied, trigger the various reactive attitudes. How should we understand these internal standards? To avoid getting sidetracked in Strawson exegesis,⁷ let me just briefly and baldly state how I understand them. I think they are fundamentally expectations for goodwill from others: expectations that others’ actions display goodwill and, conversely, avoid displaying ill will. Strawson claims that these are endemic to our responsibility practices (Shoemaker, 2013). If we could provide an informative account of what

⁷For instance, some theorists take the internal standards to be essentially communicative, such that the reactive attitudes all address demands that call for a response, for example, resentment signals that a demand was not met and calls for apology and recompense (Darwall, 2006; McKenna, 2012; Watson, 1987). But these communicative accounts face serious problems, stemming from the fact that the contexts in which the incipient communication of reactive attitudes is apt are narrower than the contexts in which the reactive attitudes are warranted (Vargas, 2016).

the expected forms of goodwill are, then we could develop a view of morally creditworthy motivation that respects the Strawsonian insight that earning and giving moral credit occur in a relational context structured by the internal standards of reactive attitudes.

Along these lines, in earlier work, I defended the view that morally creditworthy motivation consists in being guided by the appropriate attitudes of goodwill, such as respect and care, for those involved with one's actions (Rozeboom, 2017). This view makes clear how being guided by such attitudes goes beyond responding to moral reasons. Consider my central example: Charity values helping others, but she doesn't particularly care about the individuals she helps. She just really likes helping others in the same way that you might like baking a crusty loaf of bread, or playing soccer, or tending your garden. Carey, by contrast, does not intrinsically enjoy the activity of helping others, but she does care about those she encounters and, as a result, is intrinsically motivated to help them. When we imagine a situation where Carey and Charity are each drawn to help someone in need, and they do so for the same moral reason—the importance of alleviating that person's need—it seems clear that Carey will be much more deserving of moral credit than Charity. My account was meant to explain this: Carey is guided by an attitude of goodwill for the person being helped, whereas Charity is not (even though, again, she responds to the same moral reason).⁸ Returning to the Strawsonian insight, in being guided by an attitude of care for others, Carey maintains and displays goodwill for the recipients of her aid, whereas Charity, responding only to the fact that someone was in need of her aid, does not, or at least to a far lesser degree.⁹

One way of explaining my verdict here is to draw on the familiar idea that an agent deserves credit for doing what is right only if, given their motives, it is no accident that they do what is right. It is a notoriously tricky matter explaining what the relevant sense of “no accident” is (Johnson King, 2020). My suggestion is that, given the Strawsonian relational context of moral credit, the relevant idea of nonaccidentality is itself relational, that is, it concerns one's orientation toward those whose interests and rights pertain to what one is doing. It is *relationally* no accident that one does what is right if one is properly attuned to and concerned with those who are importantly involved with (or impacted by) what one does. This notion of relational

⁸ Recent theorists might instead respond to this case by returning to a modified version of the motive of duty view (Johnson King, 2020; Stratton-Lake, 2000) or a modified version of the moral reasons view (Singh, 2020), arguing that Charity's action is less creditworthy because her motives are more contingently connected to doing the right thing than Carey's, and we can only strengthen the modal connection between her motives and doing what is right through some form of the motive of duty and/or responsiveness to moral reasons. But as I explain later, the relevant notion of nonaccidentality is relational, and it entails that we can strengthen the modal connection between motives and right action only through the practical functioning of attitudes of goodwill toward others.

⁹ You might think that, to defend the idea that attitudes of goodwill are central for morally creditworthy motivation, I need a stronger claim here—the claim that agents like Charity are not at all morally creditworthy for their right actions. I disagree, because, first, the fact that there is a difference in the degree to which agents like Charity and Carey are creditworthy can plausibly be explained only in terms of their (lack of) attitudes of goodwill and, second, the limited creditworthiness of agents like Charity can be explained in terms of how their motives *partially* realize attitudes of goodwill. More on this in section 3.

nonaccidentality is one way to further clarify the goodwill-focused internal standards of the reactive attitudes. In evaluating whether others show the relevant forms of goodwill, a part of what we are doing is evaluating whether they are nonaccidentally concerned with those involved with their actions.

This idea can best be appreciated from the perspective of those who are the beneficiaries of others' morally right actions. When others treat us rightly—pay heed to our preferences, uphold their commitments to us, respond to our calls for help—what determines the degree to which we take their actions as a basis for moral credit—for increasing our levels of trust, appreciation, and gratitude toward them—is whether we think their action reflects motives that are nonaccidentally concerned with us. If someone responds to our plea for help, but only because we happen to be the nearest option for indulging their enjoyment of helping others (e.g., Charity), our appreciative response toward them will be dampened. We do not matter to them, beyond being an occasion for their beneficence. Their motives do not manifest nonaccidental concern for us.¹⁰

With this relational notion of nonaccidentality in mind, return to the contrast between Charity and Carey. Charity accidentally does what is right, because it just happens that what she likes to do—helping others—aligns with being attuned to how her action importantly impacts those involved with it. In other cases, when what morality requires of her goes beyond or conflicts with benefiting others (as when she should avoid paternalistically helping others), Charity's motives will lead her astray. Not so for Carey: she is guided by an appropriate goodwill-attitude toward others by which she is attuned to how they are importantly involved with what she does, and so she nonaccidentally does what is right (on the relational notion of nonaccidentality). Even if you do not find the Charity–Carey case intuitively compelling, it illustrates this important contrast.

At this point, you may doubt that the lesson of the Charity–Carey case extends to corporate agency. This case is relatively personal, involving a small number of agents who are familiar with one another. Corporate agents, by contrast, tend to operate in impersonal environments that involve an indeterminate number of agents who are strangers. In these environments, it is natural to think that doing what is right out of a general sense of duty—for example, paying one's taxes because it's the right thing to do—is enough for morally creditworthy motivation, even without any goodwill-attitudes. But consider why agents in these impersonal environments care about doing what is right. Is it out of respect for their fellow citizens and the political institutions that call on them to do their fair share (the relational context relevant to

¹⁰ This line of argument shares some affinity with Bernard Williams's (1981) famous "one thought too many" thought experiment, in which he suggests that a husband who, when faced with the choice of saving a drowning stranger or his drowning wife, first had to apply some general moral principle to determine that saving his wife was morally permissible or required would require one thought too many. One (but by no means the only) potential upshot of Williams's thought experiment is that morally creditworthy motivation centrally involves being directly concerned for others—from the perspective of the drowning wife, what matters is that the husband acts out of love for her. This is what Williams describes as "the thought that it was his wife" who was drowning (18). But I do not subscribe to, nor does a Strawsonian need to accept, Williams's conclusion that systematic moral theory is a "misrepresentation" (19).

paying one's taxes)? Or do they just have a bare impulse to do what is right? The former is more creditworthy, and we can see this using, again, the relational notion of nonaccidentality. Paying one's taxes simply from an impulse to do what is right, without any respect for one's fellow citizens and shared political institutions, makes it an accident that one does what is right, given how others (though strangers) pertain to one's actions. One's motive of duty is relationally untethered, disconnected from the attitudes of goodwill that make one attuned to how others are importantly impacted by one's actions. We can see clear analogues of this in corporate agency (the structure of which I discuss in [section 3](#)) when we compare companies that are merely concerned to avoid doing wrong (e.g., Google's former "don't be evil" slogan) and companies that cultivate practices and policies that situate a commitment to avoid doing wrong within a broader concern for those impacted by the companies' actions.

If the preceding line of thought is on the right track—that morally creditworthy motivation involves being guided by goodwill-attitudes for those importantly involved with one's actions—then the task of showing how companies can have morally creditworthy motives becomes much more difficult. It is unclear that companies can be guided by goodwill-attitudes. This is because such attitudes involve distinctively affective and phenomenal states. (On Darwall's [2006: chap. 6] "second-personal" account of respect, for instance, respect is tied to emotions like guilt and anger.) Even if we grant that companies can respond to moral reasons, it is doubtful they have such affective and phenomenal states (Baddorf, 2017; Sepinwall, 2017; Tollefsen, 2008).

2. FUNCTIONING FOR DIRECT CONCERN

To begin addressing this difficulty, we should first ask what attitudes of goodwill contribute to making an agent's motives morally creditworthy. How do goodwill-attitudes make it the case that (to continue using the relational notion of nonaccidentality) an agent nonaccidentally does what is right, given how others are importantly involved with their action? My conjecture is that goodwill-attitudes contribute direct concern for those involved with an agent's action. This is what Carey has and Charity lacks. Charity has no direct concern for those she aids. They slip behind her enthusiasm to alleviate their need, rather than being the point—the "end" (*Zweck*) in a Kantian sense—of her endeavoring to help them (Velleman, 1999).

To explain, suppose that we add to Charity's motivational profile the fact that she is directly concerned for those she benefits but hold everything else fixed, and so we leave open the question of whether she maintains a goodwill-attitude, like Carey. This would mean that she is oriented toward those she helps not merely as beneficiaries of an activity she intrinsically enjoys but also as individuals who are worthy of her consideration independently of how they figure into activities she enjoys. Even without knowing anything else about her motivational profile, it seems clear that her motives are now (more) nonaccidentally attuned to those in her relational context and thus (more) morally creditworthy. She now has the component of Carey's motives that counts for deserving moral credit. This suggests that what

matters about being guided by goodwill-attitudes is that one's motives involve a direct concern for those impacted by one's actions.¹¹

How exactly do attitudes of goodwill, such as respect and care, involve being directly concerned for others? I will argue they do so via their practical functioning, where "practical functioning" concerns the roles these attitudes play in practical reasoning. These are the roles that goodwill-attitudes play in shaping what reasons agents consider or exclude and what weight they place on these reasons. Realizing such practical functioning is sufficient for being directly concerned for others. Being directly concerned for others thus does not require the affective and phenomenal states that goodwill-attitudes involve.

Our next question will be whether corporate agents can realize this practical functioning. If yes, then we can set aside the issue of whether corporate agents have the affective and phenomenal states of goodwill-attitudes, for we will have shown that corporate agents can satisfy what the internal standards of the reactive attitudes most fundamentally require in calling for goodwill—that agents' motives be non-accidentally attuned to others.

Begin with the attitude of respect for persons, which we can understand broadly as the attitude of accepting someone's basic, authoritative standing as a person. This centrally, but not exclusively, involves being disposed to treat others' decisions that are relevant to one's interaction with them as preemptive in one's deliberations about how to treat them.¹² That is, given that the respected individual *R* has made a relevant decision *D*, one is disposed to i) treat the fact that *R* has made *D* as a preemptive reason not to interfere with *D* and, perhaps, to support *D*, and thus to ii) set aside a range of other reasons in favor of interfering with *D*. This attitude entails, furthermore, iii) keeping track of how *R*'s exercise of their authoritative standing as a person pertains to what one is considering doing—a kind of vigilance—and iv) giving *R* the opportunity to weigh in prior to one's decisions about personal matters that fall within the scope of *R*'s personal authority.

Realizing these four practical functions with respect to *R* is sufficient for being directly concerned for *R*. (Again, I do not deny that there is more to the practical functioning of respect for persons than the four-part functioning laid out here; it also seems to include, for example, being accountable to others [Darwall, 2006: chap. 4].) It ensures that one is reliably disposed to pay attention to *R* and take *R* into account in one's reasoning whenever one bumps up against *R*'s personal affairs, and not simply when *R* happens to pertain to one's other concerns. That is, we cannot explain *R*'s

¹¹ Now, you might press further here: why should moral creditworthiness hinge on being directly, and thus nonaccidentally, concerned for others in this way? To begin answering this question, which pushes beyond the scope of this article, recall our Strawsonian insight from section 1, that we can understand moral creditworthiness in terms of our broader practices of reactive attitudes whose internal standards call for showing goodwill. These practices seem to be justified, at least in part, by how they institute mutual recognition between persons (see Darwall, 2006; Helm, 2017; Jonker, 2020). Given this broader aim of instituting mutual recognition, it makes sense that direct, nonaccidental concern for others is the central constituent of morally creditworthy motivation.

¹² On the relevant idea of respect for persons, see Darwall (2006) and Rozeboom (2022), and on preemptive reasons, see Raz (1986: chap. 3; 2006: 1021–22) and Darwall (2013: 152–55).

counterfactually durable appearance in one's thought by pointing to one's other concerns, and we can explain it simply by pointing to the four practical functions of one's attitude of respect. It thus makes sense to think that this practical functioning is sufficient for one's being directly concerned for *R*. This means that, even though the attitude of respect for persons is also distinguished by various affective dispositions that also may be sufficient for being directly concerned with *R* (such as feeling guilt for acting disrespectfully toward *R*), they are not necessary for such direct concern.¹³

You might worry that there are objectionable attitudes (for example, paternalizing or condescending attitudes) that seem to involve being directly concerned for others in the same general way. For a corporate example, consider the paternalizing attitude that WeWork conveyed toward its employees in imposing strict vegetarian policies on them (Gelles, 2018). WeWork wanted to help its employees do their part to mitigate the climate crisis and realize better personal health outcomes. Such measures can be understood to express an attitude of direct concern, but one that is objectionable and not (very) creditworthy. I agree and, in response, need to make explicit a background assumption of my argument: the relevant forms of practical functioning derive from attitudes that are morally appropriate to their objects. There is ongoing debate about what makes attitudes appropriate to their objects (Rozeboom, 2017: 9–10), but the outcome of this debate should not alter my argument—that the practical functioning of respect for persons is sufficient for the kind of direct concern that can constitute morally creditworthy motivation—because respect for persons is a paradigmatically appropriate attitude to hold toward persons. We can thus move forward with this background condition now made explicit, that the relevant attitudes of goodwill are morally appropriate to their objects.

Let us turn now to care for persons, broadly understood as an attitude of benevolent, sympathetic concern for others.¹⁴ (This may be the same basic attitude of care we extend to nonpersons, but I will set that issue aside here.) Its practical functioning involves i) understanding what it is like for the cared-for person *C* to be in the conditions resulting from one's actions, ii) keeping track of any impediments to *C*'s well-being that these conditions present—another kind of vigilance—and, as a result, iii) desiring to remove these impediments. That is, one wants to help and not harm *C* simply because one understands what it is for *C* to face the conditions resulting from one's actions, and regardless of how doing so pertains to one's other interests and aims. This is a benevolent attitude, because it involves a noninstrumental desire to protect someone's well-being, and it is sympathetic, because it is based on understanding what it is like to be them.

Here, again, we see that the attitude's practical functioning is sufficient for being directly concerned with its object. The attitude's *C*-directed dispositions ensure that the subject of care is reliably disposed to pay attention to *C* and take *C* into account

¹³ For a similar argument about the auxiliary role of moral emotions in Kantian respect-attitudes, see Hess (2018: 78).

¹⁴ Here I am influenced by Darwall (2002: chap. 3) and the care ethics tradition, for example, Noddings (1984) and Benhabib (1985).

whenever their actions intersect with *C*'s welfare, regardless of whether *C* pertains to their other concerns. And, as was true for respect for persons, although care involves affective dispositions—feeling anxious about impediments to *C*'s well-being or feeling relieved when *C* is kept from harm, for example—these are not necessary for the functioning of care to realize a form of direct concern for *C*. The background condition on goodwill-attitudes is also satisfied by care: it is generally an appropriate attitude to hold toward persons. We have, then, another example of how an attitude's practical functioning suffices for the kind of direct concern that constitutes morally creditworthy motivation.

One concern with treating affect as unnecessary for realizing the creditworthy-relevant components of goodwill-attitudes is that my account overlooks how moral credit seems to be especially or more fully deserved by agents who realize the affective states associated with goodwill-attitudes.¹⁵ Don't we feel special appreciation, trust, and gratitude for those whose feelings, and not just decision-making, constitute direct concern for us? Yes, but I think this is only because we often view the presence of such affect as an important signal of the practical functioning of goodwill-attitudes. That is, for many kinds of agents (but not all human agents, and not corporate agents, as I discuss later), we take their affect to be a necessary concomitant of fully realizing the practical functioning of direct concern of goodwill-attitudes. The presence or absence of affect in these agents thus sensibly shapes our credit-giving responses toward them. But for agents that are incapable of experiencing some or all forms of affect but still capable of realizing the practical functioning of direct concern of goodwill-attitudes (e.g., corporate agents), the presence or absence of affective states does not sensibly shape our credit-giving reactions in this way.

3. REALIZING DIRECT CORPORATE CONCERN FOR OTHERS

My next claim is conditional: if companies can realize and guide their actions by the practical functioning of attitudes of goodwill, then companies can realize the aspects of these attitudes sufficiently for morally creditworthy motivation. My argument for this claim is a natural extension of what I argued earlier: if the reason why morally creditworthy motivation involves being guided by goodwill-attitudes is that these attitudes involve being directly concerned for others, and if the practical functioning of these attitudes, as described, is sufficient for being directly concerned for others, then realizing the practical functioning of these attitudes is sufficient for realizing the morally creditworthy aspects of goodwill-attitudes. This can be achieved even if the agent in question does not fully maintain those attitudes because they lack affective or phenomenal states. And because realizing the creditworthy aspects of goodwill-attitudes is what constitutes morally creditworthy motivation, then realizing the practical functioning of goodwill-attitudes, as described before, can suffice for

¹⁵ You might also worry that my account loses what is distinctively "second personal" about goodwill-attitudes (Darwall, 2006). But see Morrison et al. (2022) for a functionalist-friendly account of second-personal attitudes.

having morally creditworthy motives.¹⁶ It follows that if companies can realize this practical functioning, then they can have morally creditworthy motives.

If that is right, then it does not matter whether companies can maintain the affective and phenomenal states of goodwill-attitudes. Even if we are reluctant to say that companies can fully maintain such attitudes in the absence of these states, we may still allow for companies to realize the practical functioning of goodwill-attitudes sufficient for being directly concerned with others. In this way, my argument establishes a principled sufficiency condition for morally creditworthy corporate motivation: it shows what it takes for a corporate agent's motives to realize enough of the features of morally creditworthy motivation found in individual human agents—namely, those that are sufficient for the direct concern of goodwill-attitudes.

In doing so, I avoid a problem for Hess and Björnsson's (2017) somewhat parallel account of corporate blameworthiness. (Again, I am focusing on the positive correlate of blameworthiness, creditworthiness.) They provide a functionalist description of the reactive attitudes they think any blameworthy agent must be capable of experiencing, such as indignation and guilt, and they argue that the internal processes, norms, and structures of corporations can realize what they call "morally equivalent" functions of these attitudes. This is similar to how, later, I argue that companies can realize the practical functioning of goodwill-attitudes that is sufficient for having creditworthy motives. But for Hess and Björnsson, "moral equivalence" requires only that corporate agents realize the features of reactive attitudes that are sufficient for exercising rational agency (278), and they do not explain why realizing only these rational agential features is sufficient for moral blameworthiness. There is thus a gap in their account between the idea that moral blameworthiness requires the capacity for experiencing reactive attitudes and the claim that realizing the rational agential features of reactive attitudes is sufficient for moral blameworthiness. No such gap exists in my account, because the practical functioning that I claim is sufficient for moral creditworthiness derives from how our relational practice of giving moral credit responds to goodwill-attitudes—specifically, to how these attitudes involve being nonaccidentally concerned for others.

How, then, can companies realize the practical functioning of goodwill-attitudes sufficient for being directly concerned for those involved with their decisions? That is, how can the decision-making of companies be shaped in the ways that the reasoning of individual agents is shaped when they are guided by goodwill-attitudes, such as respect and care? To answer this question, I will focus on companies' decision-making procedures, without ruling out other aspects of corporate life that may also suffice for realizing direct concern (such as corporate culture, which I discuss later¹⁷). Although this is in keeping with the tradition started by Peter French

¹⁶ I say "can suffice" because, as I mention later (note 31), it matters to what extent this practical functioning extends to all of the individuals importantly involved with a corporate agent's actions.

¹⁷ For a view of corporate moral responsibility that centers on corporate culture, see Silver (2005: 287ff.). For the claim that corporate culture can help realize the kind of practical functioning I am discussing here, see Hess (2018).

(1984, 1996), who takes “corporate internal decision-making (CID) structures” as central to explaining corporate moral responsibility, I will go further. An important question that French does not fully answer is how the social-psychological organization of a company instantiates CID structures. That is, what makes it the case that a given decision-making procedure, described abstractly as a set of instructions about how to make and/or implement a certain kind of decision or policy, actually characterizes the company’s operations such that it counts as one of the company’s CID structures?

French’s (1984: 60–62) initial response to this question is suggestive but incomplete: “acceptance among the corporate personnel or the higher managerial officers.” What does this acceptance involve, and in particular, what kind of buy-in from an organization’s different members does it require? Michael Bratman (2017) argues that it can be modeled as shared policies, which involve an interlocking web of mutually supporting individual intentions (an extension of his view of small-scale shared agency). Each individual employee intends to do their part partly by way of the intended contributions of others, so as to carry out the overarching aim or policy.

But I do not think that this web of interlocking intentions is necessary (and, to be clear, Bratman thinks it is only sufficient¹⁸), given, as Scott Shapiro (2011) points out, the amount of motivational “alienation” often present in large organizations. That is, oftentimes employees comply with a corporate procedure simply to do their job and get paid and not also to support the larger functioning and aims of the organization (Shapiro, 2011: 144–49). In place of interlocking intentions, there can instead be hierarchical structures and divided labor, by which persons “intentionally play their parts, and resolve their disputes peacefully and openly” (Shapiro, 2011: 149). This ensures that “alienated participants end up acting in the same way as nonalienated ones” (Shapiro, 2011: 150).

I will focus on this kind of heterogeneous structure of procedure acceptance going forward, given the amount of motivational alienation that can arise within a company. Such a heterogeneous structure will involve the reliable implementation of rules sorting out people’s roles and authorizing some of them to make decisions on behalf of the corporate agent that others must follow, what Shapiro (2011: 169ff.) calls “authorizing” rules. And the reliable implementation of such rules can, again, result from individuals operating with varied motives. Some members, especially those in leading roles, may be individually committed to the overarching goals and principles that purport to justify the rules, but others (those who are “alienated”) may simply be committed to doing their jobs, staying out of trouble, and getting paid.¹⁹

¹⁸ In more recent work, Bratman (2022) develops a detailed theory of large-scale, institutional agency that accommodates the sort of alienation I go on to discuss. I take Bratman’s view to be largely compatible with my account. My view goes beyond Bratman’s, though, in providing an account of the functioning of goodwill-attitudes, which are pertinent to moral creditworthiness. Bratman focuses only on the corporate functional realization of intentions. These corporate intentions may play a role in constituting corporate goodwill, but we can understand their role in doing so only by applying a prior conception of the functioning of the goodwill-attitudes that matter for moral credit, such as the one I defended earlier.

¹⁹ Shapiro’s point here was anticipated by French (1996: 152) and, again, developed in detail by Bratman (2022).

Notice further that reliably upholding and enforcing a company's rules does not require perfect adherence. But it does require that the rules not be openly violated with impunity. At least some of those charged with enforcing procedures must respond with rule-prescribed sanctions for known violations.

We now have a general picture of how formal decision-making procedures can be realized as the procedures of a company. This picture already suggests how a company itself, and not just its individual members, can realize various forms of practical functioning: it allows for a company's decision-making procedures to be instantiated even if its members commit only to their part of implementation and not also to the larger procedures themselves and the corporate aims they encompass. This will help show how a company itself, and not just its individual members, can realize the practical functioning that constitutes morally creditworthy motivation.

Consider first the practical functioning of respect for persons, which I described as follows:

Respect function. Given that a respected individual R has made a relevant decision D about their personal affairs, one is disposed to i) treat the fact that R has made D as a preemptive reason not to interfere with D and, perhaps, to support D , and thereby ii) set aside other reasons in favor of interfering with D , and also to iii) keep track of how R 's further exercise of their authoritative standing pertains to what one is considering doing and, if feasible, iv) give R the opportunity to weigh in prior to one's decisions that may impinge upon the personal matters that fall within the scope of R 's authoritative personal standing.

These four practical functions can be realized by the decision-making procedures of a company. That is, we can formulate corporate procedures that instantiate these four aspects of respect's practical functioning, which can be realized within companies according to the preceding picture of corporate procedure realization.

Let us illustrate with part i) of the respect function, taking R to refer to a hypothetical company's contractors. (You might imagine Uber becoming genuinely concerned with the personal autonomy of its drivers.) Suppose that there are roles and teams within the organization whose work includes shaping policies that might impinge on contractors' personal lives, for example, policies about personal social media posting. The company may implement a rule that governs the decision-making of these roles and teams of the following form:

Procedure 1. When creating or revising policies that might impact contractors' personal lives, leave room for their personal decision-making, unless there's no other way to pursue very important company aims; proactively seek contractors' input about these policies; and do not ignore this input, unless there are very strong, morally comparable reasons for doing so.

If reliably implemented, such a procedure would realize part i) of the respect function toward a company's contractors.²⁰ Cognates of procedure 1 could be implemented for the company's other stakeholders.

²⁰ A worry here is that implementing a procedure for treating others' decisions as authoritative—as having preemptive weight and excluding countervailing considerations—is not the same thing as genuinely viewing

We could imagine similar procedures (or clusters of procedures) that would realize the other aspects of the respect function. There is a complication with part iii), because it calls for a general form of vigilance, rather than imposing a direct constraint on corporate decision-making. But this complication is easily addressed by considering how companies commonly exercise various forms of vigilance, for example, they monitor changes in consumer demand for their products. Similar information-gathering and -disseminating procedures could, at least in principle, be used for keeping track of how the individuals whose personal lives may be impacted by a company's actions tend to make decisions about the pertinent personal matters. To extend the foregoing example, just as companies regularly gather personal information from employees about their preferences and routines through surveys, focus groups, and townhall meetings, they could do something similar for their contractors. And if such information gathering were done to ensure that, in line with procedure 1, the company's contractor-impacting policies were crafted to better respect contractors' personal decision-making, then it would realize the vigilance of the respect function.

This is only a partial sketch of how a company's decision-making procedures and associated processes could realize the respect function. But because I aim only to show how it is possible for companies to realize the practical functioning of attitudes of goodwill, and not to determine the extent to which actual companies do so, I will move on to consider our next important question: can a company realize the respect function nonderivatively, that is, without its members realizing the respect function in their individual minds? Is it the corporate agent itself, or just a collection of the individuals in it, that realizes this core functioning of respect for persons?

A company's realization of the respect function does not require any of its members to individually realize any part of the respect function. This is because, as I pointed out earlier, the heterogenous set of individual attitudes that support the implementation of a company's procedures, such as procedure 1, need not involve individual attitudes that realize the respect function. For instance, perhaps following procedure 1 is a part of my job, and I am thus committed to applying this procedure in all meetings with colleagues about company policies that impact contractors, to holding my colleagues accountable for following procedure 1, and to abiding by decisions made on the basis of procedure 1. But perhaps I do so only because I am committed to receiving performance bonuses, and fastidiously implementing company policies is a key criterion in my performance evaluations. This makes me, to continue using Shapiro's (2011) terminology, an alienated participant. My intention to receive performance bonuses does not realize (any part of) the respect function in my individual psychology. But even so, it figures into a network of individual attitudes that, together, does realize this practical functioning. This wider network may be similarly composed of role-based, pecuniary

these decisions as authoritative. It merely mimics doing so. This worry is an instance of a more general concern about functionalist approaches in the philosophy of mind and action that I cannot tackle here. I am simply deploying a functionalist approach that has gained traction in standard philosophical accounts of practical reasoning (see esp. Bratman, 1987; Setiya, 2007; Velleman, 1989) to theorize about corporate moral agency. This, again, is also broadly in line with the approach taken by Hess and Björnsson (2017) and Hess (2018).

attitudes that, because of the enforcement and accountability mechanisms in place, reliably implement procedure 1.

You might respond that policies like procedure 1 must be created and enforced, at least initially, by decision makers whose individual attitudes do overlap with the respect function. Think, for instance, of the much-lauded employee-friendly policies at Barry-Wehmiller Group, which seem to be driven by the sincere compassion that Barry-Wehmiller's CEO, Bob Chapman, has for his employees (Chapman & Sisdia, 2015). Although this is a common feature of how the relevant sorts of policies are created and enforced, it is not necessary. Suppose that a company's CEO and other high-level executives hold a series of meetings to decide whether the company should adopt a more hands-on approach to promoting wellness among stakeholders, including contractors. They decide it should not, and this decision results in, among other things, the adoption of procedure 1. In so deciding, none of the executives' reasons for adopting, and now being committed to enforce, procedure 1 must stem from an attitude of respect that realizes the respect function in their individual minds.²¹ Perhaps the CEO likes procedure 1 because he thinks it will help promote more innovation among contractors, while the CFO likes procedure 1 because it will increase contractor job satisfaction, thereby decreasing the expenses caused by high turnover, and so on for the other executives. None of them need be guided by an attitude that, in their individual minds, realizes (any part of) the respect function. But once they choose procedure 1 and it takes hold through various mechanisms of reporting and accountability, including employees who are motivated by performance incentives and come to treat following procedure 1 as a part of their jobs (on this point, see Hess & Björnsson, 2017; List & Pettit, 2011), we can see that procedure 1 can be both adopted and reliably enforced in organizations whose members do not individually realize (any part of) the respect function.²²

Let us next consider the attitude of care for persons, the practical functioning of which I described as follows:

Care function. i) Understanding what it is like for the cared-for *C* to be in the conditions resulting from one's actions, ii) keeping track of any impediments to *C*'s well-being that these conditions present and, iii) as a result, desiring to remove these impediments.

²¹ You might object here that, even if the executives' reasons for choosing procedure 1 do not involve attitudes that realize (any part of) the respect function in their minds, their resulting commitment to do their part to implement and enforce procedure 1 will need to do so. But this is not necessary, for they may be motivated to do their (albeit leading) part in implementing and enforcing procedure 1 by the same kinds of role-based and pecuniary considerations that drive lower-level employees to do their (albeit subordinate) parts. I am grateful to Michael Bratman for pressing this issue.

²² You might still object that the corporate functioning is at the behest of the individuals who design and implement the policies, who are the only agents eligible to be morally creditworthy. This raises large issues that are beyond the scope of this article (as I cautioned at the end of the introduction), but notice that, given what I have established so far, corporate agents can achieve independence from the individual designers and implementers of corporate policy in two crucial ways: 1) as just discussed, the individual agents need not have the creditworthy motives realized by the corporate agent, and 2) the individual agents who, on this objection, are exclusively eligible to be creditworthy could be replaced over the course of the creditworthy corporate activity without any impact on the corporate agent's functioning.

As with the respect function, the care function contains a form of vigilance in part ii) that could be realized through a company's information-gathering and -disseminating procedures.

Suppose, for instance, that a dairy-processing company is developing a new cheese product and decides to monitor, in addition to the legally requisite safety and quality measures, how the new product more broadly impacts consumers' well-being. Is the standard amount of consumer consumption of the product compatible with a healthy diet? Do its packaging and formulation encourage overconsumption? Does the supply chain avoid unnecessary harms to consumers through pollution (e.g., waste runoff from industrial farms)? If these investigatory mechanisms feed into decision-making procedures that realize part iii)—which I discuss later—then they will help realize the care function.

You might doubt that such monitoring is sufficient for realizing part ii) of the care function, because part ii) connotes a kind of sympathetic understanding that cannot be achieved merely by way of information gathering and dissemination. But remember that part ii) simply designates an aspect of practical functioning—a part of the role that the attitude of care plays in practical reasoning. The function that sympathetic understanding of someone's condition plays in care-shaped practical deliberation is supplying information about the person's condition that pertains to how one's actions impact the person's well-being. The information-gathering and -disseminating process described earlier suffices for this.

To realize part iii) of the care function, we need to understand the practical functioning of (noninstrumental²³) desire. The desire of part iii) aims at removing impediments to *C*'s well-being (i.e., to keep *C* from harm) simply for *C*'s sake and not for some further end beyond *C* and *C*'s welfare. What are the key functional characteristics of such a desire? Here, again, we run into the issue of explaining how we can attribute an attitude to corporate agents that, in individual human agents, is characterized by affective and phenomenal states. For instance, the known satisfaction of one's desires tends to generate a pleasant feeling. As before, I want to bracket this issue, because appealing to the practical functioning of desire is sufficient for explaining how part iii), in conjunction with parts i) and ii), can constitute a form of a direct concern for others.

This functioning concerns how desire shapes action and attention—what Neil Sinhababu (2017) calls desire's "motivational" and "attentional" aspects. Desire shapes action by being what many philosophers view as the most basic, generic form of motivation. It involves being disposed to choose and pursue whatever one desires, provided that one has beliefs about how to do so (and relative to the strength of one's desires).²⁴ And desire shapes attention by disposing one to attend to the features of

²³ Part iii) derives from the fact that care involves a noninstrumental desire for the well-being of cared-for persons.

²⁴ This way of describing how desire shapes action is heavily influenced by the Humean theories of desire put forward by, for example, Michael Smith (1994) and Sinhababu (2017). But it can be construed in ways amenable to Kantians; see, for example, Tamar Schapiro's (2011) account of desire as involving "incorporation," that is, the tentative volitional endorsement of the object of one's desire. On the Kantian

one's situation relevant to satisfying one's desires. That is, one is disposed to use one's cognitive resources to keep information front of mind—accessible to a wide range of mental states and processes—whenever it is believed to be relevant to satisfying one's desires.²⁵

Why is the motivational and attentional functioning of the desire in part iii) of the care function all that is needed to help constitute a form of direct concern for others? Recall that being directly concerned for others is achieved by being reliably disposed to pay attention to them and give them consideration in one's deliberations, independently of how they figure into one's other ends and concerns. If a subject has only the motivational and attentional aspects of the desire of part iii), in addition to realizing parts i) and ii), then this will be enough for them to be reliably disposed to pay attention to *C* and to give *C* consideration in their deliberations, regardless of how *C* figures into their other ends.

How, then, can a company realize the motivational and attentional functioning of part iii)'s desire through its decision-making and information-gathering procedures? One way is by reliably implementing procedure 2:

Procedure 2. Do not seriously consider any new policies or plans if they cause significant impediments to the well-being of some group *G* (comprising persons whose welfare is monitored via the company's realization of part ii) of the care function), unless these impediments are necessary for avoiding comparable harms to some other cared-for group *H*, along with information-gathering mechanisms that proactively gather and disseminate information about how to more effectively protect the well-being of *G*.²⁶

Procedure 2 would realize the motivational aspect of the desire of part iii), for it would dispose the company to pursue the object of part iii)'s desire, provided that the company has information about how to do so. The information-gathering mechanisms would realize the attentional aspect, because these would dispose the company to gather and make accessible information that pertains to satisfying part iii)'s desire—to keep it at the front of the company's mind, as it were.

As with the respect function, a company can realize the care function without its individual members holding attitudes that realize (any part of) the care function in their individual minds. It is enough that these members commit to doing their part to implement the decision-making and information-gathering procedures as part of their jobs. They thus may be driven by role-based and pecuniary considerations that do not stem from any attitudes that realize the care function, and the procedures may be created and enforced by leaders whose reasons for doing so similarly do not stem from attitudes that realize the care function.

view, desiring will still involve being disposed to choose and pursue what one desires, provided that one has beliefs about how to do so, and relative to the strength of one's endorsement.

²⁵ So to be clear, insofar as attention involves conscious awareness, I am referring only to what Ned Block (1995) influentially terms “access consciousness,” as distinct from “phenomenal consciousness.” This will be important for thinking about how corporate agents can attend to objects through realizing the practical functioning of desire.

²⁶ This functionalist approach sidesteps the worries about attributing desires to corporations that lead Denis Arnold (2006: 282–83) to instead posit shared intentions as the basis for corporate moral responsibility.

Here you might worry again about the distance between creditworthy motivation and affect that my view allows. If neither the corporate agent nor any of its constituent members experience the affective states associated with the goodwill-attitudes whose practical functioning the corporate agent realizes, isn't the corporate agent's concern for those impacted by its actions more accidental than not? After all, given the role-based and pecuniary motivations of its members, there would be little resistance to changing the corporate agent's procedures in ways that fail to realize anything like goodwill-attitudes. But this worry applies a misplaced standard of nonaccidentality. It asks, in effect, whether the corporate agent would still be directly and appropriately concerned for others if it had different motives. For, according to my argument, the corporate agent's motives are given by its decision-making procedures. If these change, then so do its motives. We want to evaluate, instead, whether, *holding an agent's motives fixed*, it is no accident that the agent is directly and appropriately concerned for others. It is a further question how easily an agent's motives are changed, one which I do not think bears on the agent's moral creditworthiness.²⁷

We now have a picture of how a company's decision-making procedures (and associated information-gathering and -disseminating processes) can begin realizing the respect and care functions.²⁸ This is a picture of how a company can realize the practical functioning of goodwill-attitudes, which is sufficient for realizing the sort of direct concern that constitutes morally creditworthy motivation. This is because, again, in calling for attitudes of goodwill, what the internal standards of the reactive attitudes require is motivation that involves direct concern for others and is thereby nonaccidentally attuned to how one's actions impact others.²⁹

²⁷ This point follows, in part, from the fact that no ordinary agent (corporate or natural) could have morally creditworthy motives if what is required is not simply nonaccidentally doing what is right relative to one's motives (i.e., holding one's motives fixed) but also nonaccidentally having the right motives (Markovits, 2010: 209–15). Additionally, we need to preserve the distinction between acting from morally creditworthy motives and having morally desirable character traits (by which one nonaccidentally acts from morally creditworthy motives), given that, as I discuss later, agents who do not have good character traits can still, on occasion, act from morally creditworthy motives (Markovits, 2010).

²⁸ Here I am mindful of Kirk Ludwig's (2017: n8) concerns about developing a functionalist theory of corporate psychology. I avoid appealing to a simple "parity principle" between individual and corporate agents, which Ludwig sensibly criticizes, with my argument for the moral sufficiency of the practical functioning of direct concern of goodwill-attitudes. Additionally, I do not attribute a mind to corporate agents—which would involve attributing a complicated web of cognitive and conative commitments—even though I do attribute some discrete forms of practical functioning to them (specifically, the practical functioning of goodwill-attitudes). This might help further forestall Ludwig's worries, which, as he states, are about functionalist accounts of the "mind" of organizations. Now, an underappreciated challenge for attributing discrete practical attitudes to organizations without also attributing a mind to them is posed by Donald Davidson's (1982) claim about the "holism" of mental states. It states, roughly, that attributing any given mental state to an agent requires attributing a holistic background of additional mental states sufficient for constituting a whole mind. I cannot address this concern here, beyond indicating my broad agreement with Bratman's (2017, 2022) approach to doing so.

²⁹ You might wonder whether, returning to the earlier Charity–Carey case, we could say something similar about Charity. Her enjoyment of benefiting others has enough functional overlap in her reasoning with the functioning of the attitude of care to realize creditworthy motivation. But actually, there is very little functional overlap. As is clear in my discussion of the different aspects of the practical functioning of care, it is

But the term *motivation* raises an unresolved issue for us: how is it that, even if a company realizes the respect and care functions through its procedures and is directly concerned for others in the right ways, it can be *motivated* by this direct concern? For, again, what we ultimately want to determine is whether a company can do what is right from morally creditworthy motives. And for this, it is not enough to see that a company both does what is right (for example, choosing to invest in products that better protect consumer well-being) and has procedures that realize the requisite practical functioning (for example, the care function).

This is for the same reason that it is not enough for an individual to do what is right and maintain attitudes of goodwill for others. The motives behind their right action might not include their goodwill-attitudes. They might bracket those attitudes in their decision-making and act only out of self-interest, like Kant's infamous shop-keeper. Their goodwill-attitudes must figure into their motives and guide their right actions.³⁰

I think something similar is true of companies: in addition to doing what is right and having decision-making procedures that realize the practical functioning of direct concern of goodwill-attitudes, their right action must be guided in part by their implementation of those procedures.³¹ (I say "in part," because as I mentioned early on, we can allow for mixed motives; we need not insist that a company be guided only by its procedures of direct concern.) The dairy company's development of a product that protects consumers' well-being, for example, must be guided in part by its procedures that focus on keeping consumers from harm and monitoring threats to consumer well-being. Such guidance will be a matter of how the company both initiated the right course of action and ensured its completion and success. Which procedures did the action directly issue from?³² Which procedures were used to ensure that the action was fully carried out and achieved its intended objectives? If the answer to either of these questions does not refer (in part) to the procedures that realize the goodwill-derived forms of direct concern, then the company is not motivated by its direct concern for others.

At this point, it would be useful to contrast my account with David Silver's (2005) account of corporate moral responsibility, which similarly starts with a Strawsonian premise about the internal standards of our reactive attitudes toward agents. Silver

not enough merely to reason in ways that give weight to alleviating someone's needs. But it remains true that there is some minimal overlap between Charity's motives and the practical functioning of care, and this explains why Charity deserves some moral credit, although much less than Carey.

³⁰ On the relevant notion of attitude guidance, see Rozeboom (2017: 7–9).

³¹ Note that it does not yet follow from the fact that a company is guided by procedures that realize the practical functioning of goodwill-attitudes that it acts from fully morally creditworthy motives. It will matter the extent to which it realizes this practical functioning toward *all* the pertinent individuals.

³² The term "directly led" is meant to rule out what philosophers of action call "deviant" causation, which is the causation of behavior by motives through nonrational processes. At a corporate level, this would occur when a company's decision procedures systematically create an unintended but congruent side effect, for example, perhaps the existence of the decision procedures that realize the respect function have a chilling effect, leading all managers to fearfully refrain from asking questions about employees' personal lives, leading to a congruent but deviantly caused outcome.

claims that these internal standards track “moral dispositions” that help constitute the moral character of agents and by which agents are disposed to think correctly about “the appropriate way to treat persons and other things of value such as animals” (Silver, 2005: 284). On this basis, he argues that a subset of our reactive attitudes are corporate reactive attitudes, attuned to the kinds of corporate moral dispositions that are distinctively realized by what he takes to be the corporate form of moral character—organizational cultures that lead employees to respond appropriately to the interests and rights of those impacted by corporate actions (Silver, 2005: 287–88).

As far as I can tell, what I describe as the corporate practical functioning of goodwill-derived direct concern will overlap significantly with what Silver describes as corporate moral dispositions. But I think my idea of goodwill-derived direct concern is more fundamental. That is, the moral dispositions that are relevant to satisfying the internal standards of the reactive attitudes, including corporate moral dispositions, just are those that are entailed by the forms of direct concern derived from appropriate goodwill-attitudes. We need an account of the practical functioning of goodwill-attitudes to determine what the relevant moral dispositions are.

To see why, notice that there is an important aspect of the internal standards of reactive attitudes that my view can readily capture and that Silver’s view would have trouble explaining without drawing on a prior conception of goodwill-derived direct concern. Consider excuses for bad behavior (i.e., considerations that mitigate blame for bad behavior according to the internal standards of the reactive attitudes) that show, not that the stable dispositions that make up an agent’s character were better than their behavior suggested, but rather that the agent was motivated by goodwill-derived forms of direct concern. What I have in mind are examples of loutish agents, whose characters do not contain Silver’s moral dispositions, who are stirred by genuine forms of direct concern and attempt but fail to act rightly. Think of a misanthropic person who attempts on some occasion to be complimentary but ends up making a condescending remark. Observing their bad behavior, we might initially blame them, chalking the action up to their misanthropic character. But a kind of (albeit partial) excuse, which would mitigate our blame, would point out that, despite their bad character, they were acting out of concern for the person they were trying but failing to compliment. We can imagine corporate analogues of this case: a company’s environmental track record and internal culture are poor, and it rolls out a clumsy attempt to lessen environmental harms that ends up exacerbating them instead. Despite the company’s “bad character” in Silver’s sense, if the decision-making procedures that led to this particular decision realized goodwill-derived direct concern for the natural environment, this would sensibly mitigate our blame.³³

³³ Silver may reply that, if this excuse is correct, then it shows that the corporate agent has some moral dispositions after all. Fair enough, but this just reinforces the point I am making here: what is most fundamental according to the internal standards of the reactive attitudes is not the presence of moral dispositions in general but rather the presence of direct concern, as entailed by appropriate goodwill-attitudes, which in turn involves some limited moral dispositions. Note also that, in treating moral

Another apparent point of contrast that I think ends up being mostly superficial concerns Silver's emphasis on corporate culture. Even though I have not focused on corporate culture here, I agree with Silver that a culture that leads employees to consider the interests of those impacted by corporate actions can count toward the moral creditworthiness of the organization, so long as this helps realize the practical functioning of goodwill-derived direct concern.³⁴ But two caveats are worth keeping in mind. First, unlike on Silver's view, and as I stressed earlier, it does not matter for me whether a corporate culture leads its employees to act with individual attitudes of direct concern for others; all that matters is whether the culture leads employees to do their part in realizing the corporate practical functioning of direct concern. Second, there are some dangers in how corporate culture influences employees, which should make us cautious about endorsing it as a central element of morally creditworthy corporate motivation.³⁵

4. OBJECTIONS AND IMPLICATIONS

Now consider two further, important objections to my view. First, parallel to what Amy Sepinwall (2017) alleges about the inaptness of blaming companies, you might think that it is never appropriate to give companies moral credit because they are incapable of properly responding to being given moral credit, for this response would entail affective states that companies cannot maintain. Appreciating that one has received moral credit plausibly entails an affect-laden response of moral self-satisfaction that is a positive analogue to feeling remorse for being blamed. Corporate agents cannot experience this, and so they cannot be creditworthy. My reply to this worry extends the functionalist approach I have utilized throughout this article: we can identify in the appropriate credit-appreciating response a core of practical functioning that is realizable by corporate decision-making procedures—the practical functioning of moral self-satisfaction. Now, Sepinwall (2017: 19–20) worries that this sort of functionalist account will leave too much room for corporate insincerity. But I think that, so long as the full array of practical functioning of moral self-satisfaction is realized, the corporate response to receiving moral credit will be sufficiently sincere.

dispositions as most fundamental, Silver's view ends up requiring a deeper cleavage between corporate- and individual-directed reactive attitudes than my view does, because corporate moral dispositions operate very differently from how individual moral dispositions operate (Silver, 2005: 290). On my view, there is deep continuity between the practical functioning of goodwill-derived direct concern in individual agents and in corporate agents and, thus, deep continuity in how these agents are evaluated by the internal standards of reactive attitudes.

³⁴ For instance, if Lisa Herzog (2020) is correct that corporate culture operates via the "routinization" of employee behaviors, this can help achieve the functioning of goodwill-derived direct concern if what is routinized are the various components of corporate decision-making (and information gathering) that constitute such practical functioning.

³⁵ Herzog (2020: 216–17) argues that corporate culture tends to be imposed in "implicit" ways that undermine employees' ability to critically reflect on it. This prevents them from autonomously accepting the values and norms that the culture imposes. Because decision-making procedures are explicit in what they ask employees to prioritize or ignore, they do not pose the same kind of threat to employee autonomy.

A related worry about my account is that an agent can be morally creditworthy only if it is a moral agent, and moral agency requires mental capacities that go beyond realizing the practical functioning of goodwill-derived direct concern. For instance, David Rönnegard (2015: 28–29) claims that moral agents are able to be aware of their decision-making processes and resulting intentions, and corporate agents lack such practical self-consciousness.³⁶ Perhaps so, but I doubt that we should follow Rönnegard in adding requirements for moral creditworthiness that derive from some broad notion of moral agency that goes beyond the idea of acting out of goodwill-derived direct concern for others. This doubt stems from my Strawsonian approach: we start our theorizing about moral creditworthiness by considering the internal standards of the reactive attitudes, which reveal that quality of will, and specifically goodwill-derived direct concern for others, determines whether agents are morally creditworthy. There is no further notion of being a moral agent, beyond this idea of acting with goodwill-derived direct concern for others, that places requirements on moral creditworthiness. This, anyway, is what follows from my approach here.

A second, different source of concern about my view is that attributing morally creditworthy motives to companies entails giving them moral status they clearly do not have.³⁷ This objection can be summarized as follows: having moral capacities is what grounds having the full moral standing (worth, value, rights, etc.) of a person. If a company can have morally creditworthy motives, then it has whatever moral capacities ground the full moral standing of a person.

To address this issue, I want to signal my agreement with those who think we can pry apart the question of whether a company can exercise robust forms of moral agency—such as morally creditworthy motivation—from the question of whether it has the full moral standing of a person (see esp. Hess, 2018: 83–84; Hindriks, 2014: 15; List, 2018: 38–40). One reason to separate these questions is that the protections and recognition afforded by moral standing may not derive simply from our agential capacities, such as those needed to act from morally creditworthy motives, but also from our relational vulnerabilities. These vulnerabilities plausibly depend on our phenomenal awareness of our dependence on others and our associated emotional dispositions (to feel sympathy, attachment, betrayal, etc.). Without these phenomenal and emotional

³⁶ A further, related concern is about the causal grounds of corporate motivation. Given that legal compliance plays a large role in causing corporate agents to have the kinds of procedures that might realize the practical functioning of goodwill-attitudes, you might think that corporate agents tend to realize this functioning on the wrong kind of basis, one that diminishes their moral creditworthiness. I think this worry ignores how individual human agents often form goodwill-attitudes in similar ways. Our parents teach us to respect and care about others by threatening to punish us for failing to do so, for instance. So long as these attitudes eventually take hold in our minds, no longer depending on the threat of punishment to shape our thinking, they can serve as the basis of morally creditworthy motivation. I think the same point applies to the practical functioning of goodwill-derived direct concern in corporate agents. It can constitute morally creditworthy corporate motivation even if it was initially formed in response to the threat of legal punishment.

³⁷ You might think this problem quickly arises from the fact that being morally creditworthy implies having the second-personal standing to demand moral credit from others. But remember that deserving moral credit does not entail such second-personal standing; moral credit is not something that is generally owed to morally creditworthy agents (even if it is a mistake not to give them moral credit).

dispositions, companies are not relationally vulnerable and thus do not have full moral standing, even though they retain the capacity to have morally creditworthy motives—or so I am inclined to think; much more work would be needed to settle this matter fully.

Even if companies are not relationally vulnerable to us, we are vulnerable to them, and this is (one reason) why the moral quality of their motives should matter to us. To close, then, let me begin exploring the import of a company's morally creditworthy motivation (or lack thereof) for one's relationship with that company. Return to Scanlon's (2008) view of the relational "meaning" of actions as a starting point. Scanlon focuses primarily on the meaning of blameworthy actions: he claims that blameworthy wrongdoing "indicates" something about the wrongdoer's attitudes that "impairs" his relationships with his victims. Impairment is the violation of the standards implied by the "normative ideal" of a given relationship, and this violation makes it appropriate for victims to shift their reactive attitudes toward the wrongdoer away from the amicable and trusting attitudes "the relationship normally involves" (135). What does this suggest about the relational meaning of morally creditworthy actions? Scanlon does not directly say, but here is one natural way to extend his account: when someone acts from creditworthy motives, this indicates something about the agent's attitudes that reinforces her relationships with those she thereby treats well. She exemplifies the normative ideal of her relationship with them, which makes it appropriate for them to more unreservedly maintain the amicable attitudes that their relationship has standardly involved—what I have been describing broadly as trust, appreciation, and gratitude.

What does this imply about someone's relationship with a company? To fully answer that question, we would need to know what normative ideal(s) applied to the relationship. These ideals will vary widely across the range of one's possible ties to a company. Let me consider just one example to illustrate how the Scanlonian view might extend to individual-company relationships.³⁸ Suppose that you are leading an effort to improve the environmental cleanliness of public spaces in your city, including your city's central river, which remains heavily polluted from historical practices (now halted) of factories dumping industrial waste into it. You solicit the help of the companies responsible for this waste dumping, proposing that they split the cost of dredging the river. All but one of the companies declines, but the company that agrees—paper manufacturer Papyrus—decides to shoulder the entire cost of the dredging operation. The dredging operation goes well, and Papyrus follows through on its commitment, even when the operation ends up costing more than initially expected. You now have a relationship with Papyrus. To apply Scanlon's account, let us assume that your relationship is governed by a normative ideal of social solidarity, whereby agents regard one another as trusting, respected partners who share and/or are concerned about solving the same social problems. Our question is how your relationship of social solidarity with Papyrus should be modified,

³⁸ Scanlon (2008, 162–65) briefly discusses such relationships and suggests, in line with my view, that the configuration of a company's decision-making structure will help determine the level of trust it warrants from others.

according to Scanlon's account, depending on whether Papyrus had morally creditworthy motives, according to my account.

Suppose it did not. Papyrus's expenditure of resources on the dredging operation was not guided by any procedures that realize the practical functioning of goodwill-derived direct concern. Instead, its guiding procedures aimed only at improving its local reputation, to recruit and retain more reliable employees. You may still feel some gratitude to Papyrus. But this gratitude would be somewhat tempered by the knowledge that its actions were focused only on shoring up its local reputation. You would remain reserved about trusting Papyrus as a partner for future endeavors and reluctant to publicly commend Papyrus. (Although you may still have pragmatic reasons to publicly praise the company.) You would remain nearly as guarded in your admiration for, and eagerness to associate with, Papyrus as you were prior to Papyrus's funding of the river-dredging effort. This is all because, on the Scanlonian view, Papyrus's motives did not reinforce the normative ideal of your relationship of social solidarity with the company. Papyrus did not act out of direct concern for the community members whose health and safety are tied to the river cleanup effort, nor did it act out of direct concern for you, a trusting partner who shares an interest in solving an environmental problem that it helped create.

By contrast, suppose that Papyrus did act out of direct concern for impacted community members and you. That is, its decision-making procedures realized the practical functioning of goodwill-derived forms of direct concern for the community members and you, and these were what guided Papyrus' contribution to the river-dredging effort. In being so motivated, Papyrus's actions would reinforce the normative ideal of your relationship of social solidarity.³⁹ This would make it appropriate for you to be more freely interested in working with Papyrus in the future, to publicly commend Papyrus, and to more generally trust Papyrus as a partner for solving social problems. Papyrus's morally creditworthy motives would support a very different orientation to your ongoing relationship with the company compared to when its motives were not creditworthy.

This is a brief illustration of one of the relational implications of my account of morally creditworthy corporate motivation. This account gives pride of place, again, to the practical functioning of goodwill-attitudes, such as respect and care. Such functioning is sufficient for being directly concerned for those impacted by one's actions, and it can be realized by the decision-making procedures of companies.⁴⁰ This direct concern is what we care about in evaluating the relational import of agents' actions using the internal standards of the reactive attitudes, which require that agents be nonaccidentally attuned to those importantly impacted by their actions. What I have tried to provide, then, is a plausible, functionalist model of

³⁹ I have not provided, much less defended, a specification of the normative ideal of a relationship of social solidarity. But this is not needed to reach the following conclusions about the relational significance of creditworthy motivation.

⁴⁰ What about business-to-business relationships? I do not think that Scanlon's account of the relational meaning of actions applies to business-to-business relationships, because companies are not relationally vulnerable to one another (we might say "betrayable") in the ways that Scanlon's account presupposes.

how companies can exhibit the aspects of goodwill that matter when we wonder, *But did they do it for the right reasons?*

Acknowledgments

I thank audiences at an Oxford R:ETRO seminar, hosted by Alan Morrison and Rita Mota; a Zicklin Normative Business Ethics Workshop hosted by the Wharton School; a Central American Philosophical Association conference symposium session, which included very helpful comments by Ritwik Agarwal and Kurt Nutting; and the 2021 Society for Business Ethics annual meeting. I am also grateful for correspondence and conversations with Michael Bratman, Gideon Pogrud, Facundo Alonso, Samuel Asarnow, Carlos Núñez, and Julian Jonker, and for research assistance by Jocelyn Villalta-Lua and Chloe Wilson. Finally, I benefited from the careful comments and suggestions offered by two anonymous referees for *Business Ethics Quarterly* and by its senior associate editor, Denis Arnold.

REFERENCES

- Arnold, D. 2006. Corporate moral agency. *Midwest Studies in Philosophy*, 30(1): 279–91.
- Arpaly, N., & Schroeder, T. 2013. *In praise of desire*. Oxford: Oxford University Press.
- Baddorf, M. 2017. Phenomenal consciousness, collective mentality, and collective moral responsibility. *Philosophical Studies*, 174(11): 2769–86.
- Benhabib, S. 1985. The generalized and the concrete other: The Kohlberg–Gilligan controversy and feminist theory. *Praxis International*, 5(4): 402–24.
- Block, N. 1995. On a confusion about the function of consciousness. *Behavioral and Brain Sciences*, 18(2): 227–47.
- Bratman, M. 1987. *Intentions, plans, and practical reason*. Cambridge, MA: Harvard University Press.
- Bratman, M. 2017. The intentions of a group. In E. Orts & C. Smith (Eds.), *The moral responsibility of firms*: 36–52. Oxford: Oxford University Press.
- Bratman, M. 2022. *Shared and institutional agency: Toward a planning theory of human practical organization*. Oxford: Oxford University Press.
- Chapman, B., & Sisodia, R. 2015. *Everybody matters: The extraordinary power of caring for your people like family*. New York: Penguin Random House.
- Darwall, S. L. 2002. *Welfare and rational care*. Princeton, NJ: Princeton University Press.
- Darwall, S. L. 2006. *The second person standpoint: Morality, respect, and accountability*. Cambridge, MA: Harvard University Press.
- Darwall, S. L. 2013. Authority and reasons: Exclusionary and second personal. In *Morality, authority, and law: Essays in second-personal ethics I*: 151–67. Oxford: Oxford University Press.
- Davidson, D. 1982. Rational animals. *Dialectica*, 36(4): 317–27.
- di Norcia, V., & Tigner, J. 2000. Mixed motives and ethical decisions in business. *Journal of Business Ethics*, 25(1): 1–13.
- French, P. 1984. *Collective and corporate responsibility*. New York: Columbia University Press.
- French, P. 1996. Integrity, intentions, and corporations. *American Business Law Journal*, 34(2): 141–55.
- Gelles, D. 2018. Memo from the boss: You're a vegetarian now. *New York Times*, July 20.

- Helm, B. 2017. *Communities of respect: Grounding responsibility, authority, and dignity*. New York: Oxford University Press.
- Herman, B. 1993. On the value of acting from the motive of duty. In *The practice of moral judgment*: 1–22. Cambridge, MA: Harvard University Press.
- Herzog, L. M. 2020. Citizens' autonomy and corporate cultural power. *Journal of Social Philosophy*, 51(2): 205–30.
- Hess, K. 2018. Does the machine need a ghost? *Journal of the American Philosophical Association*, 4(1): 67–86.
- Hess, K., & Björnsson, G. 2017. Corporate crocodile tears? On the reactive attitudes of corporate agents. *Philosophy and Phenomenological Research*, 94(2): 273–98.
- Hindriks, F. 2014. How autonomous are collective agents? Corporate rights and normative individualism. *Erkenntnis*, 79(9): 1565–85.
- Johnson King, Z. 2020. Accidentally doing the right thing. *Philosophy and Phenomenological Research*, 100(1): 186–206.
- Jonker, J. 2020. Directed duties and moral repair. *Philosophers' Imprint*, 20(23): 1–32.
- Kant, I. 2008. *Groundwork for the metaphysics of morals* (A. Wood, Trans. & Ed.). New Haven, CT: Yale University Press.
- List, C. 2018. What is it like to be a group agent? *Nous*, 52(2): 295–319.
- List, C., & Pettit, P. 2011. *Group agency: The possibility, design, and status of corporate agents*. Oxford: Oxford University Press.
- Ludwig, K. 2017. Do corporations have minds of their own? *Philosophical Psychology*, 30(3): 269–301.
- Markovits, J. 2010. Acting for the right reasons. *Philosophical Review*, 119(2): 201–42.
- Massoud, A. 2016. Moral worth and supererogation. *Ethics*, 126(2): 690–710.
- McKenna, M. 2012. *Conversation and responsibility*. New York: Oxford University Press.
- Morrison, A., Mota, R., & Wilhelm, W. 2022. Relationships, authority, and reasons: A second-personal account of corporate moral agency. *Business Ethics Quarterly*, 32(2): 322–47.
- Noddings, N. 1984. *Caring: A feminine approach to ethics and moral education*. Berkeley: University of California Press.
- Pettit, P. 1995. The cunning of trust. *Philosophy and Public Affairs*, 24(3): 202–25.
- Raz, J. 1986. *The morality of freedom*. Oxford: Oxford University Press.
- Raz, J. 2006. The problem of authority: Revisiting the service conception. *Minnesota Law Review*, 90: 1003–44.
- Rönnegard, D. 2015. *The fallacy of corporate moral agency*. New York: Springer.
- Rozeboom, G. 2017. The motives for moral credit. *Journal of Ethics and Social Philosophy*, 11(3): 1–29.
- Rozeboom, G. 2022. Roles, Rousseau, and respect for persons. *Social Theory and Practice*, 48(4): 771–97.
- Scanlon, T. M. 2008. *Moral dimensions: Permissibility, meaning, and blame*. Cambridge, MA: Belknap Press.
- Schapiro, T. 2011. Foregrounding desire: A defense of Kant's incorporation thesis. *Journal of Ethics*, 15(3): 147–67.
- Sepinwall, A. 2017. Blame, emotion, and the corporation. In E. Orts & C. Smith (Eds.), *The moral responsibility of firms*: 143–66. Oxford: Oxford University Press.
- Setiya, K. 2007. *Reasons without rationalism*. Princeton, NJ: Princeton University Press.
- Shapiro, S. 2011. *Legality*. Cambridge, MA: Belknap Press.
- Shoemaker, D. 2013. Qualities of will. *Social Philosophy and Policy*, 30(1–2): 95–120.

- Silver, D. 2005. A Strawsonian defense of corporate moral responsibility. *American Philosophical Quarterly*, 42(4): 279–93.
- Singh, K. 2020. Moral worth, credit, and non-accidentality. In M. Timmons (Ed.), *Oxford studies in normative ethics* Volume 10, 156–81. Oxford: Oxford University Press.
- Sinhababu, N. 2017. *Humean nature: How desire explains action, thought, and feeling*. Oxford: Oxford University Press.
- Sliwa, P. 2016. Moral worth and moral knowledge. *Philosophy and Phenomenological Research*, 93(2): 393–418.
- Smith, H. M. 1991. Varieties of moral worth and moral credit. *Ethics*, 101(2): 279–303.
- Smith, M. 1994. *The moral problem*. Malden, MA: Blackwell.
- Stark, A. 1993. What's the matter with business ethics? *Harvard Business Review*, May–June.
- Stratton-Lake, P. 2000. *Kant, duty, and moral worth*. New York: Routledge.
- Strawson, P. F. 2008. Freedom and resentment. In *Freedom and resentment and other essays*: 1–28. New York: Routledge.
- Tollefsen, D. 2008. Affectivity, moral agency, and corporate–human relations. *APA Newsletter on Philosophy and Law*, 7(2): 10–13.
- Tollefsen, D. 2015. *Groups as agents*. Hoboken, NJ: John Wiley.
- Vargas, M. 2016. Responsibility and the limits of conversation. *Criminal Law and Philosophy*, 10(2): 221–40.
- Velleman, J. D. 1989. *Practical reflection*. Princeton, NJ: Princeton University Press.
- Velleman, J. D. 1999. Love as a moral emotion. *Ethics*, 109(2): 338–74.
- Watson, G. 1987. Responsibility and the limits of evil: Variations on a Strawsonian theme. In M. McKenna & P. Russell (Eds.), *Free will and reactive attitudes: Perspectives on P. F. Strawson's "Freedom and resentment"*: 115–42. Burlington, VT: Ashgate.
- Williams, B. 1981. Persons, character and morality. In *Moral luck: Philosophical papers 1973–1980*: 1–19. Cambridge: Cambridge University Press.
- Zaveri, P. 2020. Salesforce plans to double its number of Black leaders in the US by 2023. *Business Insider*, July 1. <https://www.businessinsider.com/salesforce-double-black-leaders-employees-racial-justice-2020-7>.

...

GRANT ROZEBOOM (gjr5@stmarys-ca.edu) is an assistant professor of business ethics and social responsibility in the Department of Organizations and Responsible Business at Saint Mary's College of California. His research concerns a wide range of issues in normative and applied ethics, including relational equality and the workplace, the basis of moral equality, the attitude of respect for persons, and morally creditworthy motivation. He has published in *Ethics*, *Philosophy and Public Affairs*, and the *Journal of Business Ethics* and with Oxford University Press. He received a BA in philosophy from the University of Northern Iowa and a PhD in philosophy from Stanford University.

This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.