

ARTICLE

How Much Can We Ask of Collective Agents?

Stephanie Collins

Dianoia Institute of Philosophy, Australian Catholic University, Melbourne, Victoria, Australia
Email: stephanie.collins@acu.edu.au

Abstract

Are obligations of collective agents—such as states, businesses, and non-profits—ever overdemanding? I argue they are not. I consider two seemingly attractive routes to collective overdemandingness: that an obligation is overdemanding on a collective just if the performance would be overdemanding for members; and that an obligation is overdemanding on a collective just if the performance would frustrate the collective's permissible deep preferences. I reject these. Instead, collective overdemandingness complaints should be reinterpreted as complaints about inability or third-party costs. These are not the same as overdemandingness. Accordingly, we can ask an awful lot of collective agents.

Keywords: Collective agency; collective obligation; overdemandingness; ought implies can

Introduction

Many philosophers believe there are limits to morality's demands. For example, can morality demand that I give to poverty relief to the point of marginal utility? This seems extreme. Any such demand ignores the prerogatives that derive from my personal point of view (Scheffler 1982; *cf.* Berkey 2016) or derogates the value of fairness, because I wouldn't need to give so much if everyone did their share (Murphy 2000; *cf.* Ridge 2010), or infringes my permission to pursue my projects (Cullity 2004; *cf.* Stroud 2013), or undermines my very ability to be a moral agent (Herman 2002; Thomas 2005; Thomas 2009).¹

The debate about overdemandingness has focused on individuals. Yet, recently, philosophical interest in collective agents has exploded. According to an increasing number of philosophers, collective agents include (at least) states, businesses, and nonprofits. These entities are agents insofar as—roughly—their individual constituents are *pro tanto* committed to abiding by a rational group-level decision-making procedure, which allows the group to arrive at (functional equivalents of) beliefs, preferences, intentions, decisions, and other mental states, whose content is different from members' mental states (List and Pettit 2011). These mental states constitute a “rational point of view”—a bundle of interlocking mental states from which the collective acts and that the collective seeks to maintain (Rovane 1998). Via its procedure and rational point of view, a collective can distribute roles to members that, when performed, amount to actions that are properly attributable to—“incorporated by”—the collective (French 1984). Collectives can thereby act. If the collective's decision-making procedure can process moral inputs, then it can do all this while being sensitive to morality's demands (Hindriks 2018; Collins 2019, chap. 6). I will assume this picture (or something like it) is correct. Some collectives can bear moral obligations.

This paper unites these two literatures—on moral overdemandingness and on collective agency—to answer the question: How much can we ask of collective agents? More specifically: Can collective

¹On overdemandingness generally, see Chatterjee (2004) and Chappell (2009).

agents legitimately make overdemandingness complaints that aren't equivalent to (a bundle of) individual overdemandingness complaints?² Intuitively, the answer is yes. Let us consider three examples.

First, Peter Singer has asked whether art galleries have moral obligations to donate some of their proceeds to poverty relief. He concludes they do not: "They were set up for a different purpose" (2015, 123). Intuitively, art galleries have obligations to not actively cause harm. And perhaps they have minimal obligations to make art accessible for people without much money. But extensive obligations regarding the amelioration of poverty (and other structural injustices) seem to ask too much.

Second, consider a small, family-owned gas station. Given catastrophic climate change, it would be better if this company sold sustainable energy (perhaps for electric cars). Yet the company's margins are slim, entry costs to a new industry are high, another gas station would soon take its place, and moral obligations to curtail the gas industry are—we usually think—held by regulators, consumers, and voters. These people have moral obligations to use laws or wallets to rationally induce our gas station to redirect its operations, while materially supporting its ability to do so. Requiring a small gas station to change direction under its own steam might cause the business to go under. How can such a company be expected to do that? It's too much to ask.³

Third, consider a state's purported moral obligation to extend the vote to foreigners that are affected by its laws. This might undermine the sense of political community within the state.⁴ If so, then surely the state cannot have an obligation to extend its franchise to foreigners. The weakening of political community is simply too much to ask. Or so it seems.

Can collective agents legitimately make overdemandingness complaints that aren't equivalent to (a bundle of) individual overdemandingness complaints? I argue not. Section 1 begins by characterising the relevant collectives, emphasising how their agency is distinct from members' agency. This clarifies how complaints of collectives might be nonequivalent to bundles of individuals' complaints, while explaining how individuals constitute collectives.

Section 2 considers an initially attractive proposal: an obligation is overdemanding on a collective just if it's overdemanding on at least one member. This view is inadequate because collectives have their own distinct rational points of view. Thus, high demands on one or more members—even high demands on each and every member—is neither necessary nor sufficient for a high demand on the collective. This conceptualisation cannot produce a 'yes' answer to our question.

Section 3 considers a view that makes collectives' overdemandingness complaints analogous to (rather than mutually implicative of) individuals' overdemandingness complaints: an obligation is overdemanding on a collective just if the obligation frustrates the collective's *permissible deep preferences*. Permissible deep preferences are long held by the collective, ground many of the collective's other preferences, are viewed by the collective as highly important, and are permissible when considering aspects of morality other than overdemandingness. I argue that three prominent arguments that shield individuals' permissible deep preferences from morality's demands do not have analogues for collectives. And it looks extremely unlikely that collectives' permissible deep preferences are protection-worthy in a *sui generis* way. This proposal also fails to vindicate collectives' overdemandingness complaints.⁵

Section 4 reconceptualises collectives' legitimate overdemandingness complaints as complaints that an obligation violates the collective's constitutional constraints. Legitimate collective

²Contrast Pinkert (2018), who analyses collective overdemandingness complaints that *are* equivalent to individual overdemandingness complaints. Pinkert accepts some collective overdemandingness complaints might not be equivalent (178).

³Scott Wisor (2014) argues the fossil-fuel divestment movement "essentially call[s] on [energy companies] to stop being energy companies," implying this is ridiculous.

⁴Goodin (2007) defends the "all-affected principle"; Miller (2009) offers a community-based alternative.

⁵Berkey (2019) considers proposals like those in sections 2 and 3. I'll explain how Berkey's discussion differs from mine. My positive proposal in section 4 has no analogue in Berkey's discussion.

overdemandingness complaints thus collapse into *obligation implies ability* complaints. They should be given as much credence as the latter principle. Whether that principle is true is independent of overdemandingness, so there are no legitimate collective overdemandingness complaints as such. What's more, constitutional constraints occur far less often than we might think—at least for the large organisations that dominate our social, political, and economic world—due to the diachronicity of many obligations and the informality of many collective decision-making procedures. The upshot is that we can ask an awful lot of collective agents.

I. Collectives and their members

Before conceptualising collectives' overdemandingness complaints, we must conceptualise their agency. I'll assume a collective is constituted by agents united under a group-level, rationally operated, distinct decision-making procedure (Collins 2019, chap. 6). These agents are the collective's "members." To explain, I'll work backwards through "united under a group-level, rationally operated, distinct decision-making procedure."

A collective's *decision-making procedure* takes inputs from members regarding what the collective should intend, believe, prefer, hope, regret, and so on. The procedure processes those inputs to produce decisions on what the collective intends, believes, prefers, hopes, regrets, and so on. A decision-making procedure need not be highly regimented—it can change over time, be complexly multifaceted, and be variously formal versus informal (Hess 2018, 37–38). It might include voting, committees, decrees, discretion, debates, conversations, bargains, horse-trades, compromises, and so on. Some procedures—and, certainly, some actions on their basis—require material objects, not just human members (Epstein 2015).

Agents are *united under a group-level* and *distinct* decision-making procedure when: (1) each is *pro tanto* committed (perhaps tacitly) to abide by the procedure's results; (2) the procedure's inputs, and how it processes inputs to form decisions, relies on the behaviour (e.g., votes, decrees, compliance) of the agents, while being distinct from the inputs, and procedures (if any), that any of the agents uses when deciding for themselves; and (3) the enactment of at least some of the collective's decisions requires behaviours by the agents, where those behaviours are also properly attributed to the collective. To illustrate (2), suppose a collective takes as inputs the beliefs of *all* members and processes them using the method of *conversation-based consensus*. If no member uses all those inputs, processed in that way, when deciding what she herself will believe, then the collective's inputs and procedures are distinct from those of any member.⁶

If a group's procedure is *rationaly operated*, it's operated with the aim of ensuring that current decisions accord with current inputs and past decisions. In any decision-making entity—individual or collective—there is the possibility of conflict between current inputs, past decisions, and potential decisions. A rational entity adjudicates these conflicts—deciding which past decisions will stay, which will go, and which will be revised—before it settles potential decisions. Such adjudication is rational operation.

In practice, this adjudication means the collective's current decisions are likely to depart from what many members would choose, if each member were deciding for the collective.⁷ For example, if an art gallery has a long-standing preference for displaying contemporary art, then it might be rational for the gallery to maintain this preference even if some, most, or even all current curators would prefer the gallery to display noncontemporary art. Perhaps they each want to display a different kind of noncontemporary art, so contemporary emerges as the compromise position. The collective's preference (contemporary art)—and its procedure for deciding on that preference (compromise)—departs from that of members, who would each do something different if

⁶In Collins (2019, chap. 6), I explain how underlings in hierarchical collectives are nonetheless members.

⁷List and Pettit (2002) demonstrate this formally.

determining the collective's decision. The possibility of such departure is crucial, since—section 2 will explain—it introduces the possibility of an obligation being overdemanding on a collective without being overdemanding on any member.

With collectives in hand, we can turn to our question: Can collective agents legitimately make overdemandingness complaints that aren't equivalent to (a collection of) individual overdemandingness complaints?

2. Membership overdemandingness

2.a The proposal

A simple proposal says a collective can legitimately make an overdemandingness complaint just in case the enactment of its purported obligation would be overdemanding for at least one member. Call this proposal *Membership Overdemandingness*.⁸ This is a tempting view of the introduction's examples: perhaps *curators* could complain if the Louvre devoted its proceeds to poverty relief; perhaps the *owners* could complain if Mom-and-Pop Gas devoted itself to electric cars; perhaps *citizens* could complain if Sweden extended the franchise to foreigners. If curators, owners, and citizens are "members" of art galleries, gas stations, and democratic states (which seems plausible), then perhaps the overdemandingness complaints mentioned in the introduction rise and fall with the overdemandingness complaints of members.

Our guiding question was: Can collective agents legitimately make overdemandingness complaints that *aren't equivalent to* (a bundle of) individual overdemandingness complaints? Membership Overdemandingness might seem to allow an equivalence thereby evading the question. But a biconditional doesn't entail equivalence. Perhaps a collective's overdemandingness complaint arises if and only if member(s) have overdemandingness complaints while being distinct from them. Such distinctness could arise via two routes: the collective's complaint might be *multiply realisable* in members' complaints, and/or there might be *systemic relations* between members' complaints (Elder-Vass 2007; List and Spiekermann 2013). It's worth exploring these, to see how Membership Overdemandingness might work.

Unfortunately, the multiple realisability route won't get us distinctive collective-level complaints. If the multiple realisability of some collective-level phenomenon is to render it distinct from its individual-level realisers, then the different potential individual-level realisers must be of different types, not just different tokens of the same type. Consider the analogous case in philosophy of mind: suppose my pain is always realised by C-fibres firing in my brain, though different C-fibres fire on different pain-occasions (because C-fibres degenerate and are replaced by new ones). Then, pain is equivalent to C-fibres firing. For pain to be nonequivalent to (i.e., distinct from) its neural realisers, there must be different *types* of potential neural realisers. According to Membership Overdemandingness, collective overdemandingness complaints are always realised by the same type of thing: individuals' overdemandingness complaints. This makes collectives' overdemandingness complaints equivalent to individuals' complaints. So this isn't how to understand Membership Overdemandingness.

The second route to nonequivalence is on firmer ground. This route says collectives have distinct complaints because there are *systemic relations* between members' complaints, even though the

⁸Berkey (2019, 119–21) rejects the view that collectives' overdemandingness complaints are generated by *aggregating* demandingness complaints of members. On the view I'm considering, collectives have overdemandingness complaints just if one, some, or perhaps every member *individually* has an overdemandingness complaint. Berkey rejects the aggregation view because collectives are not experiencing subjects. I'll reject member overdemandingness because it doesn't take collectives' distinctive agency seriously. Berkey briefly rejects the view I'm considering, because individuals' complaints will never be proportionate to the good collectives can do—but this assumes a controversial view on which overdemandingness is proportionate to the good an obligation could do, rather than an absolute floor no obligation can infringe.

collective's complaint rises and falls with members' complaints.⁹ The systemic relations produce a network of individuals' complaints, where the network is over and above the bundle of complaints in the network. The network is "over and above" because the relations between the complaints are also part of the network. Thus we can identify a collective's complaint with a systemic network of member complaints.

For example, perhaps Sweden's complaint against franchise-extension arises if and only if there are citizens' complaints against franchise-extension such that the citizens' complaints are related by having the same source. The source is the laws and norms of extended-franchise Sweden. This common source systemically relates the members' complaints to each other, rendering members' complaints components in a holistic network of complaints that constitutes Sweden's complaint. The network of complaints is more than the sum of its member-complaint parts. Thus Membership Overdemandingness gives rise to distinctive collective-level complaints via a tweaking of the biconditional: a collective can legitimately make an overdemandingness complaint just in case the enactment of its purported obligation would be overdemanding for *more than one member in systemically related ways*.

However, the proposed biconditional doesn't hold. A collective's obligation being overdemanding on more than one member (in systemically related ways) is neither sufficient nor necessary for that obligation being overdemanding on the collective. Even if a collective's obligation is overdemanding on *every* member (in systemically related ways), that's neither necessary nor sufficient for the obligation to be overdemanding on the collective. I'll now show this.

2.b Insufficiency

First, a collective's obligation being overdemanding on members (even with systemic relations between their complaints) is not *sufficient* for its being overdemanding on the collective. This is clearest in organisations with altruistic goals, such as nonprofit organisations pressing for justice. The collective's ardent pursuit of justice might be extremely demanding on members because members have goals other than justice (e.g., leisure). If the collective's only goal is justice, then any action members can take toward that goal will not be overdemanding for the collective—considered as an entity picked out by its distinct bundle of (functionally characterised) beliefs, preferences, hopes, regrets, and intentions, and the procedure that produces that bundle when enacted by members. Justice-pursuing actions that infringe on members' private lives are a boon to the collective: these actions pursue the collective's only goal. Yet such actions might be overdemanding for members—perhaps even for every member—with complaints being systemically related via a common cause (namely, the organisation's justice goal).

Perhaps the organisation's obligation does dissolve if that obligation is overdemanding on one, some, or every member. But that dissolution should not be conceptualised as the obligation's being *overdemanding* on the organisation. Instead, if the organisation's obligation dissolves due to its overdemanding effects on members' private lives, then the obligation dissolves due to detrimental effects on third parties. Given that (ex hypothesi) members' leisure is not a goal of the organisation, those detrimental effects are not effects on individuals qua members. They are, instead, effects on individuals qua third parties, considered external to the organisation. If an obligation of any agent has overly detrimental effects on external parties, then the obligation may dissolve. But it's false that the external agents' complaints make the obligation *overdemanding*. For example, if my keeping a promise to my friend would cause you great suffering, then your suffering may dissolve my obligation. But your suffering isn't sufficient for my obligation being overdemanding on me. Things are analogous for members' private suffering and the dissolution of organisations' obligations.

One might say: if the organisation valued justice, then it would recognise that member burnout is detrimental to justice. It would therefore value members' well-being. Thus, infringements on

⁹This route requires that more than one member has a legitimate complaint.

members' well-being would be internal costs to the organisation (not external costs on third parties), and high infringements on members' well-being would indeed be sufficient for an organisation's overdemandingness complaint. Membership Overdemandingness would be vindicated.

In reply: for this to create a *legitimate* group-level overdemandingness complaint, we would need an argument for upholding the organisation's justice goal. After all, according to the objection, the organisation values members' well-being only as a means to justice. If the organisation values members' well-being intrinsically, then we'd need an argument for upholding that goal, too. I'll consider such arguments in the next section, under the heading of upholding organisations' permissible deep preferences.

A different defence of Membership Overdemandingness runs as follows. Perhaps members' overdemandingness (e.g., lack of leisure) is a cost internal to the collective, and therefore sufficient for a collective overdemandingness complaint, because members are not really external agents to their collectives. When a collective has an obligation, this plausibly entails obligations for at least one member (Collins 2019, chap. 7). Call these *membership obligations*. The costs of membership obligations (i) must be imposed upon members by themselves (which is a feature of overdemanding obligations [McElwee 2017]) and (ii) are not held by members qua third parties (so the obligations' demands are internal to the collective). The sufficiency half of Membership Overdemandingness appears vindicated.¹⁰

To respond, we must distinguish the different parts of an individual's life: parts internal to, and external of, her collective. When an individual discharges her overdemanding membership duty, she inhabits one part, and imposes costs upon herself in another part. Although the cost imposer is internal to the collective, the cost experiencer (the part that wants leisure) is external. Here I follow Carol Rovane's ontology, under which "when human beings achieve rationality together at the level of the whole group, this tends to produce a certain kind of *rational fragmentation* in their lives, because not all of the thoughts and actions associated with their brains and bodies proceed from the same point of view—some proceed from the group agent's point of view, while others proceed from a point of view that is somewhat smaller than human size" (2014, 1665; similarly Rovane 1998; Rovane 2004).

Once we separate the cost imposer from the cost experiencer—though both are housed within one human life—we see that the cost experiencer remains external to the collective, even though the costs are self-imposed from the perspective of the whole human life. But the costs are not self-imposed by the collective upon the collective, or by the individual-qua-member upon the individual-qua-member. The complaints of the individual-qua-private-person are insufficient for complaints of the collective because this particular part of the individual's agency is external to the collective.

Thus, a collective obligation's being overdemanding on "a member" is triply ambiguous: Is it overdemanding on member-qua-member, member-qua-private-person, or member-qua-whole-human? The member-qua-member is the only one whose complaints are truly internal to the collective. But the member-qua-member pursues the collective's goals. So, any complaints made from within this role must be complaints that the *collective's* goals are infringed, not complaints about leisure. But section 3 will argue that infringements of a collective's goals don't produce legitimate overdemandingness complaints. Meanwhile, complaints of members-qua-private-persons or members-qua-whole-humans are external to the collective. Member overdemandingness is insufficient for collective overdemandingness.¹¹

¹⁰I thank an anonymous reviewer for this suggestion.

¹¹We might say the same about individuals' overdemandingness: perhaps costs experienced by me-qua-consumer do not create overdemandingness complaints for me-qua-charitable-giver (or for me-qua-whole-private-human) because me-qua-consumer is separate from me-qua-charitable-giver. (I thank an anonymous reviewer for this.) In response: the pressure towards rational unification is stronger across different aspects of an individual's personal life than it is across that life and the roles the individual plays in collective agents. The different aspects of my personal life are internal to me-qua-whole-private-human in a way that no aspect of my personal life is internal to (say) my university. Costs to me-qua-consumer can therefore create demandingness complaints for me-qua-whole-private-human. Undoubtedly, this issue warrants more discussion than I have space for.

2.c Nonnecessity

Member overdemandingness is also unnecessary for collective overdemandingness. Complaints of members-qua-private-persons and members-qua-whole-humans can contradict the goals of a complaining collective, as the justice example demonstrates. So it would be strange indeed if these external parties' complaints were *necessary* for a collective's complaints. Meanwhile, complaints of members-qua-members are complaints made with the collective's goals in view. Such member complaints might be necessary for collectives' complaints, but assessing the *legitimacy* of those complaints requires us to assess whether the collective's goals should be protected. That's the task of the next section.

This might look suspicious. Above I said collectives' obligations entail members' obligations and that member actions are incorporated by collectives. Plausibly, when members act within, because of, and while performing their role, their actions are incorporated by the collective (Collins 2018). If that's right, where could collectives' overdemandingness complaints possibly come from if not from the member actions that constitute collective actions?

From the frustration of the collective's preferences. Suppose all Sweden's members become indifferent to whether Sweden extends its franchise abroad, including all costs this would impose on them. Nonetheless, suppose (*arguendo*) that such franchise extension would undermine Sweden's sense of community. Then extending the franchise would undermine one of Sweden's preferences, namely, retaining its sense of community. It would be overdemanding on *Sweden itself*, even if members individually endorse the action (and the resulting loss of community). This example provides a concrete demonstration that member overdemandingness is unnecessary for collective overdemandingness. To attempt to vindicate collectives' overdemandingness complaints, we must turn to the collective's preferences.

3. Permissible deep preferences

3.a The proposal

The preceding section emphasised that collectives have distinct preferences, whose content can depart from all members' preferences. Given this, perhaps an obligation places a demand on a collective just in case it frustrates one of the collective's preferences. And an obligation is *over-demanding* on a collective just in case the frustrated collective preference is a "permissible deep" preference. Call this proposal *Deep Preferences*.

A deep preference is long held by its bearer, grounds many of the bearer's other preferences, and is held by the bearer as amongst the bearer's most important preferences. This is a high bar for a preference to meet, including objective and subjective elements. The height of the bar is intentional: the more conditions a preference must meet to count as deep, the more compelling is the idea that it's overdemanding to frustrate such preferences. We should understand *permissible* as *permissible when considering aspects of morality other than overdemandingness*. This prevents us from begging any questions about which preferences are rendered permissible by overdemandingness complaints, while capturing the thought that morally abhorrent preferences (such as racist preferences) cannot serve as legitimate complaints against obligations that would frustrate those preferences.¹²

Deep Preferences has three attractions. First, it respects the previous section's result: member overdemandingness is neither necessary nor sufficient for collective overdemandingness. A collective can hold a permissible deep preference that no member holds, and every member can hold a permissible deep preference that the collective doesn't hold. We saw this with the art gallery that prefers to display contemporary art while all members prefer the gallery to display noncontemporary art.

Second, deep preferences seem present in the introduction's examples: an art gallery's deep preference is to display art; a gas station's deep preference is to sell gas; a state's deep preference is to

¹²This "permissible" caveat follows Berkey (2019, 124n13).

retain a sense of community. Perhaps these preferences aren't permissible when considering aspects of morality other than demandingness. Arguably, the gallery's preference distracts from more morally important goals, the gas station's preference is positively harmful (even if it also benefits motorists), and the state's preference disrespects foreigners affected by its laws. But—the advocate of Deep Preferences might respond—that simply shows that these specific examples aren't legitimate complaints. Still, a collective's *permissible* deep preferences should be protected. In fact, Deep Preferences provides a good error theory of the introduction's examples: we are tempted to grant these overdemandingness complaints because collectives' deep preferences are at issue; we err in assuming the deep preferences are permissible when considering aspects of morality other than demandingness. So Deep Preferences does well at addressing these examples.

A third attraction is that deep preferences seem operative in the overdemandingness complaints of both individual and collective agents. This attractively unifies our theory of overdemandingness: our theory can apply to collectives just as it applies to individuals. An individual's deep preferences might include "being an excellent teacher and researcher, being a dedicated parent, maintaining lifelong friendships, or seeking new cultural experiences," while a collective's deep preferences might include "searching for oil" or "worshipping [and] proselytising."¹³ Perhaps all these preferences block moral obligations via overdemandingness (Collins and Lawford-Smith 2016, 52–53).

The third attraction of Deep Preferences is its Achilles' heel. By raising the prospect of a unified overdemandingness theory across agent types, it invites us to apply to collectives our justifications for protecting individuals' permissible deep preferences. If the justifications don't apply, this is bad news for collective overdemandingness complaints. To assess this, I will consider whether three prominent justifications for protecting individuals' permissible deep preferences apply to collectives. These proposals come from Garrett Cullity, Liam Murphy, and Samuel Scheffler and Bernard Williams. For argument's sake, I assume these justifications satisfactorily protect individuals' deep preferences.¹⁴ Still, they don't extend to collectives. Perhaps there is a *sui generis* justification for protecting collectives' deep preferences, but my discussion will leave that looking far-fetched.

2.b Reciprocal obligations

Garrett Cullity (2004) argues morality cannot infringe individuals' permissible deep preferences because we judge that individuals are sometimes obliged to help other individuals pursue similar preferences. For example, if I can help you contact a long-lost friend at little cost to myself, then I have a duty to do so. This intuition, Cullity argues, is on par with intuitions used to defend demanding moral conclusions, such as Singer's (1972) intuition that one must sacrifice one's shoes to save a drowning toddler, from which Singer argues that one must sacrifice a lot to ending poverty. According to Cullity, my obligation to help you contact your friend implies you are permitted to contact your friend. But if you're permitted to contact your friend, consistency requires that I'm permitted to pursue similar projects. In general, individuals are permitted to pursue whichever projects others would have obligations to help them pursue. This includes at least some permissible deep preferences.

The problem is this. We don't tend to judge that collectives are obliged to help one another pursue permissible deep preferences, or that individuals are obliged to help collectives pursue permissible deep preferences. That is, Cullity's envisaged reciprocal obligations of project-assistance (with concomitant permissions of project-pursuit) don't include collective project-pursuers or collective project-assisters.

Of course, some collectives have reciprocal legal obligations. For example, sometimes tech companies have obligations to share code and large firms have obligations not to compete with

¹³Collins and Lawford-Smith (2016, 52) give these examples the label *constitutive ends*, but the ends are not *constitutive* in the sense I develop in section 4. I therefore use the label *deep preferences*.

¹⁴For arguments against this, see Berkey 2016; Ridge 2010; Stroud 2013.

small businesses.¹⁵ But these are not reciprocal obligations to help one another pursue projects. Instead, these obligations are grounded in social utility considerations: by requiring code sharing or restricting competition, we bolster gross domestic product, or human happiness, or so on. And our intuitions about these obligations do not have the same “simply intuitive” status as Cullity’s friend-helping or Singer’s toddler-saving. Cullity’s argument relies on that status.

Likewise, sometimes collectives are obliged to help *individuals* pursue projects. Perhaps my employer is obliged to provide me parental leave to help me pursue my project of raising children. But, again, this can’t grant my employer permissions via Cullity’s strategy. The strategy applies only when the projects of the assister and assistee are similar. It’s about reciprocity: I’m permitted to pursue projects like those I’m obliged to help others pursue. But the kinds of projects collectives are obliged to help individuals pursue (parenthood, leisure, work–life balance) aren’t the kinds of projects collectives can pursue because collectives aren’t biological creatures. Reciprocal obligations don’t protect collectives’ permissible deep preferences.

2.c Fairness

Liam Murphy (2000) provides another route to protecting individuals’ permissible deep preferences: fairness.¹⁶ In brief: if everyone was willing to do what morality required, then a fair division of labour wouldn’t require individuals to sacrifice deep preferences. The deep preferences of compliers are infringed only because noncompliers refuse to do their fair share. Murphy’s strategy allows us to protect compliers’ permissible deep preferences if, and because, the demand to forgo them would treat compliers unfairly and disrespectfully relative to noncompliers.

Do collectives deserve fairness and equal respect? In the literature on collective intentionality, agency, and responsibility, almost all philosophers are “normative individualists”: although collectives might have intentionality, agency, and responsibility, only individuals are of intrinsic moral worth.¹⁷ Even those who endorse rights for collectives view those rights as a necessary means to some end, such as social utility or collectives’ accountability (Smith 2018). And even those who argue collectives deserve moral consideration do not require that we treat them fairly, or with equal respect, relative to individuals (Silver 2019).

That said, even without a requirement to treat collectives fairly and with equal respect relative to *individuals*, there might be a requirement to do so relative to *other collectives*. This is compelling when two collectives are (in some sense) the same type. For example, it might be problematically unfair, or unequally respectful, if funding bodies applied higher teaching or research standards to your university than my university. And—unlike collectives’ obligations to one another under Cullity’s proposal—this needn’t be understood as a means to gross domestic product or human happiness. Perhaps it’s intrinsically wrong to treat universities differentially in these ways.¹⁸ Could the need for fairness and equal respect protect collectives’ permissible deep preferences, shoring up their overdemandingness complaints?

I suggest not. Under Murphy’s proposal, individuals’ permissible deep preferences are protected conditional on the contingent, partly empirical fact that fair shares don’t infringe individuals’ permissible deep preferences. This is less plausible for collectives. Consider oil companies: treating one oil company fairly (and equally respectfully) relative to others plausibly means demanding that *all* oil companies redirect their operations. Likewise for states extending the franchise. The same even goes for art galleries: if it’s fairness *within* each collective-type that matters, then we need some independent story about why fairness *amongst art galleries* should be expected to preserve the

¹⁵I thank Hannah Tierney for this point.

¹⁶Murphy’s goal arguably wasn’t to protect individuals’ permissible deep preferences, but his argument provides a strategy for this. Murphy replaces overdemandingness with unfairness.

¹⁷For example, List and Pettit (2011), Hess (2013), Wringer (2014), Berkey (2019).

¹⁸I thank William Tuckwell and Alexander Dietz for pressing this.

display of art. Fairness and equal respect might spell the (fair and equal) end of all art galleries. Fairness and equal respect within each collective-type does not imply nondemanding obligations.

But perhaps morality demands fairness and equal respect across all collective-types. If so, some collectives' deep preferences might be protected. But, I suggest, morality does not demand fairness and equal respect across all collective-types: we typically think, for example, that states are more sacrosanct than corporations or art galleries. If there is a sense of fairness that demands fairness across all collectives, it's a sense of fairness that permits highly differentiated (and, sometimes, highly demanding) obligations across different collective-types. The upshot: to protect collectives' permissible deep preferences, we need something more robust than fairness or equal respect for collectives.

3.d Projects and points of view

A third justification provides such robustness. It says deep preferences are intrinsically valuable from the perspective of ground projects or the personal point of view. According to Bernard Williams (1981, 12–13), ground projects are “closely related to [one’s] existence,” “give meaning to ... life,” and provide “a reason for living.” According to Samuel Scheffler (1982, 56, 77), the personal point of view is the perspective from which “projects are undertaken,” “plans are developed,” “events are observed,” “life is lived,” and “a locus relative to which harms and benefits can be assessed.” Perhaps most importantly, the personal point of view is “a source for the generation and pursuit of personal commitments and concerns that is independent of the impersonal perspective” (1982, 57).¹⁹

For Williams and for Scheffler, legitimate overdemandingness complaints derive from the intrinsic value of ground projects or the personal point of view, both of which are partly constituted by deep preferences. For Scheffler, this value permits individuals to give *more weight* to their personal point of view than the impersonal point of view. For Williams, ground projects seem to provide a *hard constraint* against moral theories that require individuals to make decisions purely from the impersonal point of view. Either way, the weight or constraint are not conditional on contingent empirical facts about what needs to be done to make the world better, as under Murphy’s view. If collectives have ground projects and personal points of view, then perhaps Scheffler’s and Williams’s stories apply to them.

Before concluding this, we need to ask why ground projects or personal points of view matter. Several philosophers have argued they matter only insofar as their bearers are phenomenally conscious, which rules out collectives (Hess 2013; Pasternak 2017; Berkey 2019, 126). Yet Williams and Scheffler do not evoke consciousness. They do not emphasize that humans have negative qualia when their ground projects or personal points of view are stymied; they do not emphasise “what it is like” to have ground projects and points of view; they do not discuss feelings. Williams even distances project pursuit from feelings: ground projects are a cornerstone of his rejection of utilitarianism, which (in Williams’s view) focuses too much on producing good feelings (1963, 112–13; 1981, 8). Instead, ground projects and personal points of view are connected to *agency*: as reflected in the quotes above, ground projects and personal points of view provide reasons, meaning, projects, plans, assessments, commitments, and concerns.

Thus, we need not rely on consciousness to resist applying Williams and Scheffler to collectives. There’s a different reason for resistance: collectives’ ground projects and personal points of view are not produced in the right way. They are produced via a procedure that is thrust upon the collective from the outside when the collective is established. To be sure, the collective’s procedure is *maintained* internally via members-qua-members’ commitments, inputs, and actions. But internal maintenance can happen only once the procedure is established. The establishment of a collective’s

¹⁹I thank Daniel Muñoz for enhancing my understanding of Scheffler.

agency—and the establishment of the procedure by which its agency is exercised—is a process enacted upon the collective by individuals (or other collectives) who, at that establishment moment, should be conceived of as nonmembers (though perhaps soon-to-be members). A collective’s decision-making procedure is established by agents other than the collective itself.

Contrast individuals. When individuals are created, they are not instantly agents. Individuals’ agency develops through a complex interplay between the individual and her environment. The creators (parents) lack the power to intentionally impose particular procedures, so parents cannot shape individuals’ projects or points of view in the way collectives’ creators can. An individual may lack decision-making procedures altogether—certainly, her projects and point of view are not driven by decision-making procedures designed by others and imposed upon her from the outside when she is created.²⁰ This is reflected in Williams’s emphasise on *convictions* (1963, 116). Humans’ convictions are arrived at by no procedure. Contrastively, collectives have convictions only via a procedure that was established, and is maintained, by others.

So individuals have fuller agency over their ground projects and points of view than do collectives. Full agency seems to be at the core of Scheffler’s and Williams’s proposals. If fullness of agency is what matters, then we have weaker reason to value collectives’ ground projects and personal points of view (and the deep preferences that partly constitute those projects and points of view) than we have to value individuals’.

At this point, one might wonder whether collectives’ agency is so attenuated that they cannot bear moral obligations at all.²¹ If so, the answer to my title question is “nothing”—not, as I’ve said, “an awful lot.” To respond, we must consider the function of moral obligations. From the first-personal perspective (when one has an obligation), obligations paradigmatically function as fitting inputs into one’s decision-making about what to do or be. From the second-personal perspective (when one is owed an obligation from another), obligations paradigmatically function as demands or claims. From a third-personal perspective (when another owes an obligation to someone else), obligations paradigmatically function as advice or guidance to the obligation bearer. All three functions presuppose merely that the obligation bearer has the ability to make decisions that are sensitive to obligations.²² Collectives have that ability. So they are agents in a full enough sense to bear obligations—even if the ability was imbued in them by others such that they lack the complete fullness of agency at issue in the Williams–Scheffler view.

Could there be a fourth justification for protecting collectives’ deep preferences? The prospects are grim. We’ve seen that almost all philosophers deny collective consciousness (or its value), that collectives’ agency is made and maintained in a less than fully autonomous way, and that fairness and equal respect are uphill routes to nondemanding collective obligations. We’ve seen that our intuitions about collectives’ role in our moral world don’t endorse preference protection. I conclude that collectives’ permissible deep preferences are up for moral grabs.²³

4. Constitutional constraints

4.a The proposal

Yet collectives’ moral obligations are sometimes limited. These limitations provide a reinterpretation of our intuitions about art galleries, gas stations, and states.

²⁰Others heavily influence individuals’ deep preferences, but others cannot entirely impose the procedure by which individuals’ preferences arise, or the constitutional constraints we find in collectives (see section 4).

²¹I thank an anonymous reviewer for this.

²²I expand this in Collins (ms.).

²³There are other options. For example, Richard Chappell (2019) argues a moral theory is overdemanding if it requires too much “willpower.” If the rationale for preserving willpower doesn’t rely on consciousness, autonomy, fairness, or equal respect, then this could apply to collectives. But Chappell says willpower exertion contains a “distinctive phenomenology” (2019, 5)—tying willpower to consciousness. I suggest we’ll get similar results for other options.

To begin, notice that collectives have agency over only the part of the world regarding which their decision-making procedures can process inputs and make decisions. That part of the world may not include everything regarding which there are moral principles; a collective's decision-making procedure may thus be constrained. It might be constrained in two ways: by its *form* (e.g., voting, committees, discretion, bargains, norms, etc.) and by its *fundamental ends* (i.e., the ends the collective cannot but pursue, at least within a particular time period). I'll use "constitutional constraints" to refer to the aspects of a decision-making procedure's form and fundamental ends that impose limits on the collective's domain of agency within some time period. (This point about the narrow *domain* of [some] collectives' agency is separate from section 3's point about the *origins* of [all] collectives' agency.)

My suggestion is that constitutional constraints place limits on collectives' abilities. We therefore violate "obligation implies ability" when we ask a collective to do what it is constitutionally constrained from doing. Collectives' legitimate overdemandingness complaints transmute into complaints that the collective is *unable* to perform the obligation. But inability is distinct from overdemandingness. The result is that overdemandingness per se is never an excuse for a collective not to perform an obligation. The closest we get are complaints that obligations violate constitutional constraints. These are best understood as complaints that "obligation" implies "ability" and the collective lacks the relevant ability.²⁴ Call my suggestion *Constitutional Constraints*.

We can contrast constitutional constraints with deep preferences. Recall: a deep preference is long held, undergirds other preferences, and is viewed by the collective as important. But many deep preferences are overridable in the collective's practical reasoning. A collective's constitution, by contrast, is the framework within which preferences are formed, networked together, and selected as important. The constitution also contains some preferences, but these are stronger than deep preferences: the preferences in the constitution target *fundamental ends*, which are ends the collective *cannot but* pursue, at least within a particular timeframe.²⁵

For example, suppose some friends establish a basketball team for the local competition. This happens by one friend asking the others if they want to comprise a team, the others saying yes, and one doing the registration. This establishes the collective and, with it, the decision-making procedure's form and fundamental end. The procedure's *form* is conversation-based consensus. The procedure's *fundamental end* is playing basketball. This is the end the collective cannot but pursue, at least in the short term. This example demonstrates that a constitution doesn't necessarily have legal status, such as articles of incorporation for a corporation.

Of course, in the long term, the team could use conversation-based consensus to ditch basketball and start a knitting circle (though it would potentially "kill" itself in the process, with the same people establishing a new and distinct collective). And the team could use conversation-based consensus to change its decision-making procedure—consensually deciding to appoint a dictator, for example. But such changes take time. Morality cannot demand that a collective change its ways overnight. So the team cannot have an obligation to knit blankets—still less to do so via dictate—within any less time than it'd take to use its current procedure to move toward that. In small groups, this might seem negligible. But in large, complex, and multifaceted groups, such changes can take weeks, months, or even years—as when a state gradually shifts along the political spectrum, or a business gradually shifts into new markets.

²⁴There is an enormous literature arguing that "obligation implies ability" is true of individuals (e.g., Vranas 2007; Littlejohn 2012; Schwan 2018). I assume it's an open question whether it's true of collectives (Lawford-Smith [2012] argues it is), but I assume this is the most plausible way to understand collective overdemandingness complaints. I say "obligation implies ability" rather than "ought implies can" because the former is more plausible than the latter.

²⁵Are fundamental ends covered by Williams's ground projects or Scheffler's personal point of view? I don't think so. For the Williams–Scheffler view to be about overdemanding obligations rather than inability, agents must be *able* to override their ground projects or depart from their personal points of view. But collectives are *unable* to override their fundamental ends (at least within a given timeframe).

4.b Explaining the intuitions

How does Constitutional Constraints explain the intuitions that an art gallery need not give to poverty relief, a small gas station need not overhaul itself, and a state need not extend its franchise abroad? Like Deep Preferences, Constitutional Constraints suggests an error theory. Our intuitions derive from the fact that it looks like a constitutional constraint would be violated. In each case, this appearance is somewhat (but only somewhat) misleading. A Constitutional Constraint may well block the obligation from being performed in the short term. But the relevant timeframe is not short term: obligations to give to relieve poverty, sell sustainable energy, and extend the franchise plausibly operate over the long term. The relevant collectives probably cannot do these things within weeks or even months. But collectives often have diachronic abilities: they are unable to perform some action now, but they are able to take some action now that, if taken, would enable them to perform the action later (Jensen 2009).²⁶ If the action would still be valuable at that later time, then the obligation is not ruled out by a constitutional constraint. Our intuitions, then, result from a kind of myopia in which we assume that a *present* constitutional constraint cannot be changed or that short-term inability implies long-term inability.

Do some collectives have long-term constitutional constraints, which would block even obligations that call for long-term change? Long-term constitutional constraints are exceedingly rare. A long-term constraint via a procedure's *form* would imply that some action cannot arise from any possible future procedure that the current procedure could be used to adopt. Decision rules almost always contain rules of change, meaning the procedure's form is malleable and will almost never amount to a long-term constraint.²⁷ Likewise, most collectives are set up to be able to change their *ends*, meaning they lack long-term fundamental ends (except perhaps the end of being rational). That said, taken literally, constitutions under which a collective's purpose is "including but not limited to X" imply that the collective really cannot drop X as a purpose (though other purposes can be added).²⁸

This might seem too demanding on collectives. If long-term constitutional constraints are rare, and if many obligations concern long-term actions or outcomes, then collectives' legitimate complaints about obligations are also rare. A less demanding view might say an obligation violates a constitutional constraint only if the collective doesn't *require* or *expect* or *usually find* that kind of obligation as an input into its procedure, or if the obligation conflicts with its *current* procedures (rather than procedures it could come to have at a later time when the obligation would still have value), or if the obligation conflicts with its *current* fundamental ends. Collectives that face these weaker types of "constraints" do seem to be constituted as inept or unreliable regarding the relevant obligations.²⁹ Perhaps they are constitutionally incapable of responding to those obligations, even if those obligations make long-term demands.

However, consider two points. First, this permissive standard is not used with humans. Your racist colleague is not let off the hook purely because he doesn't currently require himself to take seriously his obligation not to be racist, or doesn't expect that he has obligations not to be racist, or doesn't usually find the obligation to be racist as a consideration in his decision-making. He's not let off if his current manner of making decisions only considers the interests of certain racialized groups, or if he is currently unable to let go of racism (as is plausibly the case for implicit biases: these cannot simply be willed away in the short term). These facts do not render him unable to fulfil his obligation not to be racist in the long term. Assuming it would still be valuable for him not to be racist later, once he has changed his racist dispositions, then he has an obligation to do so.

²⁶Pinkert (2014, 196–201) analyses such abilities for nonagential groups.

²⁷I thank both anonymous reviewers for pressing the rarity of long-term constitutional constraints.

²⁸I thank an anonymous reviewer for raising this wording.

²⁹Hindriks (2018) gives such an account of collectives' moral agency.

Second, consider our intuitions about wrongful collectives. If it is constitutionally possible for the mafia (understood as criminal) to end its own existence in the long term—the most demanding obligation of all—then we should demand that it do so in the long term. Whether this is constitutionally possible for the mafia is an empirical question not addressed here. I suggest that if we find ourselves resisting such an obligation, then this resistance has arisen because we tacitly believe one or both of two things: first, that the mafia cannot use its decision-making procedures to end its existence (even in the long term); second, that there is little point in asserting the mafia's obligation because it is unlikely to perform it. The first of these is an empirical issue on which my argument is neutral. The second can easily be debunked: that an entity is *unlikely* to perform its obligation does not entail that the entity is excused from the obligation or lacks the obligation. So I leave it an open question whether the mafia has this obligation depending on what it is able to do.

If the proposal still seems too demanding, this is likely because some collective obligations are unduly costly on members-qua-private-persons. This provides a supplement to the above error theory. The above error theory suggested that we have overdemandingness intuitions when collectives are unable to perform obligations in the short term, though we err in overlooking that they can perform them in the long term. Additionally, perhaps our intuitions are influenced by the thought that it's overdemanding on members-qua-private-persons if their collective performs the obligation. As emphasised above, such a thought is not about overdemandingness.

Constitutional Constraints can thereby explain the intuitions from the introduction. What's more, it shares the other two virtues I attributed to Deep Preferences. First, it respects the result that collectives' complaints do not rise and fall with the complaints of members-qua-private-persons or members-qua-humans. Constitutional Constraints respects this result because a collective might be constitutionally incapable of producing an outcome within a timeframe, even though its members could produce that outcome in that timeframe by acting *ultra vires*.

Second, Constitutional Constraints allows for a unified theory across collective and individual agents. Several accounts of individuals' overdemandingness appeal to psychological inability.³⁰ A collective's constitution can be viewed as its psychological enablers and constrainers: the constitution (combined with inputs) determines what ends up in the collective's rational point of view. The rational point of view is constituted by (functional equivalents of) mental states. Insofar as a collective's rational point of view is psychological, the constitution that enables and constrains that point of view generates psychological abilities and inabilities for the collective. So psychological ability is at issue for both individual and collective overdemandingness, if a psychological inability-based theory of individual overdemandingness is correct.

That said, the relevant psychological inability manifests differently across individual versus collective agents. First, individuals do not have the kind of fundamental ends some collectives do. Individuals might have the fundamental end of rationality or abiding by Kantian categorical imperatives (Korsgaard 1996), but these are far less domain-specific and task-oriented than "playing basketball." Second, individuals plausibly lack decision-making procedures like conversation-based consensus, so these won't produce individuals' constraints. But this doesn't mean we can ask more of individuals than of collectives; likely, individuals' psychological abilities are also restricted in ways collectives' are not because of individuals' embodiment and emotion. Psychological inability arises in both types of agent, but it manifests differently.

4.c Underdemanding?

Some might find it underdemanding that constitutional constraints remove some collective obligations. Forget about Mom-and-Pop Gas: Aren't I letting, say, ExxonMobil off the global-

³⁰Hooker 2000, 165–66; Lichtenberg 2004; Louise 2009; Swanton 2009. Does this mean individual overdemandingness complaints are properly reinterpreted as inability complaints? Yes, if these theories of individual overdemandingness are correct. But I left it open that one of section 3's accounts is true of individuals.

warming hook, at least in the short term? No. ExxonMobil's diachronic climate-related obligations are plausibly highly demanding. Those obligations will require action in the short term. To show that ExxonMobil can bear a diachronic obligation to move entirely into sustainable energy, we just need to show that its constitutional form and fundamental ends don't rule out this change in the long run. There are reasons to think such change is within ExxonMobil's abilities.

First, consider ExxonMobil's forms. In characterising collectives, I said that informal norms and culture are part of a collective's decision-making procedure. They are part of the constitutional form. These informal procedures will often be the initial mechanism by which changes occur: water-cooler chat, post-meeting complaining, informal reason-giving, and so on, amongst managers and shareholders. ExxonMobil may be capable of using these informal forms of decision-making to respecify how it enacts its more formal decision-making forms (for example, by revising them to include a wider range of stakeholders).

Second, what about ExxonMobil's fundamental ends? Isn't ExxonMobil's fundamental end "selling oil"? Not necessarily, at least not in the long term. "Supplying energy" may be a more accurate specification of ExxonMobil's longer-term fundamental end; "selling oil" may be merely a deep preference that is one way of specifying that fundamental end. This is suggested by ExxonMobil's recent forays into biofuels,³¹ and even wind and solar power.³² These forays are evidence that, all along, ExxonMobil faced no long-term constitutional constraint against holding sustainable energy as an end. The forays provide retrospective evidence of its always-latent diachronic abilities. If ExxonMobil's longer-term fundamental end is "supplying energy," not "drilling oil," then there's no reason to think it can't supply *only* sustainable energy.

What's more, even in the rare cases where constitutional constraints let collectives entirely off the hook, other nearby agents will likely have obligations instead. These nearby agents include members-qua-private-persons, different collectives composed of those individuals, individual nonmembers, or collective nonmembers. These agents may have obligations to alter ExxonMobil's constitution from the outside, to reign ExxonMobil in, or (at the extreme) to "kill" ExxonMobil.

The same goes for other collectives that are rotten to the core. Suppose (simply *arguendo*) the mafia is indefinitely constitutionally constrained from ceasing to engage in criminality. Suppose, however, that if members act outside their role, for example by ratting, sabotaging, or subverting, then they could bring the mafia down. The individuals are candidates for obligations to perform these actions. If there are several such individuals, they might have obligations to act together, and bear obligations as a collective with an entirely different constitution from that of the mafia itself. Individual nonmembers might have obligations to assist in take-down efforts, as possible. And collective nonmembers (state bodies or other organisations) might have obligations to do what *they* can. There are, under this scenario, many obligations related to ending rotten organisations. The obligations are simply not held *by* those organisations (again, assuming there's a genuine long-term constitutional constraint and assuming obligation implies ability).

5. Conclusion

When we take seriously the idea that collectives are distinct agents—made of individuals, but not identical to a collection of individuals—we get a new perspective on collective overdemandingness. First, an obligation's being overdemanding on one or more member is neither sufficient nor necessary for that obligation's being overdemanding on the collective. Second, we cannot appeal to a collective's permissible deep preferences to work out which obligations are overdemanding for it, because we lack a good story about why those preferences block obligations. This suggests that

³¹<https://www.exxonmobil.com.au/en-au/research-and-innovation/advanced-biofuels>

³²<https://www.forbes.com/sites/johnparnell/2018/11/28/renewable-energy-laggard-exxonmobil-to-power-oil-exploration-with-wind-and-solar>

collective overdemandingness complaints are a chimera. Instead, we should reconstrue such complaints as complaints that obligations violate their constitutional constraints, at least as those constraints are currently configured. This is ultimately an appeal to the principle that obligation implies ability. Some readers might want to deny that principle. If so, those readers' moral theories can ask anything they want of collective agents. And those who endorse the principle can still ask an awful lot of collective agents, especially in the long term.

Acknowledgments. For helpful feedback, I thank audiences at "Demandingness in Practice" at the University of Münster in 2018, "Collective Agency and Group Responsibility" at the University of Vienna in 2019, the Australasian Association of Philosophy conference at the University of Wollongong in 2019, "Feasibility and Poverty" at the Australian National University in 2019, and "Group Duties" at University College Dublin in 2019. For comments on written drafts, I thank Caroline Arruda, Jelena Belic, Matej Cibik, and two anonymous reviewers.

Funding Statement. Research for this article was funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No. 740922, ERC Advanced Grant "The Normative and Moral Foundations of Group Agency." I did further work on the article while receiving financial support under the Australian Research Council's DECRA scheme (project number DE200101413).

Stephanie Collins is an associate professor of philosophy at the Australian Catholic University. She is the author of *Group Duties: Their Existence and Their Implications for Individuals* (Oxford University Press, 2019) and *The Core of Care Ethics* (Palgrave Macmillan, 2015).

References

- Berkey, Brian. 2016. "The Demandingness of Morality: Toward a Reflective Equilibrium." *Philosophical Studies* 173: 3015–35.
- Berkey, Brian. 2019. "Collective Obligations and Demandingness Complaints." *Moral Philosophy and Politics* 6 (1): 113–32.
- Chappell, Richard. 2019. "Willpower Satisficing." *Noûs* 53 (2): 251–65.
- Chappell, Timothy, ed. 2009. *The Problem of Moral Demandingness*. Hampshire, England: Palgrave Macmillan.
- Chatterjee, Deen, ed. 2004. *The Ethics of Assistance*. Cambridge: Cambridge University Press.
- Collins, Stephanie. 2018. "Who Does Wrong When an Organisation Does Wrong?" In Hess, Ignieski, and Isaacs 2018, 197–220.
- Collins, Stephanie. 2019. *Group Duties*. Oxford: Oxford University Press.
- Collins, Stephanie. ms. "Able but Not Obligated: Lessons from Non-Agential Groups."
- Collins, Stephanie, and Holly Lawford-Smith. 2016. "Collectives' Duties and Individuals' Duties: A Parity Argument." *Canadian Journal of Philosophy* 46 (1): 38–58.
- Cullity, Garrett. 2004. *The Moral Demands of Affluence*. Oxford: Oxford University Press.
- Elder-Vass, David. 2007. "For Emergence: Refining Archer's Account of Social Structure." *Journal of the Theory of Social Behaviour* 37 (1): 25–44.
- Epstein, Brian. 2015. *The Ant Trap*. New York: Oxford University Press.
- French, Peter. 1984. *Collective and Corporate Responsibility*. New York: Columbia University Press.
- Goodin, Robert E. 2007. "Enfranchising All Affected Interests, and Its Alternatives." *Philosophy and Public Affairs* 35 (1): 40–68.
- Herman, Barbara. 2002. "The Scope of Moral Requirement." *Philosophy and Public Affairs* 30 (3): 227–56.
- Hess, Kendy. 2018. "The Peculiar Unity of Corporate Agents." In Hess, Ignieski, and Isaacs 2018, 35–60.
- Hess, Kendy. 2013. "If You Tickle Us . . . : How Corporations Can Be Moral Agents without Being Moral Persons." *Journal of Value Inquiry* 47 (3): 319–35.
- Hess, Kendy, Violetta Ignieski, and Tracy Isaacs, eds. 2018. *Collectivity*. Lanham, MD: Rowman & Littlefield.
- Hindriks, Frank. 2018. "Collective Agency: Moral and Amoral." *Dialectica* 72 (1): 3–23.
- Hooker, Brad. 2000. *Ideal Code, Real World*. Oxford: Oxford University Press.
- Jensen, Mark. 2009. "The Limits of Practical Possibility." *Journal of Political Philosophy* 17 (2): 168–84.
- Korsgaard, Christine. 1996. *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Lawford-Smith, Holly. 2012. "The Feasibility of Collectives' Actions." *Australasian Journal of Philosophy* 90 (3): 453–67.
- Lichtenberg, Judith. 2004. "Absence and the Unfond Heart: Why People Are Less Giving than They Might Be." In Chatterjee 2004, 75–100.
- List, Christian, and Philip Pettit. 2002. "Aggregating Sets of Judgments: An Impossibility Result." *Economics and Philosophy* 18: 89–110.
- List, Christian, and Kai Spiekermann. 2013. "Methodological Individualism and Holism in Political Science: A Reconciliation." *American Political Science Review* 107 (4): 629–43.
- List, Christian, and Philip Pettit. 2011. *Group Agency*. Oxford: Oxford University Press.
- Littlejohn, Clayton. 2012. "Does 'Ought' Still Imply 'Can'?" *Philosophia* 40 (4): 821–28.
- Louise, Jennie. 2009. "Moral Demands and Not Doing the Best One Can." In Timothy Chappell 2009, 185–200.

- McElwee, Brian. 2017. "Demandingness Objections in Ethics." *Philosophical Quarterly* 67 (266): 84–105.
- Miller, David. 2009. "Democracy's Domain." *Philosophy & Public Affairs* 37: 201–28.
- Murphy, Liam. 2000. *Moral Demands in Nonideal Theory*. New York: Oxford University Press.
- Pasternak, Avia. 2017. "From Corporate Moral Agency to Corporate Moral Rights." *Law and Ethics of Human Rights* 11 (1): 135–59.
- Pinkert, Felix. 2014. "What We Together Can (Be Required to) Do." *Midwest Studies in Philosophy* XXXVIII: 187–202.
- Pinkert, Felix. 2018. "When Are Collective Obligations Too Demanding?" In Hess, Ignesi, and Isaacs 2018, 175–96.
- Ridge, Michael. 2010. "Fairness and Non-Compliance." In *Partiality and Impartiality*, edited by Brian Feltham and John Cottingham, 194–222. Oxford: Oxford University Press.
- Rovane, Carol. 1998. *Bounds of Agency*. Princeton, NJ: Princeton University Press.
- Rovane, Carol. 2004. "What Is an Agent?" *Synthese* 140: 181–98.
- Rovane, Carol. 2014. "Group Agency and Individualism." *Erkenntnis* 79: 1663–84.
- Scheffler, Samuel. 1982. *The Rejection of Consequentialism*. Oxford: Oxford University Press.
- Schwan, Ben. 2018. "What Ability Can Do." *Philosophical Studies* 175 (3): 703–23.
- Silver, Kenneth. 2019. "Can a Corporation Be Worthy of Moral Consideration?" *Journal of Business Ethics* 159: 253–65.
- Singer, Peter. 1972. "Famine, Affluence, and Morality." *Philosophy & Public Affairs* 1 (3): 229–43.
- Singer, Peter. 2015. *The Most Good You Can Do*. New Haven, CT: Yale University Press.
- Smith, Leonie. 2018. "The Curious Case of Ronald McDonald's Claim to Rights: An Ontological Account of Differences in Group and Individual Person Rights." *Journal of Social Ontology* 4 (1): 1–28.
- Stroud, Sarah. 2013. "They Can't Take That Away from Me: Restricting the Reach of Morality's Demands." In *Oxford Studies in Normative Ethics*, vol. 3, edited by Mark Timmons, 203–34. Oxford: Oxford University Press.
- Swanton, Christine. 2009. "Virtue Ethics and the Problem of Demandingness." In Timothy Chappell 2009, 104–22.
- Thomas, Alan. 2005. "Reasonable Partiality and the Agent's Personal Point of View." *Ethical Theory and Moral Practice* 8 (1): 25–43.
- Thomas, Alan. 2009. "Consequentialism, Integrity and Demandingness." In Timothy Chappell 2009, 123–47.
- Vranas, Peter B. M. 2007. "I Ought, Therefore I Can." *Philosophical Studies* 136 (2): 167–216.
- Williams, Bernard. 1963. "A Critique of Utilitarianism." In *Utilitarianism: For and Against*, by J. J. C. Smart and Bernard Williams, 27–52. Cambridge: Cambridge University Press.
- Williams, Bernard. 1981. "Persons, Character, and Morality." In *Moral Luck: Philosophical Papers 1973–1980*, by Bernard Williams, 1–19. Cambridge: Cambridge University Press.
- Wisor, Scott. 2014. "Why Climate Change Divestment Will Not Work." *Ethics and International Affairs*. <https://www.ethicsandinternationalaffairs.org/2014/why-climate-change-divestment-will-not-work>.
- Wringe, Bill. 2014. "May I Treat a Collective as a Mere Means?" *American Philosophical Quarterly* 513 (3): 273–84.