

Structural Injustice and Massively Shared Obligations

ANNE SCHWENKENBECHER 

ABSTRACT *It is often argued that our obligations to address structural injustice are collective in character. But what exactly does it mean for ‘ordinary citizens’ to have collective obligations vis-à-vis large-scale injustice? In this article, I propose to pay closer attention to the different kinds of collective action needed in addressing some of these structural injustices and the extent to which these are available to large, unorganised groups of people. I argue that large, dispersed, and unorganised groups of people are often in a position to perform distributive collective actions. As such, ordinary citizens can have massively shared obligations to address structural injustice through distributive action, but, ultimately, such obligations are ‘collective’ only in a fairly weak sense.*

Introduction

According to Corwin Aragon and Alison Jaggar, ‘structural injustice is an emergent property of social practices’ where ‘typically unplanned and often unforeseen interaction of a variety of systemic factors’ generate ‘nested networks of constraints and opportunities’.¹ According to Iris Marion Young, structural injustice exists

when social processes put large groups of persons under systematic threat of domination or deprivation of the means to develop and exercise their capacities, at the same time that these processes enable others to dominate or to have a wide range of opportunities for developing and exercising capacities available to them.²

Structural injustice, then, is collectively caused in that it results from and is constituted by everyday actions of millions of people operating within unjust social institutions.³ Examples of such injustice include sweatshop labour and the precarious employment conditions in many developing-world garment factories,⁴ global poverty,⁵ or migrant domestic labour.⁶ We can easily see that environmental harms will regularly fit the above description, too. Many structural injustices are *global* moral challenges.

But instead of focusing on how structural injustice is *caused* by or constituted through collective behaviour, this article will look at our obligations to *address* such injustice and potentially remedy it. It is often argued that the responsibility for addressing structural injustice is in an important sense *shared* or *collective*.⁷ In this article, I discuss in what sense we can share moral obligations⁸ to address structural injustice. As such, the focus is not on retrospective ascriptions of responsibility and blameworthiness for, or complicity in, existing unjust social institutions and practices,

but on our prospective responsibility, that is, our moral obligations to change that status quo.⁹

Aragon and Jaggar take issue with the notion of collective responsibility for structural injustice: ‘It is not hard to argue that most people today are somehow connected with unjust global processes. But does this mean we are all responsible for all forms of global structural injustice?’¹⁰ While they focused specifically on Young’s work,¹¹ I think their observation applies to the debate on large-scale collective obligations more generally. It is fair to say that the oft-repeated view that structural injustice can only be remedied by collective action and that our obligations to do so are in some sense ‘collective’ is notoriously vague and in dire need of qualification. It is my aim in this article to provide such qualification.

In particular, existing notions of collective obligation fail to engage with the literature on collective action and epistemology. That is why, in this article, I propose to pay closer attention to the different kinds of collective action and knowledge needed in addressing some of these structural injustices and the extent to which these are available to large, unorganised groups of people. I argue that large and unorganised groups of people are often in a position to perform distributive collective actions, but that ultimately, our obligations to perform such actions will be ‘collective’ only in a fairly weak sense. As such, the aim of the article is to reboot and reorganise an important but lamentably muddled debate, but also, ultimately, to push for a novel distinction between weak and strong collective obligations and the significant conceptual and practical differences between them.

Different Types Of Collective Action Problems

For each of us, there are three distinct ways in which we can take action vis-à-vis these large-scale problems. We may

- (1) Take direct action where we can individually make a difference; or
- (2) Act as members of organised groups; or
- (3) Act as constituents of unorganised collectives.

This article focuses on the third type of action and the collective obligations that ‘ordinary citizens’ can have qua constituents of unorganised collectives (as opposed to group agents). As constituents of unorganised collectives, ordinary citizens can perform cooperative joint actions or distributive actions – more on these terms later – but such groups have no formal, enduring decision-making structure. Why do I focus on (3)? Few of us have the power and influence to individually make a direct difference to the kind of problems I am concentrating on here, as per (1). Regarding (2), clearly, group agents such as states, national and international organisations, and corporations and, therefore, their members have duties to address global challenges. Many more of us can play a role as members of organised groups such as universities, corporations, governments, and their agencies and other organisations. However, these agents regularly fail to act on their obligations. Meanwhile, ‘ordinary citizens’ can still effect change in the world, albeit only collectively. The third type of action is available to all of us, and it can potentially be very impactful.

I deliberately refrain from focusing on the ‘global affluent’ only – that is, those ordinary citizens who can easily afford to make monetary (or in-kind) contributions. Iris Marion Young argued that even those who are victims of injustice have some responsibility to fix the problems they are facing.¹² I will not take sides in this debate but instead simply refer to ‘ordinary citizens’ as including all capable moral agents, keeping in mind that obligations are not necessarily evenly distributed amongst them.

Most importantly, Young’s point that we ought to fight structural injustice by ‘joining together with others in collective action’¹³ needs qualifying. To start with, we need to distinguish different ways in which outcomes can be collectively brought about by two or more agents. *Genuine collective* action is not to be confused with the *aggregation* of actions. Let me explain each in turn: the latter describes outcomes produced by the *aggregation* of independent individual actions where these actions are *not intended* to produce that outcome, such as causing climate change through cumulative greenhouse gas emissions. This kind of problem is often referred to as a collective action problem in economics and politics. I am not using ‘collective action’ in this – very broad – sense here. In any case, my focus in this article is on what I call *genuine collective action*.

Genuine collective actions are either

- (a) cooperative (interdependent) *joint actions*,¹⁴ or
- (b) *distributive actions* resulting from cumulative individual contributions towards a shared goal,¹⁵ or
- (c) actions of *group agents* or incorporated groups, where the group’s agency is not reducible to the agency of its individual members or the group has an identity over and above that of its members.¹⁶

Since I am interested in obligations that individuals outside incorporated groups hold, I will leave (c) aside. The remainder of this section clarifies the difference between cooperative joint actions (a) and distributive actions (b). Let me stress that I am not suggesting that these two types of action exhaust the whole spectrum of genuinely collective action in unorganised collectives nor that they will always occur in the pure form described here. Rather, in their pure form they represent two opposite ends of a continuum. There will be hybrid forms, too. What I am aiming to show in the next section is that these different types of collective action have very different implications for potential collective obligations to address structural injustice. But beforehand, let me give a more detailed account of the differences between them.

Dancing tango, playing a duet, or lifting a table together are instances of what I call *cooperative joint action*. In contrast, people donating one million dollars to a charity for disaster relief, or neighbours keeping a local beach clean by taking turns in collecting rubbish, would count as *distributive actions*.

Cooperative joint action is highly interdependent collaboration between individuals, and many joint actions cannot be performed by one individual agent as a matter of principle, such as aforementioned duet playing or tango dancing.¹⁷ More importantly, however, individual contributory actions cannot be performed in isolation, even though they may be performed consecutively, as in coauthoring a book or an article, where coauthors take turns.¹⁸

In contrast, distributive collective action results from the individual actions of two or more agents that are intended as contributions towards a joint endeavour (or shared goal). Such actions can in principle – though often not literally – be performed by one individual agent. A group (or set) of agents' cumulative ability to produce distributive action merely requires that a sufficiently great number of individual agents have the ability to perform individual actions towards a shared goal and that they each have reason to adopt that goal.¹⁹

Naturally, it will more often be the case that a group (or set) of agents have the ability to perform a distributive collective action than a cooperative joint action, since conditions for joint ability are more demanding in the latter case.

For distributive collective action, the epistemic conditions are much weaker. First-order *shared* knowledge of a joint or shared plan²⁰ will often suffice: a proposition *P* is shared knowledge among agents *x*, *y*, and *z* if each of them knows *P* (or else it is shared knowledge among them *to a degree* depending on how many of them know *P*).

One might think that some knowledge of others' intentions (concerning the shared goal) is required for this action type: how can I intend to reduce my carbon footprint *with a view to the shared end of mitigating climate change*, for instance, if I do not know whether others also intend to reduce theirs? After all, for mitigating climate change it is necessary that a large number of agents reduce their emissions, not just me. I suggest that we need not know that *specific* others share our intention in order to form our own, but we merely need to have a belief that it is *likely* that *some* persons share our intentions (that is, have the same intentions as we do). Obviously, the more certainty we have about others' intentions to contribute, the stronger will be our reasons to also adopt such intentions.

For example, if I stop eating meat because current animal husbandry practices are bad for the environment, then this may form part of a distributive action. Similarly, people may be reducing their carbon footprint in order to contribute to mitigating climate change.²¹ Importantly, where distributive actions take place there is often some kind of shared, publicly available plan, which ascribes individual roles to secure the collective end. Public health campaigns may promote distributive action on important issues, for instance, increasing vaccination rates. Individual contributions to distributive action tend to be straightforward and simple.²²

Many collective actions, however, are significantly more complicated than that. They require precise levels of (often sustained) interaction between participants. People in unstructured groups can perform such cooperative joint actions to a point. Usually, organised (or structured) agents will be best at performing such tasks, and unorganised groups might transform into structured group agents when confronted with such challenges. The more intricate the level of organisation in such a group is, the more closely will it resemble a group agent. (But I am not concerned with that part of the spectrum of genuine collective action here.)

For cooperative action to succeed, agents must usually have some second-or-higher-order knowledge of a shared plan and others' intentions. This will often be the case where agents can directly communicate with one another, and these facts are 'out in the open', so to speak. The strongest form of interdependent knowledge of this kind is common knowledge,²³ but weaker forms of interdependent knowledge²⁴ may suffice, depending on the action in question. For many types of cooperative action, for instance, it may be sufficient that plans and intentions are *public* knowledge amongst

agents. By that I mean a type of interdependent higher-order knowledge that is less demanding than common knowledge: a proposition P is public knowledge among agents x , y , and z if it is true that (a) most of them know P , and (b) most of them know (a).²⁵

The two types of actions discussed – cooperative action and distributive action – are ideal types of collective action by unorganised groups. Some collective actions may fall in between these categories or display features of both. However, roughly, they represent two fundamentally different ways in which agents in unstructured groups can jointly produce outcomes (or perform actions). The important thing to note is that many instances of structural injustice will require a variety of agents and groups of agents to take various forms of remedial collective action.

Acting On Global Poverty – An Example

The aforementioned distinctions are important because they will ultimately help us better understand the kind of collective obligations we may have to address structural injustice. To illustrate this, let me have a quick look at some of the arguments made in relation to our (collective) responsibility to address one of the most pressing problems of structural injustice, global poverty. In my view, authors rarely give enough credit to the diversity of collective actions required to address this problem. As such, arguments for collective obligations to address poverty tend to be too general and, as a result, lack plausibility.

One observed tendency is to focus predominantly on distributive action solutions to global poverty, for instance, when it is discussed as a problem of individual donations towards charity.²⁶ However, I think it is misleading to approach the problem of global poverty primarily in this way. Both types of collective action are required in addressing this problem. In my view, the topography of obligations concerning complex, large-scale problems is complex and multifaceted, rather than neat and tidy.²⁷ For instance, a significant part of the problem of global poverty are unjust global trade agreements and unjust global financial institutions.²⁸ Taking action on both issues requires genuine cooperative action at a global level, ideally by institutional agents. Short-term poverty and disaster relief, in contrast, can be achieved by distributive collective action. Cumulative donations to charities (and other contributions to their work) may count as distributive actions in support of existing cooperative action.

Consequently, I also want to caution against approaches that portray ‘our’ duty to address global poverty mainly as a duty to perform a cooperative action. Just like the previously mentioned view, these approaches paint an unrealistically simplistic picture of the moral landscape. One of these overly reductive approaches is what I call the *joint-rescue-analogy view*. According to this view, just like two individuals can have obligations to jointly rescue a drowning child from a pond if it takes two people to do so, we can also have collective obligations to help those suffering from poverty globally. ‘We’ usually means ‘the global community’ or ‘the affluent’. Variations of this view have been defended by Garrett Cullity²⁹ and Bill Wringe, for instance.³⁰

In *The Moral Demands of Affluence*, Cullity writes that ‘we stand under a collective moral requirement of beneficence to help needy people through aid agencies’.³¹ He claims that the life-saving analogy for donating to poverty relief as first put forward by

Peter Singer³² should be understood as grounding a collective rather than an individual requirement ‘given the way that contributions to aid agencies are cooperatively pooled’.³³ While I agree with the gist of his argument, I believe that the thought experiment underpinning it is unsuitable. In arguing against Singer’s view of obligations to assist the global poor as direct individual obligations to save lives, Cullity provides his own variation of Singer’s original shallow-pond case – the two-person ‘rescue case’. He uses this rescue case to motivate the collective requirement:

A collective requirement of beneficence can exist when it is possible for us to help someone together, even if none of us could have helped on her own. . . . If someone is drowning in front of you and me, and can be rescued only by using a winch mechanism that requires two people to operate, then it is obvious that we are morally required to help him, even if neither of us could do so single-handedly.³⁴

A similar argument has been made by Bill Wringe, who defends the idea of collective obligations by appealing to small-scale collective action scenarios: two people must alert the owner of their building of growing damage to the roof and they can only do this together, because only one of them can type the email and only one of them knows the technical details to describe the damage accurately.³⁵ Wringe argues that they are collectively obligated to report the damage and extrapolates that the same kind of collective obligations arises with regard to global poverty.³⁶ His main argument is that members of collectives that are not agents – such as the two people in the building with the leaking roof, but also the ‘global collective’ – can organise themselves to become agentive.

Neither Wringe nor Cullity gives an account of what it means for a group to have collective ability, but both implicitly assume that the kind of collective ability (and therefore also the kind of collective action and obligation) involved in the small-scale scenario must be the same or at least similar enough to the global case.³⁷ This is surprising, because a number of factors influencing the group’s ability to organize itself and to jointly act that are present in the small-scale scenario are lacking in the global case.

These factors usually include the level of epistemic cohesion in the group – the members’ ability to communicate with a view to developing a shared plan towards some desired end and allocating individual contributory actions, as well as the level of interdependent group knowledge. Such group knowledge is often considered to be the basis for shared or collective intentions – plural intentions that according to many scholars are characteristic of (cooperative) joint activity.^{38,39} In any case, more sophisticated forms of interdependent knowledge will usually increase people’s ability to act collectively.

Large-scale collective action scenarios will usually lack the epistemic simplicity that characterises small-scale cooperation scenarios invoked in the literature, where the solution to a problem is obvious to the individual present. Further, in rescue cases and similar scenarios used in the literature to motivate the idea of collective obligations, the agents’ willingness and ability to assist is usually instantly evident – people can communicate effortlessly and reassure each other in the process of performing an action or producing an outcome. Also, there is a certain moral simplicity to these small-scale cases: they tend to involve one-off efforts, which pose a negligible burden

on contributors while there exist no obvious competing duties. In short, I believe it is misleading to model large-scale (or even global) collective action and obligations on cases of small-scale one-off cooperation. The type of collective action that is available at large scale to agents in unorganised groups differs starkly from that of small-scale scenarios, because the nature of such groups and their collective ability will be very different.

Large-Scale Distributive Actions

But we need not approach the question of collective ability of unorganised groups of individuals (passers-by, commuters, the global affluent, ordinary citizens, etc.) as one of ability for cooperative action in the narrow sense described above. Many of those who defend collective obligations assume very different – much weaker – criteria for collective ability and also a weaker concept of collective action.

The notion of collective action underlying Felix Pinkert's view of joint duties is more closely aligned with what I have called distributive action. Just like Wringer and Cullity, Pinkert thinks that unorganised groups of agents can hold moral duties, provided they have joint ability. Pinkert distinguishes between mediate and immediate joint ability to do x .⁴⁰

Immediate joint ability to do x exists where there is one *salient* possible collective pattern of actions that would constitute x -ing, and every agent in the relevant group believes what is in fact her part to be her part.⁴¹ In cases where there is no such salient pattern and the individual parts are not immediately known to the members of the groups, or else, they do not (yet) have true beliefs concerning their parts in x , they could still hold joint duties if they have *mediated* joint ability – that is, the capacity to establish immediate joint ability: 'agents can already be jointly able to perform actions which first require them to coordinate'.⁴²

In other words, if there is some preliminary step to doing x , which establishes their immediate joint ability to do x , which is obvious to them and which they have the ability to take, then a group of agents have mediated joint ability to do x . This means that they may have a joint duty to do x even if there is no salient pattern of action to start with. On Pinkert's account, then, immediate joint ability (and a joint duties) can be present even if there is no communication between the different members of a group, let alone something akin to interdependent, interlocking intentions.⁴³

Pinkert's criteria will often be met where simple, distributive actions are concerned. Interestingly, though, he does not think that the global affluent (or humanity as a whole) are under a collective obligation to remedy global poverty, for example. This is because he believes that it would take us too much time to establish immediate joint ability to actually save people from dying. That is, we would take too long to coordinate in order to save all the lives that could be saved. I believe that this conclusion, too, falls prey to an overly simplistic view of our obligations vis-à-vis large-scale injustice: global poverty is a complex problem, and as such the obligations to remedy it will be diverse and complex, held by different kinds of agents. It seems misguided to assume that there could be only *one* all-encompassing collective obligation and one set of conditions triggering that obligation.

While the root causes of global poverty are best addressed by organised group agents, the unstructured global collective of ‘ordinary’ citizens can, however, play a supporting role and take up some of the slack left by those agents. Contra Pinkert, I argue that we can have multiple collective obligations to address large-scale (even global) injustice. With regard to global poverty, there do exist salient patterns of collective action, and there is some shared and public knowledge as to how we can individually contribute. One way to discharge our obligations is via financial or other contributions to existing organisations which fight the causes and mitigate the impacts of poverty.⁴⁴ Through their activities and information campaigns, these organisations make certain patterns of group action and individual contributory actions supporting their work salient.

In practical terms, this means that those wanting to foster collective action must ensure that information is spread in the group in the right way. To the extent that information on either strategy – fighting the root causes of global poverty or mitigating its impacts – is readily available, ordinary citizens are under a duty to jointly support existing efforts to fight poverty and to therewith perform collective actions in the distributive sense. Donations, for example, can be seen as cumulative contributions towards this kind of collective goal.

As mentioned previously, one particularly important factor for a group of agents’ collective ability is the knowledge or beliefs that group members share. Even though I do not agree with Pinkert’s ultimate conclusion regarding our obligations vis-à-vis global poverty, he – correctly, in my view, – highlights the significance of epistemic conditions for joint ability. Pinkert emphasizes that it must be ‘obvious’ to the agents how to play their part or that a collective solution must be ‘salient’ for them to have immediate joint ability, but he does not provide much detail on these epistemic requirements.⁴⁵

Scott Shapiro’s account of massively shared agency – collective agency of very large, loose groups – is more concrete: He emphasizes that ‘shared intentional activity is activity guided by a shared *plan*’.⁴⁶ Shapiro’s notion of a plan in some sense mirrors Pinkert’s concept of a ‘pattern’ and corresponding ‘roles’: a ‘plan’ contains a specific end as well as specifying how to achieve that end, that is, which actions individual group members are to take in order to achieve that end. So, what does it mean for a plan to be *shared*?

A plan is shared by a group to J when (1) the plan was designed, at least in part, for the members of the group so that they may engage in the joint activity J and (2) each member accepts the plan.⁴⁷

Accepting a plan does not entail that each member know the full content of the plan, but the content is publicly accessible for those who wish to find out. Further, according to Shapiro, shared activity requires that each member intentionally follows their part of the plan and that this fact – as well as the fact that there is a shared plan – is common knowledge.⁴⁸

I believe that Shapiro’s proposal is best understood as an account of distributive action.⁴⁹ But, depending on what exactly Shapiro means by ‘common knowledge’ (and he does not specify in the text), his epistemic condition may well be too strong to capture the kind of massively shared agency that he is after and which we are focused on here. According to the so-called iterative definition, ‘[a] proposition *P* is

common knowledge in this sense if and only if everyone in the group knows that *P*, everyone knows that everyone knows that *P*, and so on, for any iteration of “everyone knows that”.²⁵ Clearly, if we adopt the iterative definition of common knowledge, we have to accept that the global affluent, or members of loose and dispersed groups more generally, cannot engage in shared activity of this kind. Group members may be justified in having first- or even second-order (*de dicto*) beliefs about others’ knowledge of the shared plan and their respective intentions. But if there is no direct communication between group members, the iteration will stop somewhere.

Shapiro’s account is more plausible with a weaker interpretation of the common knowledge condition. Perhaps all we need in order to jointly follow a shared plan is *public knowledge* of the shared plan. Public knowledge obtains where most people (or a large enough subset of those who can contribute to the shared end) know the plan and most people (or a large enough subset) are aware of that fact.

Yet even with this modification, Shapiro’s account is more demanding than Pinkert’s in terms of what beliefs and attitudes group members need to share (or have in common) in order to jointly act. Rather than arguing that one account is preferable to the other, I suggest that different types of collective action problems will require different types of joint ability and differing levels of shared beliefs among group members. Further, as I will indicate below, stronger normative (including epistemic) connections between group members have implications for the obligations they can have in common.

In sum, I have discussed three issues so far: (1) the different *kinds* of actions that are required in the fight against structural injustice, focusing mainly on the problem of global poverty, (2) what it means for large, unorganised groups of people to have the ability to perform a distributive collective action, and, in particular, (3) what group members need to know in order to be able to act. The upshot was that in order for such unstructured groups to have the required collective ability, there needs to be at least one collective pattern of action, which, if realized, will fix the problem in question, and this pattern is salient (or obvious) to group members. This includes that they have accurate beliefs with regard to the role that they need to play in order for the collective action to succeed. This pattern can take the shape of a ‘publicly known shared plan’. The public nature of the plan will make it the case that group members have some reason to believe (*de dicto*) that there are indeed other agents sharing that plan and collective end. Importantly, knowledge has to be spread in the group in the right way in order for people to have the collective ability to address a joint necessity problem. I will come back to this issue in the next section, when discussing the relevance of normative links between group members.

However, so far I have said little about what it means to share obligations, or to have collective obligation, as such. This will be discussed in the next section.

Massively Shared Collective Obligations?

To start with, let me briefly locate our discussion within the larger context of the debate on collective obligations. With regard to concrete, clearly circumscribed distributive action problems, can large, dispersed groups such as the global affluent have ‘massively shared’ collective obligations? Or, in other words, can we – ordinary citizens

– *share* obligations to address certain types of problems that are part of or result from problems of structural injustice? I have already shown that large and dispersed collectives can meet one necessary criterion for such obligations: cumulative ability, provided there exists a publicly known shared plan or an otherwise salient pattern of action.⁵⁰

There are four different ways in which obligations have been understood as ‘collective’ in the literature. Before I describe these, I will briefly mention the problem of *joint necessity*, which motivates many accounts of collective obligations.

This is a feature of actions (and outcomes) that cannot be performed (or produced) by one person on their own, but require at least two people in order to be realized. Joint necessity is *analytic* where it is part of what it means to do x that x is done by at least two people, as in ‘getting married’. It is *circumstantial* where, as a matter of fact (but not as a matter of principle), an action (or outcome) cannot be performed (or produced) by one person alone, for instance, if it takes two or more people to lift a heavy object.⁵¹

Where it takes two or more agents to address some morally significant problem, as in Cullity’s rescue case above, it seems that the capacity principle, or ‘*ought*’ implies ‘*can*’, is not met. If neither agent can resolve the problem, it appears as though neither has a duty to do so, which is counterintuitive. Hence the question is how we can ascribe obligations to entities that are not agents in their own right or else how to ascribe obligations to individuals that they cannot discharge on their own. The options discussed in the literature are as follows:

- (4) Suggesting that the two agents in the collective rescue case⁵² have an obligation as a group. Two readings are possible:
 - (a) **As obligations of group agents.** The group consisting of the two agents holds a moral obligation to rescue the person. However, this reading is implausible because it is ontologically overcommitted. At the moment of encountering the drowning person, the two random passers-by are not plausibly described as a group agent. Further, under none of the standard accounts of moral group agency do they become such an agent in the process of collaborating either. Further, even if they did transform into a group agent proper, the potential obligation to do so would encounter the same difficulty described above: forming a group agent is itself a joint-necessity problem.
 - (b) **As obligations of groups that are not agents.** The two agents are not a group agent in the strict sense, but they can hold an obligation as a group just the same. This reply violates the so-called agency condition that only moral agents can hold obligations. Some have suggested that the agency condition is implausible.⁵³
- (5) **As joint obligations of two or more individual agents.** ‘Being jointly obligated’ is considered a plural moral predicate⁵⁴ which attaches to sets of two or more individual agents.⁵⁵ This requires a plural reading of the capacity principle where [a and b jointly ought to do x] implies [a and b jointly can do x].⁵⁶
- (6) **As obligations that agents hold individually, but the content of which references a collective end.** Each of the two agents has obligations to contribute to collective action or to establishing such action.⁵⁷ The problem of this

interpretation is that, ultimately, no agent or group of agents has a duty to perform the joint action or to produce the collective outcome. Some scholars are happy to bite this bullet.⁵⁸ This is a distributive reading of collective obligations.

It would take up too much space to discuss the above readings and their differences in detail.⁵⁹ But let me suggest that they need not be mutually exclusive, and that, depending on circumstances, the nature of collective obligations may vary. All of the above will yield the conclusion that individuals in unstructured collectives have *some* obligations to contribute to improving large-scale collective action problems.⁶⁰

Examples of the kind of distributive actions I am interested in include donations towards charities that aim at sustained political and economic change, individual behavioural change where our standard practices cause collective harm to those who are already disadvantaged (as in reducing our carbon footprint or purchasing certain consumer goods), and similar types of distributive actions. But in what sense would such obligations still be ‘collective’?

Normative primacy of the collective level: As mentioned above, one reason for construing obligations as collective is to account for the fact that, with regard to problems of structural injustice (or, more generally, large-scale collective action problems), no individual’s actions taken in isolation will fix the problem (joint necessity). Rather, it is the collective ability of a set of agents to produce an outcome or perform an action, which is seen as grounding such obligations. Further, in many cases, an individual’s omission or inaction may not even make a difference for the worse, namely where there are more potential contributors to solving a problem than minimally necessary, that is, when we are then faced with a *wide joint necessity* case.⁶¹ For instance, wide joint necessity obtains if in the rescue case above there are five bystanders but only two are required in order to save the child. If our obligations depended on our unique ability to make a difference to an outcome, then none of the bystanders in this case would have an obligation to intervene. After all, if they fail to act there are still others who could help the drowning person. This dilemma is avoided if, as several scholars proposed, the collective level is considered primary and the question of individual obligations as secondary.⁶² Note that this move is not available to the third interpretation of collective duties mentioned above.

Framing the idea of normative primacy of the collective level in terms of moral deliberation, one might say that a collectively optimal pattern of action can give individuals group-based (or pattern-based) reasons for action even where their individual actions make no direct causal contribution to the realization of that pattern, for instance when the outcome is overdetermined.⁶³ Alternatively, one might say that an obligation is collective when individual agents have reason to deliberate from the point of view of the group.⁶⁴ In other words, the normative primacy of the collective level could be interpreted as metaphysical, functional, or deliberative. I suggest calling obligations of this kind ‘weakly’ collective obligations.

Normative links between agents: Another way to understand the *collective* character of moral obligations is to think of them as generating special normative links between members of said collectives. I suggest to call this a ‘strongly collective’ understanding of such obligations. For a set of agents to collectively have an obligation may require them to take some responsibility for the success of the collective action, such as coordinating the joint activity, and generating the kind of group knowledge required for

the group members to be able to fulfil their contributory actions. Further, it may mean that group members have to take up the slack left by others, essentially requiring them to contribute beyond what their ideal fair share would have been. Finally, group members may also be linked by considerations of fairness in the distribution of contributory roles.⁶⁵ All this, however, presupposes that certain epistemic conditions hold within the group, for example, that information is shared in the right way.

Our obligations to perform the kind of distributive actions that I have been mainly focusing on would be weakly collective in that the collective level has normative primacy. However, normative links between group members will bolster the collective character of obligations. This second aspect of (some) obligations' *collective* character may be crucial in putting both rational and psychological pressure onto agents to actually discharge those duties. Such links will strengthen moral reasons and moral resolve, potentially turning pro tanto into all-out duties and making certain collective actions not just epistemically salient but normatively salient.

Such normative links can strengthen moral reasons in favour of collective options by making collective patterns of action salient over individual patterns of action. When facing collective action problems, individual deliberating agents are regularly torn between individually efficacious actions and collectively available options. In other words, there is a tension between choices that make a direct difference and those whose success depends on how others choose.⁶⁶ Adequate information flow within groups concerning other agents' choices can make specific group patterns salient as well as lowering the risks involved with picking collectively available options generally. Group members who can directly communicate can jointly work out strategies for optimal action via team reasoning.⁶⁷ In epistemically more tightly knit (and yet unstructured) groups, compliant members are more likely to know about others' defections, giving them an opportunity to pick up the slack left by the defectors in order to secure the group goal. Arguably, under certain conditions, group members, even in unstructured groups, have obligations to take on more than their ideal fair share of costs when discharging obligations.⁶⁸

The greater a group's epistemic cohesion the more likely is there to be a robust sense of group membership and identity. We have learned from research in social psychology that a sense of group membership will enhance cooperative behaviour.⁶⁹ It is therefore desirable, from the point of view of combating structural injustice via large-scale collective action, to distribute information in the right way so as to generate a stronger sense of group identity and mutual accountability.

Ultimately, most moral agents will hold a variety of different – individual and collective – obligations to contribute to addressing structural injustice. Some of these obligations will focus more narrowly on cooperative actions, some on distributive actions, and many on actions that will fall in between these two types. Further, I think Young had it right when she suggested that the degree of individual agents' obligations (or responsibility) is determined by a number of factors including the power they have to make a difference, the degree to which they benefit from the injustice ('privilege'), the degree to which they are impacted by it ('interest'), and finally the ability for organising collective action, 'the relative ease with which people can organize collective action to address an injustice can be a useful decision principle'.⁷⁰ I would add, though, that our obligations also depend on the type of problem we are attempting to address, the kind of action that is required to remedy it, and the kind of contribution we can make

to it. Many of us will not be in a position to initiate collective action, but merely to contribute to existing endeavours. This, too, is a way of discharging our collective obligations. In sum, I believe that our moral obligations vis-à-vis complex moral problems are best understood as a polycentric cluster of collective and individual obligations. My focus, however, has been on a particular subtype of our obligations to address structural injustice, namely those that are most easily had – obligations to contribute to large-scale distributive action, or *massively shared obligations*.

Conclusion

Our starting point was the oft-made claim that structural injustice can only be remedied by collective action and the view that the responsibility or obligation to do so is shared among many people. I suggested that this view is too general and is in need of qualification. This is in part because it fails to take into account the differences between the kinds of actions we can perform together with others and the types of group-based knowledge required for successful collective endeavours.

I then distinguished between two generic types of collective action: cooperative action and distributive action, whereby the former is a more sophisticated, interdependent type of joint action, while the latter is collective action in a looser sense, with individuals intentionally contributing to a shared goal. Distributive action, I argued, is the type of large-scale collective action most readily available to ‘ordinary people’. I showed that large, dispersed collectives (such as the global affluent for instance) do have joint ability with regard to such actions, provided certain epistemic conditions are met. I argued that, depending on the problem at hand, people in unorganised collectives need first-order (and possible higher) knowledge of a shared plan.

Finally, massively shared obligations, that is, ordinary citizens’ obligations to address large-scale moral problems through distributive action, are collective only in a very weak sense. Against that backdrop, it is easy to see that strengthening normative and epistemic links between agents in unorganised groups would change the nature of their obligations and would, in fact, increase normative pressure to contribute to collective action. Practically, this means that fostering collective action may require strengthening the sense of group identity and membership, spreading information in ways that enough group members have access to publicly available shared plans that make salient the required individual contributory actions and providing feedback to the group on members’ actions and contributions. Research in social psychology demonstrates that it is a sense of group membership combined with information on positive, that is, compliant, behaviour of others that increases people’s willingness to contribute to collective causes.⁷¹ Strongly collective obligations, where agents have some responsibility for the success of the group action and for others’ actions and levels of knowledge are more desirable from the point of view of addressing structural injustice, but harder to generate. Importantly, strengthening those normative links is an action that may not readily be available to agents in large and dispersed unorganised groups.

Anne Schwenkenbecher, College of Arts, Business, Law and Social Sciences, Murdoch University, 90 South St, Murdoch, WA, 6150, Australia.
 a.schwenkenbecher@murdoch.edu.au

Acknowledgments

This article went through many incarnations, and numerous colleagues have helped me to sharpen my argument. I am grateful for comments on earlier versions of this article from Ned Dobos, Toni Erskine, Christina Friedlander, Alberto Giubilini, Niels De Haan, John Kleinig, Holly Lawford-Smith, David Lumsden, Alejandra Mancilla, Andrew Moore, Julian Savulescu, Leonie Smith, Joe Ulatowski, and Suzanne Uniacke. Further, I enjoyed instructive discussions with audiences at the Australasian Association of Philosophy Conference, Oxford University's St Cross Ethics Seminar, University of Otago, University of Waikato, Massey University, the Centre for Applied Philosophy and Public Ethics, University of Warwick, Mancept Political Theory Workshop, and Melbourne University. Finally, I would like to thank the School of Arts at Murdoch University and the Oxford Martin School at Oxford University for funding support.

NOTES

- 1 Corwin Aragon and Alison M. Jaggar, 'Agency, complicity, and the responsibility to resist structural injustice', *Journal of Social Philosophy*, 49, 3 (2018): 442.
- 2 Iris Marion Young, *Responsibility for Justice* (New York: Oxford University Press, 2011), p. 52.
- 3 Elizabeth Ashford, 'Severe poverty as a systemic human rights violation' in Gillain Brock (ed.) *Cosmopolitanism Versus Non-Cosmopolitanism: Critiques, Defenses, Reconceptualizations* (Oxford: Oxford University Press, 2013), pp. 129–155; Young 2011 op. cit.
- 4 Iris Marion Young, 'Responsibility and global labor justice', *Journal of Political Philosophy*, 12, 4 (2004): 365–388. <https://doi.org/10.1111/j.1467-9760.2004.00205>; Young 2011 op. cit.
- 5 Ashford op. cit.
- 6 Aragon and Jaggar op. cit.
- 7 Todd Calder, 'Shared responsibility, global structural injustice, and restitution', *Social Theory and Practice*, 36, 2 (2010): 263–290; Anne Schwenkenbecher, 'Joint duties and global moral obligations', *Ratio*, 26, 3 (2013b): 310–328. <https://doi.org/10.1111/rati.12010>; Bill Wringe, 'Needs, rights, and collective obligations', *Royal Institute of Philosophy Supplement*, 80, 57 (2005): 187–207; Bill Wringe, 'Global obligations and the agency objection', *Ratio*, 23, 2 (2010): 217–231; Bill Wringe, 'From global collective obligations to institutional obligations', *Midwest Studies In Philosophy*, 38, 1 (2014): 171–186. <https://doi.org/10.1111/misp.12022>; Iris Marion Young, 'Responsibility and global justice: A social connection model', *Social Philosophy and Policy*, 23, 1 (2006): 102–130; Ashford op. cit.; Young 2004 op. cit.; Young 2011 op. cit.
- 8 I am using 'responsibility' (in the forward-looking sense) and 'moral obligation' synonymously.
- 9 Of course, retrospective responsibility is a *source* of prospective responsibility or obligations. However, my focus is on the concept and possibility of collective obligations as such.
- 10 Aragon and Jaggar op. cit., p. 446.
- 11 In particular, they "worry that the breadth and vagueness of Young's social connection model dilutes the moral urgency for any particular individual to do anything specific" (Aragon and Jagger op. cit., p. 466).
- 12 Young 2011 op. cit.
- 13 Young 2006 op. cit., p. 123.

- 14 Michael E. Bratman, *Shared Agency: A Planning Theory of Acting Together* (New York: Oxford University Press, 2014); Philip Pettit & David Schweikard, 'Joint actions and group agents', *Philosophy of the Social Sciences*, 36, 1 (2006): 18–39.
- 15 Kirk Ludwig, *From Individual to Plural Agency: Collective Action I* (Oxford: Oxford University Press, 2016). To be clear, Ludwig's account of collective agency is inclusive of cooperative joint actions as portrayed by Bratman (op. cit.) and Pettit & Schweikard (op. cit.). However, his examples more closely resemble what I call distributive actions.
- 16 Toni Erskine, 'Assigning responsibilities to institutional moral agents: The case of states and quasi-states', *Ethics and International Affairs*, 15, 2 (2001): 67–85; Christian List & Philip Pettit, *Group Agency: The Possibility, Design, and Status of Corporate Agents* (Oxford; New York: Oxford University Press, 2011); Deborah Tollefsen, *Groups as Agents* (Cambridge, UK: Polity Press, 2015).
- 17 I focus only on joint-necessity cases – actions (or outcomes) that can only be successfully performed (or produced) by two or more agents. Of course, beyond those cases, there are many actions we can do either on our own or with other people (like going for a walk). The term 'joint necessity' was coined by Holly Lawford-Smith in 'The feasibility of collectives' actions', *Australasian Journal of Philosophy*, 90, 3 (2012): 453–467.
- 18 My notion of cooperative joint action is modelled on (or at least inspired by) accounts of shared or joint agency as proposed by Michael Bratman (op. cit.) or Philip Pettit and David Schweikard (op. cit.), although it is weaker than the latter and more inclusive than the former. The notion of distributive action I employ is inspired by but not identical to Kirk Ludwig's somewhat more minimal account of collective action (Ludwig op. cit.).
- 19 This last condition excludes cases where agents are ignorant of either the shared goal or other facts that might lead them to adopt that goal. For instance, a group of agents who have never heard about climate change do not have the ability to perform some distributive collective action aimed at reducing global warming. They can, of course, still take actions to that effect. But these would be aggregate actions, not genuinely collective actions.
- 20 A 'plan' specifies the steps necessary for achieving an 'end' or 'goal'.
- 21 Thomas Dietz, Gerald T. Gardner, Jonathan Gilligan, Paul C. Stern & Michael P. Vandenberg, 'Household actions can provide a behavioral wedge to rapidly reduce US carbon emissions', *Proceedings of the National Academy of Sciences*, 106, 44 (2009): 18452–18456. <https://doi.org/10.1073/pnas.0908738106>.
- 22 With increasing complexity of contributors' interactions, the collective action will become more 'cooperative' rather than 'distributive' (except in cases with a correspondingly intricate shared plan – here interaction can be complex yet distributive).
- 23 Common knowledge itself comes in different versions: the strongest is the iterative definition of common knowledge, whereas Lewisian common knowledge is weaker (e.g. see Oliver Roy & Anne Schwenkenbecher, 'Shared intentions, loose groups, and pooled knowledge', *Synthese* (2019): 1–19).
- 24 By that I mean second- or higher-order knowledge. First-order knowledge of p in a group has group members know p . Second-order knowledge has group members also know that the other member(s) know(s) p , and so on.
- 25 See e.g. Oliver Roy and Anne Schwenkenbecher op. cit.
- 26 Garrett Cullity, *The Moral Demands of Affluence* (Oxford: Clarendon Press, 2004); Liam B. Murphy, *Moral Demands in Nonideal Theory* (New York: Oxford University Press, 2000); Peter Singer, *The Life You Can Save: Acting Now to End World Poverty* (Melbourne, Vic: Text Publishing, 2009).
- 27 In this regard, I believe my approach complements Young's rather than merely critiquing it.
- 28 Thomas Pogge, 'World poverty and human rights', *Ethics & International Affairs*, 19, 1 (2005): 1–7. <https://doi.org/10.1111/j.1747-7093.2005.tb00484.x>
- 29 One might wonder why I object to Cullity's work both for focusing on distributive solutions *and* because it portrays our collective obligations as requiring cooperative action. My point is that in *The Moral Demands of Affluence*, Cullity mainly focuses on distributive action (donations) but at the same time defends his 'collective requirement' to engage in distributive action by appealing to a cooperative joint-action scenario.
- 30 A version of this view seems to also be present in Andrew Schroeder, 'Imperfect duties, group obligations, and beneficence', *Journal of Moral Philosophy*, 11, 5 (2014): 557–584. <https://doi.org/10.1163/17455243-4681020>
- 31 Cullity is talking about beneficence here, and not the (potentially stronger) duty to remedy an injustice. I do not think that it matters for my argument here (Cullity op. cit., p. 65).

- 32 Peter Singer, 'Famine, affluence, and morality', *Philosophy and Public Affairs*, 1, 3 (1972): 229–243.
- 33 Cullity op. cit., p. 65.
- 34 Cullity op. cit., p. 61.
- 35 Bill Wringe, 'Collective obligations: Their existence, their explanatory power, and their supervenience on the obligations of individuals', *European Journal of Philosophy*, 24, 2 (2016): 472–497. <https://doi.org/10.1111/ejop.12076>
- 36 Wringe 2010 op. cit.
- 37 This problem is also present in Young's work, in my view, for instance, when she claims that we need to fight structural injustice by "joining together with others in collective action" (Young 2006 op. cit., p. 123).
- 38 Raimo Tuomela, *Social Ontology: Collective Intentionality and Group Agents* (New York: Oxford University Press, 2013); Bratman op. cit.; Pettit & Schweikard op. cit.
- 39 Scholars disagree on the strength of interdependent knowledge requirements for collective action and intention. For example, Bratman's (op. cit.) and Pettit and Schweikard's (op. cit.) accounts contain a common-knowledge condition. Olle Blomberg has shown that common knowledge is not necessary even for accounts that are built on a notion of interlocking plural intention, such as Bratman's. See Olle Blomberg, 'Common knowledge and reductionism about shared agency', *Australasian Journal of Philosophy*, 94, 2 (2016): 315–326.
- 40 Felix Pinkert, 'What we together can (be required to) do', *Midwest Studies in Philosophy*, 38, 1 (2014): 187–202. <https://doi.org/10.1111/misp.12023>
- 41 Pinkert op. cit., p. 194.
- 42 Pinkert op. cit., p. 197.
- 43 Bratman op. cit.
- 44 There may also be options for ordinary citizens to fight injustice by making discerning consumer decisions *en masse*. I will not discuss this option, but I merely point to the complexity of the issue as debated in Eivind Jacobsen & Arne Dulsrud 'Will consumers save the world? The framing of political consumerism', *Journal of Agricultural and Environmental Ethics* 20 (2007): 469–82.
- 45 Along similar lines, Tracy Isaacs invoked a 'clarity condition' for ascribing obligations to members of loose (in her terminology: putative) collectives: "Only when the course of action presenting itself is clear to the reasonable person is it accurate to think in terms of the collective obligations of putative groups". See Tracy Lynn Isaacs, *Moral Responsibility in Collective Contexts* (Oxford University Press, 2011) p. 152.
- 46 Scott Shapiro, 'Massively shared agency', in Manuel Vargas & Gideon Yaffe (eds.) *Rational and Social Agency: The Philosophy of Michael Bratman* (Oxford; New York: Oxford University Press, 2014), pp. 277.
- 47 Shapiro op. cit., p. 278.
- 48 Shapiro op. cit., p. 277.
- 49 In fact, he criticizes the way most philosophers of action approach collective agency as having 'largely concentrated on analyzing shared activities among highly committed participants. . . . This restriction, however, has rendered these theories inapplicable to instances of massively shared agency'. (Shapiro op. cit., p.258)
- 50 That is, my argument contradicts views that purport to show that loose groups do not have collective ability at a global scale, including my own previously published argument (Schwenkenbecher 2013b op. cit.) as well as Pinkert's (op. cit.) on the one hand, and views that purport to show that we have such ability, but which base that claim on a flawed analogy between global and small-scale collective action problems on the other hand (Cullity op. cit.; Wringe 2010 op. cit.).
- 51 Anne Schwenkenbecher, 'Collective moral obligations: 'We-reasoning' and the perspective of the deliberating agent', *The Monist*, 102, 2 (2019): 152.
- 52 This refers to the two-person shallow-pond scenario by Cullity discussed earlier.
- 53 Wringe 2010 op. cit.
- 54 A plural nonmoral predicate is 'playing a duet', for instance.
- 55 Pinkert op. cit.; Schwenkenbecher 2013b op. cit.
- 56 Schwenkenbecher 2013b op. cit.
- 57 Sean Aas. 'Distributing collective obligation', *Journal of Ethics and Social Philosophy*, 9, 3 (2015): 1–23; Stephanie Collins, 'Collectives' duties and collectivisation duties', *Australasian Journal of Philosophy*, 91, 2 (2013): 231–248.
- 58 Holly Lawford-Smith, 'What 'we'?', *Journal of Social Ontology*, 1, 2 (2015): 225–249.

- 59 For a detailed discussion, see Anne Schwenkenbecher, 'Making sense of collective moral obligations: A comparison of current approaches', in Tracy Isaacs, Kendy Hess, & Violetta Igneski (eds.), *Collectivity: Ontology, Ethics and Social Justice* (Rowman and Littlefield, 2018), pp. 109–132.
- 60 Most scholars agree that group-level obligations only exist where there are members' obligations, but I do not have the space to discuss the question. Let me point out, although, that Virginia Held argued that in cases where finding a collective solution to an urgent problem is difficult, we may hold the group responsible (in the retrospective sense) for failing to solve the problem, but not the individuals. See, Virginia Held, 'Can a random collection of individuals be morally responsible?' *The Journal of Philosophy*, 67, 14 (1970): 471–481.
- 61 Schwenkenbecher 2019 op. cit.
- 62 Tracy Lynn Isaacs, *Moral Responsibility in Collective Contexts* (Oxford University Press, 2011); Anne Schwenkenbecher, 'Joint moral duties', *Midwest Studies in Philosophy*, 38, 1 (2014): 58–74. <https://doi.org/10.1111/misp.12016>; Wringe 2016 op. cit.
- 63 Christopher Woodard, 'Group-based reasons for action', *Ethical Theory and Moral Practice*, 6, 2 (2003): 215–229. <https://doi.org/10.1023/A:1024458623591>; Christopher Woodard, 'Rationality and the unit of action', *Review of Philosophy and Psychology*, 2, 2 (2011): 261–277.
- 64 Schwenkenbecher 2019 op. cit.
- 65 Cullity op. cit.
- 66 Schwenkenbecher 2019 op. cit.
- 67 Raul Hakli, Kaarlo Miller, & Raimo Tuomela, 'Two kinds of we-reasoning', *Economics and Philosophy*, 26, 3 (2010): 291–320; Robert Sugden & Natalie Gold, 'Theories of team agency' in Fabienne Peter & Hans Bernhard Schmid (eds.) *Rationality and Commitment* (Oxford: Oxford University Press, 2007).
- 68 Anja Karnein, 'Putting fairness in its place: Why there is a duty to take up the slack', *Journal of Philosophy*, 111, 11 (2014): 593–607; David Miller, 'Taking up the slack? Responsibility and justice in situations of partial compliance' in Carl Knight & Zofia Stemplowska (eds.) *Responsibility and Distributive Justice* (Oxford: Oxford University Press, 2011), pp. 230–245; Anne Schwenkenbecher, 'Bridging the emissions gap: A plea for taking up the slack', *Philosophy and Public Issues*, 3, 1 (2013a): 273–301; Cullity op. cit.
- 69 Cristina Bicchieri, *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms* (New York: Oxford University Press, 2017); Emma F. Thomas, 'The role of efficacy and moral outrage norms in creating the potential for international development activism through group-based interaction', *British Journal of Social Psychology*, 48, 1 (2009): 115–134. <https://doi.org/10.1348/014466608X313774>; Emma F. Thomas, Craig McGarty & Kenneth I. Mavor, 'Aligning identities, emotions, and beliefs to create commitment to sustainable social and political action', *Personality and Social Psychology Review*, 13, 3 (2009): 194–218. <https://doi.org/10.1177/1088868309341563>.
- 70 Young 2006 op. cit., p. 28
- 71 Bicchieri op. cit.