

Abilities and Obligations: Lessons from Non-agentive Groups

Stephanie Collins

This is the final draft of a paper whose definitive version will be published in *Erkenntnis*.

Abstract

Philosophers often talk as though each ability is held by exactly one agent. This paper begins by arguing that abilities can be held by groups of agents, where the group is not an agent. I provide a new argument for—and a new analysis of—non-agentive groups’ abilities. I then provide a new argument that, surprisingly, obligations are different: non-agentive groups cannot bear obligations, at least not if those groups are large-scale such as ‘humanity’ or ‘carbon emitters.’ This pair of conclusions is important, since philosophers who endorse large-scale non-agentive groups’ abilities almost universally endorse their obligations. More importantly, the twin arguments (one for abilities, one against obligations) make the following novel contribution: abilities imply agency-involving explanations, while obligations imply action-guidance. This general conclusion should be of interest beyond social ontology.

Introduction

Groups are a good testing ground for agentive and normative concepts. This is particularly true of ‘non-agentive groups’—groups constituted by agents, that are not themselves agents. Non-agentive groups are a good testing ground because they pull us in two opposed directions.

First, these groups are constituted by agents. This tempts us to shoulder the group with agentive baggage: abilities, intentions, actions, beliefs, desires, and so on.¹ If we're giving a group agentive baggage, it seems easy to give it normative baggage: obligations, responsibility, reasons, and so on.² It's tempting to apply all our agentive and normative concepts to them.

Second, though, non-agentive groups are not themselves agents. This tempts us to remove the baggage, unpack it, and distribute its contents to members—so that we talk about what all, most, or some members intend, believe, should do, did wrong, and so on.³ It's tempting to refrain from applying our agentive and normative concepts to these groups.

Consider the carbon-intensive global economy and resultant catastrophic global warming. Who, if anyone, is able to make the economy more sustainable? Who, if anyone, acts to maintain carbon-intensity, with relevant beliefs and desires? Who, if anyone, has an obligation to transform this? Who must take responsibility? Who will be blameworthy if transformation doesn't happen?

It's tempting to answer each question with 'all of us together': humanity is able to transform the economy; humanity acts to maintain it; humanity tacitly believes it's permissible;

¹ Their abilities are endorsed by Pinkert (2014), Aas (2015), Collins (2019), Schwenkenbecher (2021), and Wringe (2020a), whom I discuss below; their intentions and actions are endorsed by Bratman (2014); their beliefs by Mathiesen (2006).

² Their obligations are endorsed by Björnsson (2014), Pinkert (2014), Wringe (2016), and Schwenkenbecher (2021); responsibility by Feinberg (1968), Held (1970), Jackson (1987), May (1987), Isaacs (2011); reasons by Parfit (1984, sec. 26) and Dietz (2016).

³ List (2014) defends this for beliefs and desires; Lawford-Smith (2015) and Collins (2019) for obligations.

humanity desires it continue; humanity has an obligation to change it; humanity must take responsibility; humanity will be to blame. Such assertions seem necessary to capture the scale of the problem, which cannot be solved by any agent acting alone.

It's also tempting to unpack that baggage and distribute its contents: to say a particular holidaymaker is able to fly less or a specific diplomat is responsible for their contributions to the failed Copenhagen summit in 2009. Such attributions seem necessary to guide action and demand recompense. These attributions don't relate to 'global warming,' but to specific contributions (flying, negotiating). So, if we take the distributive strategy, we end up with some different abilities, actions, obligations, and so on, than if we apply our concepts to humanity.

It's even more tempting to do both: give humanity the agentive and normative burdens of global warming *and* individuals the agentive and normative burdens of their contributions. This satisfies both sides of our ambivalence.⁴

Given all this temptation, it's worth examining each concept separately. Maybe some apply only to groups, maybe some apply only to individuals, maybe some apply to both. To differentiate these, we must proceed on a concept-by-concept basis. This article focuses on abilities and obligations. These are important for three reasons. First, ability is agentive while obligation is normative: differentiating them might illuminate the two categories more generally. Second, regarding problems like global warming, a concern with intentions, beliefs, desires, blame, and (pro tanto) reasons looks frivolous: more important is what we are *able* and *obligated* to do. Third, in social ontology, ability and obligation are almost universally

⁴ Cripps (2013) posits both collective and individual climate-related obligations. Isaacs (2011) argues individual-level and group-level responsibility can co-exist.

presumed to “pattern”: either non-agentive groups can have *both* abilities and obligations; or they can have *neither*.⁵ A non-patterned approach has been neglected.⁶

I argue large-scale non-agentive groups have abilities, but lack obligations. If this is correct, the concept-by-concept strategy is vindicated. The argument makes another contribution: it tentatively maps the general difference between agentive and normative concepts, suggesting normative concepts guide agents while agentive concepts explain outcomes agentively. If this is correct, there is (surprisingly) more scope for applying agentive concepts than normative concepts to non-agentive groups.

Overall, this paper makes two main contributions. First, I provide new arguments for treating obligations and abilities *asymmetrically* in application to non-agentive groups. Second, my twin arguments (one for abilities, one against obligations) imply the following: abilities concern agency-involving explanations, while obligations concern action-guidance. These implications hold interest beyond social ontology.

Section 1 examines ability. Philosophers often talk as if each ability is held by exactly one agent.⁷ I provide a new three-part argument for non-agentive groups’ abilities. I critique

⁵ Lawford-Smith (2015) argues for the latter. Wringer (2005; 2020a), Isaacs (2011), Björnsson (2014), Pinkert (2014), Aas (2015), Miller (2015, 2020), and Schwenkenbecher (2021) argue for the former. (While Miller (2001, ch. 8) denies collective moral responsibility, his other work seems to view joint abilities and obligations as patterning.)

⁶ I have previously argued against patterning (Collins 2019), but I will problematise those arguments and provide new ones.

⁷ van Inwagen 1983, 8-13; Cross 1986; Brown 1988; Mele 2003; Vihvelin 2004; Fara 2008; Mandelkern et al 2017; Maier 2018.

other philosophers' analyses of non-agentive groups' abilities, before offering a new analysis. I outline eight virtues of my analysis. The virtues of my analysis—combined with my new three-part argument—place us on firmer footing for defending non-agentive groups' abilities, compared with others' recent defences of them.⁸ Section 2 argues non-agentive groups cannot bear obligations. There are two premises here. First, non-agentive groups cannot make decisions. Second, an entity can bear an obligation only if the entity can make a decision. I provide a defence of these.⁹ Surprisingly, while humanity likely is able to alleviate global warming, it has no obligation to do so.

Before beginning: what makes a group non-agentive? I will take up Carol Rovane's account of group agency, on which a group agent exists when there is a commitment to achieving overall rational unity at the group level.¹⁰ Achieving rational unity means arriving at a consistent and interlocking bundle of beliefs, preferences, hopes, and fears, and so on (what Rovane calls 'mental episodes'). A non-agentive group is a group in which there is *no existing commitment* to arrive at such a bundle (though it is not necessary, for agency, that such a bundle has already been achieved). I am neutral on whether the commitment in question must be *one commitment* held at the group level (as in Margaret Gilbert's 'joint commitments'¹¹) or whether members might be severally committed to group-level rational unity and thereby generate

⁸ See fn. 5.

⁹ I used these premises in earlier work (Collins 2019), but that defence was critiqued extensively by Björnsson 2020; Blomberg 2020; Schwenkenbecher 2020; Wringe 2020b. I will explain two large gaps in my earlier defence (see fn. 39), will provide a different defence.

¹⁰ Rovane 1998, 8, *passim*.

¹¹ Gilbert 1989.

group agency (as, arguably, under Raimo Tuomela’s ‘we-mode’ account¹²). I assume a group is non-agentive unless such a commitment already exists; the advent of such commitments brings group agents into existence.

Now, there is a difficult question of when and how such commitments arise—particularly in the small-scale rescue cases that permeate the literature on collective obligations, but also in large movements such as vegetarianism or Anonymous. (Since, arguably, at least some large social movements contain commitments to group-level rational unity.) I focus only on those groups that all theorists agree are non-agentive: large-scale groups that clearly lack the Rovane-style commitment. These are groups like ‘humanity,’ ‘carbon emitters,’ or ‘global economy participants.’¹³ The focus on such groups distinguishes my discussion from many in the literature. Such groups, nonetheless, have high politico-socio-economic importance.

1. Ability

Broadening Agentive Concepts

Philosophers often talk as if each ability belongs to one agent. A standard picture runs as follows.¹⁴ There are dispositions—including this vase’s disposition to break when dropped,

¹² Tuomela 2013.

¹³ Shapiro (2014) analyses “massively shared agency,” but he utilises Bratman’s analysis, on which *shared agency* does not imply *a group agent*. Aas (2015, 2) reasons from individuals’ obligations, to groups’ obligations, to a group agent. However, it’s doubtful groups like ‘emitters’ have the kind of individual obligations Aas uses to start his reasoning.

¹⁴ Maier 2014; van Inwagen 1983, 8-13. I endorse this picture, barring interpretations of it which say each ability belongs to just one agent.

that pill's disposition to cure headache, your disposition to understand English, and my disposition to speak English. Some dispositions are held by agents (e.g., you or me); others are held by non-agents (e.g., vases or pills). Dispositions held by agents are powers. There are two types of power: first, powers that range over unintentional behaviours, including your power to understand English (assuming understanding is a non-action, for example because you cannot choose not to understand); second, powers that range over intentional behaviours, including my power to speak English (which I can usually choose not to exercise).¹⁵ Only the last of these are *abilities*. I say abilities range over "intentional behaviours," rather than "actions," because I assume abilities also range over cultivated habits, intentional omissions, and other behaviours that result from intentions without, perhaps, being actions.

Given this picture, one might think each ability is held by one agent. But that cannot be right. As Barbara Vetter says, "I have the ability to see; together, you and I have the ability to see each other. If Hannah has the ability to play the piano and Jane has the ability to play the flute, then Hannah and Jane together have the ability to play a duet for flute and piano."¹⁶ Presumably, these statements are true even if the group lacks a commitment to group-level rational unity (Hannah and Jane have the ability even if they have never met). For a larger-scale example, consider a diversely-motivated and spontaneous mob that storms the Capitol. Such a mob lacks a commitment to rational unity, but is able to storm the Capitol. Or consider an even more fragmented group: market participants, who have the ability to standardise milk prices via 'spontaneous order' or the 'invisible hand.' (Hayek 1969; Smith 1776) Or consider the ability of humanity to make Earth uninhabitable.

¹⁵ The understanding-versus-speaking example is van Inwagen's (1983, 10).

¹⁶ Vetter 2015, 105. As Vetter describes the duo, they might never have met, so they fail even the thinnest theory of group agency.

Importantly, the outcomes these groups are able to produce are more than the conjunction of the outcomes each member is able to produce if acting in isolation. Consider storming the Capitol. Jacob is able to play his part, Joe is able to play his, and so on. But the mob's storming the Capitol requires more than the conjunction of these parts. It requires that these parts happen *in a synchronised way*. Jacob cannot ensure Joe acts synchronically; nor can Joe ensure Jacob does. Jacob lacks the ability to produce the mob's storming of the Capitol—as does Joe, and everyone else. It is only mob members *together* that are able to synchronise their actions. Only the mob as a whole is able to storm the Capitol.

Likewise for standardised milk prices and planetary uninhabitability. My ability to buy milk—or pollute by flying—depends upon numerous other agents in global supply chains, whose parts I must take advantage of in the right way, if I am to do my part in standardising milk prices or making the planet uninhabitable. Specifically, I must buy milk or flight bookings off these other agents, *while* they are selling milk or flight bookings. The relations between our parts makes the outcome that the group produces more than the mere conjunction of the outcomes the members could each produce if they performed their parts in isolation. The outcome emerges only via complex relations amongst individuals' parts, through which those parts are networked. The *network of related individual parts* is a group-level phenomenon. This group-level phenomenon gives us reason to posit group-level abilities, because the parts are intentional behaviours and abilities range over intentional behaviours. That said, a group-level ability does not *automatically follow* from a group-level network of individual parts. My claim is rather that the presence of a group-level network of individual parts, which produces some outcome, gives us *one reason* to posit a group ability to produce that outcome.¹⁷

¹⁷ Elder-Vass (2007) argues emergent effects arise out of relationality, which gives reason for non-reductive realism about the network that causes the emergent effect. In earlier work

Multiple realisability is second reason. Humanity's ability to make the planet uninhabitable would remain if there were different types of individual abilities. If I ceased having money to travel by plane, but someone else acquired the ability to cut down trees, then humanity would remain able to destroy the planet. So there are multiple ways of 'realising' humanity's ability to destroy the planet. By including multiply realisable properties in our ontology, we can use our ontology to avoid pointlessly detailed descriptions of our world (in which we list all abilities of all individuals), we avoid giving long lists of disjuncts when stating the mechanisms by which something might happen, and we can generalise across our world and similar nearby worlds (worlds in which the same property, differently realised, exists). So, multiple realisability gives us a second reason to posit humanity's ability.¹⁸

A third reason for positing non-agentive groups' abilities is the causal-explanatory power that we would achieve, were we to posit those abilities. Consider the group 'market participants.' We can partly explain the standardisation of milk prices via the abilities, greed, unthinkingness, etc of individual members. But the standardisation of milk prices receives a more coherent, unified, and proportionate explanation if we say the market participants were able to produce it; after all, the individual attributes were not exercised in isolation. By positing a group-level ability, we can better explain how the outcome came about.

Relationality, multiple realisability, and causal-explanatory power give us three reasons to posit abilities of non-agentive groups. Again, group abilities do not *follow automatically* from these three reasons. Rather, each reason provides some support (and together they provide

(Collins 2019, 81-85), I discussed relationality as a reason to posit groups' abilities, but I didn't distinguish multiple realisability and explanatory power as I will below.

¹⁸ Putnam (1967) famously used multiple realisability to argue against psycho-physical reductionism. Sawyer (2005), amongst others, extended it to the social.

fairly weighty support) for positing group abilities. Indeed, other philosophers have argued that non-agentive groups have abilities.¹⁹ But those arguments have proceeded by giving examples. The three reasons just mentioned provide us with a new—more theoretically-motivated—argument.

Still, an opponent might insist we're not talking about *abilities* when we say 'the mob is able to storm the Capitol.' Instead, we're saying something is *possible* or *likely*, as in 'Krakatoa has the ability to erupt' or 'this pill has the ability to cure headache.' After all, 'ability' is multiply ambiguous, including agentive and non-agentive uses. We're examining 'ability' as an agentive concept. And, one might continue, agentive concepts can be applied only to agents.

Philosophers who have advocated non-agentive groups' abilities have not answered this objection. I answer: we can apply agentive concepts to non-agentive groups, because those are constituted by agents and the concepts concern the intentional behaviours of those agents. This distinguishes non-agentive groups from Krakatoa or pills. We should retain the distinction between agentive and non-agentive concepts. But we should ask why the distinction matters. I suggest the distinction matters because we care about the difference between possible-via-agency and possible-via-non-agency. This distinction helps us communicate whether intentional behaviours were behind some past or present outcome, or could come to be behind a potential future outcome. That is, the distinction between abilities and non-ability dispositions helps us explain and understand *the mechanisms by which states-of-affairs can come about*. The distinction between agentive and non-agentive dispositions has not been discussed by those who defend non-agentive groups' abilities, yet it helps us reply to opponents who worry that non-agentive groups' abilities imply inanimate objects' abilities.

¹⁹ See fn. 1.

For explaining and understanding our world, it matters whether the mechanisms that lead to outcomes are agentive (abilities) or non-agentive (non-ability dispositions). To illustrate, imagine a passage cutting through a forest. Perhaps the passage is a quirk of nature: the plants happened to grow leaving a passage. Or perhaps the passage resulted from thousands of people, uncoordinated over time, taking that path through the forest—gradually clearing away plants. One is not the result of agency; the other is. We care which is the case. We want to capture the difference in our conceptual repertoire. This is so even though the group is non-agentive and no member was able to clear the passage alone. Since the concept ‘ability’ can do this for us, we should resist Krakatoa’s or pills’ abilities.

Importantly, the individuals’ intentional behaviours of forest-clearing were *related in a not-merely-conjunctive way* (not too far apart in time), and those actions were *multiply realisable* (across worlds in which the passage resulted), such that positing the group’s ability *helps us to explain* the passage. The italicised points reflect the three reasons given above, so they matter for the conclusion that there is a group-level ability. These points will be reflected in my below analysis of non-agentive’ groups abilities.

Still, someone sceptical of non-agentive groups’ abilities might say the forest passage was the emergent effect of the agency of each member of the group, and the exercise of their individual abilities. Such a sceptic has agency firmly within their explanatory picture.²⁰ But their explanatory picture neglects the relationality, multiple realisability, and unified explanation that we find in groups’ abilities. The sceptic’s explanation is just a conjunction of individual-level abilities. Alternatively, if the sceptic aims to capture these features (relationality, multiple realisability, and scope for unified explanation), then the sceptic gains strong reason to posit a group-level agential explanation: an ability.

²⁰ I thank an anonymous reviewer for pressing this.

My argument doesn't gratuitously open the floodgates to any old ability of any old non-agentive group. The upcoming analysis will ensure this. Take the 'ability' of my mother and me to hear each other. When we are nearby, we cannot refrain from exercising this ability via choice or will (we have tried). Hearing each other is akin to understanding a billboard: we cannot choose not to, so it is not an ability. These examples contrast with our ability to make eye contact, which we can refrain from exercising by closing our eyes. Our hearing each other does not find its best and most unified explanation in positing a group-level ability that arises out of the non-conjunctive relations between the multiply-realizable exercises of our *intentions*, wills, or choices (including the habits, omissions, skills, and other doings that result from these). In this way, our hearing each other differs from our making (prolonged) eye contact—and Capitol stormings, standardised milk prices, and planetary uninhabitability. Abilities are exercised via agency. Thus, ability is still agentive, even once we allow non-agentive groups' abilities. Again, philosophers who endorse non-agentive groups' abilities have not made this reply to possible critics.

If abilities produce outcomes via agential mechanisms (such as intentions, wills, or choices), then the axiom 'ability is an agentive concept' should be interpreted as 'ability is applicable *either* to an agent *or* to an entity that is constituted by agents.' This interpretation of ability's status as an 'agentive concept' has a nice triple upshot: it accommodates the intuitive examples of groups' abilities, it respects the axiom that agency underpins abilities, and it prohibits us from attributing abilities to volcanoes or pills. The next task is to provide an analysis of non-agentive groups' abilities that accommodates the intuitive examples and the three reasons for positing the abilities, without overinflating the concept.

Proposed Analysis

Other philosophers have advocated, and analysed, non-agentive groups' abilities. None have provided the arguments and distinctions introduced above. Still, with those arguments and distinctions in hand, perhaps an existing analysis will serve us well. Unfortunately, each existing analysis has limitations.²¹

First, Felix Pinkert suggests

agents *aa* are immediately jointly able to ϕ if and only if there is exactly one salient possible collective pattern of actions of the relevant *aa*-s that constitutes *aa* ϕ -ing, and which is such that every relevant agent believes of the action which is her part in that pattern that she needs to perform this action if *aa* is to ϕ .²²

This won't work for Capitol stormings, milk prices, or environmental destruction: there are numerous (perhaps numerous salient) patterns of action that can produce these outcomes—not just one salient pattern. The outcomes are multiply realisable. Strangely, Pinkert's analysis countenances collective abilities only if the ability is *not* multiply realisable in members' abilities. Yet an absence of multiple realisability is usually a reason for elimination of the higher-level concept. Moreover, Pinkert's analysis requires each agent to *believe* that she *needs* to *act* in some way if *the group is to ϕ* . However, for two of my examples (milk prices and planetary inhabitability), each italicised clause is too strict: an agent might have no beliefs

²¹ I thank an anonymous reviewer for encouraging discussion of these analyses. Isaacs (2011) and Björnsson (2014) also provide examples of groups' abilities, but they don't provide analyses (individually necessary and jointly sufficient conditions).

²² 2014, 194. Pinkert also analyses 'mediated' and 'recursive' joint abilities, but the analysis just quoted is built into them.

about her contribution; still less the belief that she needs to perform it; still less that it is an action (which it might not be: it might be a habit, skill, omission, etc); still less that it relates to the group outcome (which she might not conceptualise at all).²³

Second, Anne Schwenkenbecher proposes:

Two agents a and b have a joint ability to do x if a has individual ability to do x_a and b has individual ability to do x_b where x_a and x_b produce x , both actions are compossible, and a and b are in principle capable of intentionally combining them.²⁴

This states a sufficient condition, but Schwenkenbecher labels it a *definition*, so it seems intended as sufficient and necessary (an analysis). I will interpret Schwenkenbecher's definition as expandable to any number of agents (not just two). With these adjustments, Schwenkenbecher's analysis has virtues over Pinkert's: it doesn't require exactly one salient pattern (x_a and x_b can be types or disjunctions), nor that agents have specific beliefs about their contributions.

However, Schwenkenbecher's agents must be "in principle capable of intentionally combining" their contributions; "they must be in a position to *intentionally perform their contributory action as a contributory action*."²⁵ This is too strong. We can imagine instances of standardised prices or environmental destruction where the outcome will occur only if people don't *intentionally combine* their contributions: the outcome might depend on agents

²³ Schwenkenbecher (2021, 48-51) raises similar worries, though she interprets Pinkert as not requiring that individuals conceptualise the outcome. Wringe (2020a) raises worries for Pinkert's belief requirement, though without providing his own analysis.

²⁴ 2021, 54; similarly 2021, 15.

²⁵ 2021, 15a, emphasis original.

thinking self-interestedly and disregarding the upshots of the combination. The group can exercise the ability, as long as agents don't *intentionally* combine their contributions. (This is analogous to my ability to fall asleep, which I can exercise only by doing so unintentionally.) Contra Schwenkenbecher, the agents' contributions needn't even be *actions*: they could be behaviours that result from habit, as long as they are ultimately explained by the performer's intentions, will, choices, or similar.

Third, Sean Aas follows the conditional account of individuals' abilities, on which '[t]o be *able* to ϕ is to be such that one would (or would be likely or disposed to) ϕ if one *tried* to ϕ ."²⁶ Aas aims to demonstrate by examples that uncoordinated groups can try to do things. Again, though, not all exercises of abilities involve tryings. Consider the Capitol storming, standardized milk prices, and planetary uninhabitability. The Capitol storming might involve tryings, though it's doubtful all mob members have a clear sense of *what* they're trying to do. Milk prices and planetary uninhabitability, however, involve many participants who have never thought about market-level or environment-level outcomes: they're just trying to buy milk or go on holiday. Moreover, Aas says a group's trying can "be, simply" all individuals trying.²⁷ Elsewhere, I've noted that such an 'identity' theory of group-level and individual-level phenomena is reason to eliminate the group-level phenomenon from our ontology.²⁸ My analysis will offer no such 'identity' theory.

That said, Aas's target is *those group abilities that are necessary for group obligations*. Aas might countenance a different analysis for broader purposes—but Aas doesn't provide such an analysis. Further, Aas may well be right that the 'ability' at issue in 'obligation implies

²⁶ 2015, 15, emphasis original.

²⁷ 2015, 16.

²⁸ 2019, 69.

ability' requires that the ability can be willingly exercised. (Schwenkenbecher might also defend her analysis using this point.) Still, ability *simpliciter* does not require that the ability can be willingly exercised.²⁹

Finally, elsewhere I have proposed:

A non-agential group is able to produce an outcome X at a time t just in case (1) each member has an individual ability at t to perform actions that contribute to X; and (2) given that enough members exercise the abilities in (1) at t, each will do their contributory part of a pattern of behaviours that will robustly secure X in the absence of defeaters.³⁰

Again, this restricts itself to *actions* of members. But individuals are *able* to do non-actions (like fall asleep), if those behaviours result from intentions (like the intention to exercise and get enough sunlight). The same should hold for non-agentive groups. Also, the temporal constraints contained in my earlier analysis limit that analysis to time-slice abilities: the contributions must be performable when the outcome occurs. My upcoming analysis will allow temporal separation between exercises and outcomes. That said, my analysis will borrow my earlier formulation “part of a pattern of behaviours that will robustly secure X”—for reasons explained below.

There is a bigger problem for both Aas’s analysis and my earlier analysis: they use a conditional structure. That is, an ability is exercised if, given some condition, an outcome occurs. This makes these analyses susceptible to familiar problems with conditional analyses

²⁹ I thank an anonymous reviewer for this distinction between ability *simpliciter* and ability *in obligation-implies-ability*.

³⁰ 2019, 71.

of abilities: if the condition's occurring would 'mask' or 'fink' the ability, then the entity counterintuitively lacks the ability.³¹ This problem is considered fatal by many metaphysicians working on individuals' abilities.

In light of the foregoing, I propose the following simple analysis of non-agentive groups' abilities:

A non-agentive group has the ability to produce outcome X just in case each member has the ability to perform behaviours that are a part of a pattern of behaviours across the group that will robustly secure X.

This analysis doesn't demand a particular pattern of actions; doesn't require beliefs of members; is not an identity theory; doesn't demand actions that are co-temporaneous with the outcome; and doesn't have a conditional structure. My analysis requires explanation, which will uncover eight virtues.

First, I analyse non-agentive groups' abilities via members' abilities. I assume the correct analysis of members' abilities will refer to their intentions, wills, choices, or similar; thus, abilities are about intentional behaviours. This is a first virtue: the analysis captures groups' abilities' dependence on the agency of their members. I argued that abilities capture agency in our explanations, so agency (of members) is in the analysis.

Second, the analysis captures all three reasons for including a group's ability in our ontology: first, relations between members matter (the analysis mentions a *pattern* of behaviours); second, the exercise of the group's ability is multiply realisable in the exercise of members' abilities (the analysis mentions *a part of a pattern*); third, the analysis attributes explanatory power to the ability (since the pattern *secures* the outcome). The first virtue reflects

³¹ Maier 2014.

the dependence relation between groups' abilities and members' abilities; the second virtue denies a reduction or identity relation.

Third, the analysis captures the intuitive examples of group abilities. Each mob member is able to perform behaviours that are part of a pattern of behaviours across the group that will robustly secure the storming of the Capitol. Each market participant is able to perform behaviours that are part of a pattern of behaviours across the group that will robustly secure standardised milk prices. Each participant in the global economy is able to perform behaviours that are part of a pattern of behaviours across the group that will robustly secure a destroyed environment. Members need not conceptualise their behaviours as contributing.

Fourth, the analysis captures the way in which supra-groups inherit the abilities of their sub-groups. Suppose Anna and Berta are two independently operating investigative journalists who together have the ability to bring down the government. (Perhaps they each know a different scandalous fact about the government and if both of these facts were revealed, then the government would be brought down.) Now consider the group consisting of Anna, Berta, and Carol. Carol is another journalist who knows nothing scandalous. The three-person group has the ability to bring down the government, since each can play their part in a relevant pattern of actions. For all these patterns, Carol's part will be staying out of the way. Furthermore, in this example, humanity—the large-scale kind of non-agential group with which my argument is concerned—has the ability to bring down the government. It might be pointless or unhelpful to state this in many contexts, but it is nonetheless true.³²

Fifth, the analysis captures unexercised abilities. Consider my ability to jump on the spot right now. I won't exercise that ability, but our analysis of individuals' abilities should

³² I thank two anonymous reviewers for pressing me on this point, one of which provided the journalists example.

include it. Likewise, consider humanity's ability to produce 'everyone on the planet, who can touch their nose, touching their nose at 4.00pm next Tuesday.' Each member of humanity is able to do their part of a pattern of behaviours across the group that will robustly secure the nose-touching outcome. It might seem ludicrous to posit that ability.³³ But this is just because it's *unlikely* enough people will do their part. If they gained reason to do so (say, if aliens beamed a universally-accessible message: 'touch your noses or face planetary destruction!'), then the ability would not be in question.³⁴ The inclusion of such examples is the account's fifth virtue. This is important for political purposes: it prevents us from claiming inability when we simply lack will.

Sixth, the analysis excludes examples that could happen only by fluke. Imagine fifty random philosophers at a conference. Is the group able to produce 'each member hitting the bull's eye when playing darts tonight'? No. Each member's relevant part is 'throwing towards the bull's eye.'³⁵ Each member is able to perform their part. However, those parts will not *robustly* secure the outcome. In most nearby worlds, at least one philosopher will miss. The outcome is modally fragile, even if it occurs. So there's no ability.³⁶

³³ Stemplowska 2016, 276-7.

³⁴ Similarly Pinkert 2014, 196; Schwenkenbecher 2021, 57.

³⁵ I assume 'hitting the bull's eye' is not an intentional behaviour, but an outcome of intentional behaviour; if I'm wrong, then it's false that each member has the ability to their part, so the group's ability still rightly gets excluded.

³⁶ The darts example and modal conception of 'flukiness' come from Southwood and Wiens' (2016) argument that 'actual' does not imply 'feasible.' I suggest that—for the same reasons Southwood and Wiens give regarding feasibility—'actual' does not imply 'was able.'

Seventh, the ‘robustness’ requirement implies there weren’t abilities to produce many actual events *in the way they were produced*. For example, there was no ability, in early 1989, to produce the outcome ‘the Berlin Wall falling on 9 November 1989.’ That outcome is too specific. There was no group such that each member doing their part would *robustly* produce it. Yet there probably was the ability to produce ‘the Berlin Wall falling.’ We can be more or less coarse-grained in our characterisation of outcomes, resulting in more or less robust securing of them when members exercise the relevant individual-level abilities, resulting in a group-level ability, or not. This sensitivity to precision is a seventh virtue.

The scalarity of robustness produces the eighth virtue: a scalar account of groups’ abilities. A group is *more able* to produce some outcome, the more robust the production of that outcome is across various possible scenarios in which the members exercise the relevant abilities. Thus, the East German people in 1989 were *more able* to produce the fall of the Berlin Wall than were the East German people in 1969, just in case the 1989 group was a *more robust* securer than the 1969 group. This raises the question: at what point does ‘less robust’ become ‘insufficiently robust to attribute an ability’?³⁷ It is not possible to give a numerical answer to this. It is likely that the cut-off point will be a contextual matter, that varies with the stakes of the outcome in question. It is further likely that it will sometimes be vague whether there is sufficient robustness to attribute an ability (just as it may be vague whether someone is sufficiently hairless to attribute baldness). The difficulty of drawing a sharp line here does not mean that there are no clear cases on either side of the line.

Given these virtues of my analysis—alongside the considerations presented in the previous sections—I conclude that large-scale non-agentive groups have abilities. While others’

³⁷ I thank an anonymous reviewer for asking this.

analyses capture much that is important, this proposed analysis is best suited to our explanatory, predictive, and aspirational purposes.

2. Obligations

It might seem a short step from abilities to obligations. Indeed, almost all philosophers who posit groups' abilities posit their obligations.³⁸ But ability and obligation are importantly different. I will argue:

1. Non-agentive groups cannot make decisions.
2. An entity can bear an obligation only if the entity can make a decision.

So,

3. Non-agentive groups cannot bear any obligation.

The next two sub-sections defend (1) and (2) in turn.³⁹

³⁸ See fn. 5.

³⁹ I have addressed these premises in earlier work (Collins 2019, ch. 3), but those earlier arguments are contestable in two important places. First, my previous argument for (1) was that there is no group-level decision-making in cases with Gilbert's 'plural subjects' or Bratman's 'joint intentions.' But non-agentive groups are far more diverse than these two group-types. My (2019) argument for (1) was therefore radically incomplete. In the present paper, I will demonstrate an absence of group-level decision-making in three broad types of non-agentive group, which plausibly subsume all non-agentive groups in which one might think there is decision-making; this argument is far more encompassing than my earlier argument. Second, my earlier argument for (2) relied on a contestable idea of 'moral worth,' whereby (1) morally worthy behaviours must derive from decisions and (2) fulfilling duties

Non-agentive Groups Cannot Make Decisions

I will consider three broad cases in which one might think a non-agentive group has made a decision: aggregation, debate, and conditional commitment. Recall: I here endorse Rovane's theory, on which group agency arises once there is a commitment to achieving overall rational unity at the group level. In the cases below, assume there is no such commitment.

First, aggregation. Here, each member makes an identical decision concerning group behaviour. The decision might be "to aim at 'humanity halts climate change'," or "to do my bit towards the outcome 'humanity halts climate change'." These decisions are *to do* something the member does (aiming; one's bit), though the decision *concerns* something the group does (halting climate change). Or each member might decide 'to halt climate change': each might decide *to do* something the group does.⁴⁰

Suppose each human decides to halt climate change or to aim at 'humanity halts climate change.' You might think *humanity* decided to perform that behaviour. We can also imagine interdependence between members' decisions. Perhaps one person raises a (metaphorical) banner saying "humanity will halt climate change," which causes others to endorse the banner, each being pulled into their decision by the sight of those who have gone before, and each being kept steadfast by those who come after. Here, the relations amongst members partly explain their respective decisions—though, *ex hypothesi*, they lack a commitment to group-level rational unity. Above, relationality provided reasons for realism about group abilities.

entails moral worth. This has generated multiple counterexamples and counterarguments (see fn. 9). My argument for (2) will rely on no such concept.

⁴⁰ Núñez 2019 argues that it is rational for individuals to intend 'social' ends in this way

However, there are three reasons no group decision occurs in such cases. First, the ‘group decision’ is not multiply realisable: it can be realised only by all members making the relevant individual-level decision. We could assert that a majority or supermajority would suffice for a group decision (since these can be multiply realised)—but there is no motivation for this assertion, if the group doesn’t have a rule stating that a majority or supermajority suffices for a group decision. Such a group-level rule implies a commitment to rational unity (though the commitment to rational unity can also exist without such group-level rules).

Second, although there’s unity (because similarity) to members’ decisions, ‘humanity decided’ doesn’t hold greater explanatory power than ‘each human decided.’ The former explanation doesn’t operate over a wider range of possible worlds, or go deeper into the causal genesis in the actual world. In fact, ‘each human decided’ goes further into the causal genesis—it tells us *how* humanity decided—and should be explanatorily preferred for that reason. Positing a group-level property doesn’t provide explanatory purchase, unlike for abilities.

Third and most importantly, the determiner of each decision—the entity with discretion, choice, or power over that decision—is the individual whose decision it is. None of the individuals’ decisions is controlled for by humanity.⁴¹ Of course, other individuals—and the relations between them—*influence* each individual’s decision. But each individual has final say over their own decision. Aggregation cases are not decisions of non-agentive groups.

Debate is another case. Suppose each member makes a decision with the same content—as in aggregation—but via some members convincing or coercing others. Here, some *intentionally pressure* others into the decision; they don’t just happen to influence. Debate is a

⁴¹ Control is also emphasised by Lawford-Smith (2015), who suggests control-over-X is necessary for obligation-over-X. This is unnecessarily strong.

relational, multiply realisable, and explanatorily powerful group-level phenomenon. If the pressure-wielders know the others are pliant, then they might wield *partial control* over others' decisions. If the resultant decisions regard something the group does, perhaps it's a group decision.

These cases require care: debate is prevalent in group agents, yielding group-level decisions. Consider a parliament or job committee. Members give and receive reasons, each attempting to convince the others to aim at a group behaviour ('passing this bill' or 'shortlisting this candidate'). But we are concerned with non-agentive groups, in which there is no commitment to overall rational unity at the group level. To ensure the debating group remains non-agentive, imagine an internet forum where members are discussing what each will do, without commitment to group-level unity. If the debate leads each member to ultimately decide to aim at 'the group halts climate change,' perhaps the group has decided to halt climate change.

Again, though, the decisions are made by individuals. Perhaps some members—as individuals—are accountable for their role in the formation of others' decisions. But to think *the group* makes one decision (or makes each individual's decision) is to undermine the authority each has over her own decision. And even if some cede their decision-making authority, they do not cede it to the group. Imagine a member says 'I'll do whatever others in this forum do.' She means *she* will decide to do whatever a many (or most) others decide for themselves to do. It's not that the group decides whatever many (or most) group members decide, such that she has decided to follow *the group's* decision.⁴²

Third and finally, perhaps *conditional commitments* enable non-agentive group decision-making. Here, each member makes a commitment to each other, where the

⁴² List (2014) says likewise about beliefs and desires.

commitment is to do one's part in a pattern of group behaviour, on condition that others commit likewise. Each has full-blown discretion—licensed by the others—over the behaviours the others perform. David Velleman explains:

when I say 'I'll go for a walk if you will' and you say, 'Then I will' ... each of [the two statements] determines the speaker's behavior, and represents itself as determining it, only in conjunction with the other's statement. ... the causal powers of these statements are in fact interdependent, as the statements themselves represent. Hence the behavior of each is settled—and is represented as being settled—only by both of us together. ... Each of us places his behavior under the joint control of both, by issuing an intention that's conditional on the other's intention. ... The result is that each of us conditionally settles, and is represented as conditionally settling, one and the same set of issues—namely, how both of us will behave—and we thereby categorically settle those issues together. (1997, 48)

My going for a walk is caused by the combination of our conditional commitments; likewise, your going for a walk. According to Velleman, each walk is caused by the combined conditional commitments in the way actions are typically caused by decisions: the combined conditional commitments are a representation that causes behaviour by representing itself as causing it.

This doesn't make for a group decision.⁴³ If it did, we would have an overabundance of group decisions. Suppose I'm sitting in traffic and decide 'I'll turn left if the driver ahead of

⁴³ Velleman would perhaps agree. He adds two conditions, discussed below. Even with those conditions, it's not clear Velleman believes *the group* decides. But someone might use his example to argue this.

me decides to turn left.’ My turning left is not caused just by my conditional decision; nor is it caused just by the other driver’s decision to turn left. My behaviour is caused by both decisions together. Yet to characterise my turning as a *group decision* implicates unwitting drivers in random decisions of which they’re not aware. Things are not improved if the other driver decides ‘I’ll turn left if the driver behind does,’ and watches for whether I signal a left turn. Even in this symmetrical case, neither his decision nor mine are made by both of us together.

Velleman emphasises that the statements are (i) public between us and (ii) made one in response to the other (1997, 38, 47-8). These conditions are perhaps necessary for rendering our behaviour rational:⁴⁴ if the drivers are unaware of each other’s conditional commitments, then there’s no explanation if one initiates turning. If they’re rational, then only the two statements together—made publicly as statement-and-response—cause their behaviour as decisions typically cause behaviour. These conditions prevent the account from over-generating group decisions.

The question becomes: if behaviour is caused by a string of events (my expressing a conditional commitment, your comprehending that, your expressing a like commitment in response, me comprehending that, me deciding to do what I conditionally committed to do, you deciding to do what you conditionally committed to do), should that string of communications be counted as one thing? Intuitively, no. The proposed thing is not multiply realisable and doesn’t wield explanatory power (though the relations amongst its components do matter for its effects). More importantly: supposing the string of communications is one thing, is that thing a *decision*? Again, no. The string of communications is amongst the causes of each person’s behaviour. But not all causes of behaviour are decisions to perform the behaviour. For Velleman, a decision is a “representation that causes behavior *by representing*

⁴⁴ Velleman doesn’t say this.

itself as causing it.”⁴⁵ Absent the kind of rational unity that would seem to amount to group agency, it’s doubtful the string of communications represents itself.

There may be a fourth mechanism for purported decision-making in non-agentive groups. It’s difficult to see how such a mechanism wouldn’t involve a commitment to producing rational unity at the group level. Alternatively, perhaps ‘non-agentive group decision’ is a primitive, which cannot be understood in terms of its underlying mechanisms. But since we can understand these cases via individual decisions, an appeal to such primitives looks question-begging. I have offered a presumptive argument for believing non-agentive groups cannot make decisions. Until that argument is undermined, our theorising should proceed accordingly.

Obligation-bearing Entails Decision-making Capacities

Why think obligation-bearing entails decision-making capacities? Think of the creatures to which we take what Peter Strawson called the ‘participatory stance’ or what Stephen Darwall called the ‘second-personal stance.’ The participatory or second-personal stance involves addressing others with “orders, requests, claims, reproaches, complaints, demands, promises, contracts, givings of consent, commands,” and so on while presupposing “authority and accountability relations” between us and them.⁴⁶ From within the participatory or second-personal stance, we view ourselves as having a valid moral complaint if we address entities in such ways and are ignored. We adopt “the attitudes and reactions of offended parties and

⁴⁵ 1997, 36, emphasis added.

⁴⁶ Darwall 2006, 8.

beneficiaries; ... such things as gratitude, resentment, forgiveness, love, and hurt feelings.”⁴⁷ When not engaged in this stance, we treat entities as things to be “managed or handled.” Strawson called this the ‘objective’ stance.

There are plausibly numerous differences between the participatory or second-personal stance and the objective stance. One important difference is this: the forward-looking actions of the participatory stance presuppose that the entities we address can make decisions. By ‘forward-looking actions,’ I mean actions by which we require some future action from another, including (from Darwall’s list) the actions of orders, requests, claims, demands, contracts, and commands ⁴⁸ We perform these actions towards most adult humans under various circumstances. We also perform these actions towards corporations, states, and occasionally some non-human animals. We do not perform any these actions towards infants, most non-human animals, volcanoes, pills, or steel.

Now, such interpersonal addresses are not all there is to obligations. Some obligations bear upon agents independently of the participatory stance. As Darwall puts it, “[s]omeone could acknowledge a norm or reason ... of whatever priority or stringency, without yet acknowledging anyone’s authority to demand that he comply with it.”⁴⁹ How would we third-personally assess (as opposed to second-personally interact with) a creature that couldn’t make decisions, if the creature had the possibility of behaving in accordance with moral obligations

⁴⁷ Strawson 1974, 5.

⁴⁸ My emphasis on the forward-looking aspects of the participatory stance differs from my earlier arguments (Collins 2019), where I focused on backward-looking judgments of ‘moral worth,’ and without use of the second-personal or participatory stance. My earlier use of ‘moral worth’ was more traditionally Kantian, rather than contractualist, in flavour. See also fn. 39.

⁴⁹ Darwall 2006, 10.

of this non-participatory kind? We would judge that the world would go better if the creature were to behave in accordance with the obligation. We would, if we could, try to cause the creature to behave that way, or cause the creature to become such that it could decide to behave that way. But, until the creature could make decisions, we would not attribute it an obligation.

Perhaps those who can't make decisions can bear obligations—they're just excused from complying. This distinction—between being excused from an obligation and not bearing it at all—can be sketched via the notion of a moral remainder. If a creature is excused from complying, there is a moral remainder: agent-regret,⁵⁰ or apologies, or recompense become appropriate forward-looking replacements for the obligation. These replacements for obligations do not arise for creatures that can't make decisions. This suggests these creatures never have the obligation, rather than having it but being excused.

One might object that what matters for obligation-bearing is not decision-making, but motivation-having: to be an obligation-bearer, one must be able to be motivated by obligation. Perhaps non-agentive groups can have obligation-sensitive motivations even if they can't make decisions. In reply, it's true that motivation-having doesn't imply decision-making. Yet I am sceptical of non-agentive groups as motivation-havers, for similar reasons to why I am sceptical of their decisions.⁵¹ I am also sceptical that we would hold the second-personal stance towards (or make obligation-affirming third-personal assessments of) any entity that could be motivated

⁵⁰ Described by Williams 1985.

⁵¹ List (2014) argues against non-agentive groups' beliefs and preferences.

by obligation but could not make decisions accordingly. Decision-making capacities therefore seem crucial to obligation-bearing, even if motivation-having is also crucial.⁵²

Theoretical considerations also suggest obligation-bearing entails decision-making capacities. Consider the function of obligations. Ascriptions of obligations demand that agents do certain things. Their function is to guide action.⁵³ Of course, we don't always know what our obligation demands. Unknown obligations cannot guide conduct. But those who know our unknown obligations can enable the function of obligations when they tell us about our obligation. Thus, our obligations guide not just our conduct, but also the conduct of second parties (those to whom we owe obligations, who might demand them from us) and third parties (those who know our obligations, and might advise us about them). They play this role by prompting our decisions.

The idea that others might demand our obligations, or advise us about them, reiterates the connection between obligations and decision-making. When we demand something from someone, or advise someone to do something, we leave the decision to them. A demander or advisor does not determine her target's decisions. Both presuppose the target can make decisions that are sensitive to the demand or advice. (These practices usually presuppose that

⁵² Björnsson (2014) argues individuals can share an obligation, because the group can have "motivational sensitivities." But he says "the sensitivities required need not be sensitivities of the group, understood as something beyond the ... individuals, as opposed to sensitivities of the individuals within the group." (2014, 117) This makes his proposal more reductionist than the one I'm considering.

⁵³ This is implicit in arguments from political philosophers that 'idealised' theories of justice are not obligation-generating *because* they're not action-guiding (Gheaus 2013, pp. 456-7; Wiens 2014).

the advisee can decide sensitively to *the very obligation being demanded or advised*—but my argument does not assume this.)

To summarise: first, standardly-assumed obligation-bearers can make decisions, while standardly-assumed non-bearers cannot; second, the idea that obligation-bearing entails decision-making capacities captures the action-guiding function of obligations and rationalises their role in demands and advice. But non-agentive groups cannot make decisions. Thus, they cannot bear obligations.

My aim in this half of the paper has been to forward a new argument that non-agentive groups cannot bear obligations. I lack space to consider all possible objections. But before concluding, it's worth considering one alternative proposal. This proposal says members of non-agentive groups can have *joint* obligations, under which there is a one-to-many relation between obligations and bearers—not a one-to-one relation, as in the group-held obligations I have been attacking.⁵⁴ Crucially, joint obligations target *individuals'* decision-making. So, proponents can agree with premises (1) and (2), while asserting that non-agentive groups' members can jointly hold obligations. Proponents of joint obligations can agree with me that obligations entail decision-making capacities—but, the proponent might say, joint obligations can arise via *members'* decision-making capacities.

However, joint obligations of large-scale and clearly non-agentive groups—such as humanity's joint obligation to halt climate change—do not fully respect the sense in which obligations presuppose decision-making capacities. It's true that members of humanity have decision-making capacities. However, the obligation 'to halt climate change' cannot serve the decision-guiding functions I outlined above. If this imperative were blasted from a giant

⁵⁴ Pinkert 2014; Schwenkenbecher 2021. I thank an anonymous reviewer for raising this.

loudspeaker by God or aliens, it could not be an unrefined input into first-personal decision-making. Considering the demand or advice to ‘halt climate change,’ any singular member of humanity will naturally feel the demand overestimates the scope of their agential powers. The content of a massive joint obligation like ‘halt climate change’ is too far-removed from individuals’ decision-making capacities, for the joint obligation approach to be plausible in such large-scale cases. The issue is that the *content* of the purported joint obligation outstrips the decision-making (or reasons-responsive) powers of the members, in large-scale cases. So we cannot posit the joint obligation merely on the basis of members’ decision-making capacities, in such large-scale cases.

To reiterate: my focus has been large-scale groups like ‘humanity’ and ‘emitters.’ In these, I suggest, joint obligations are insufficiently decision-guiding to deserve the name ‘obligation.’ Prominent advocates of joint obligations can agree. Anne Schwenkenbecher argues for joint obligations in morally, epistemically, and cooperatively simple examples.⁵⁵ But, she writes, “ordinary citizens’ obligations to address large-scale moral problems [like poverty] will often be collective only in a very weak sense, in that the collective level has normative primacy in *determining the content of their individual obligations*.”⁵⁶ My argument agrees: in large-scale cases, obligations are individuals’, though group-level facts might determine their content.

However, another joint obligation advocate, Felix Pinkert, writes: “we ... need to ask under which conditions a given presumed joint obligation can apply to a given collection. These conditions, I contend, are simply the conditions under which the agents together are jointly able to do what they supposedly jointly ought to do.” (2014, 190) By arguing for the decision-

⁵⁵ 2021, 92-93.

⁵⁶ 2021, 152, emphasis added.

guiding function of obligations in this section, combined with my earlier account of group abilities, I have disputed Pinkert's contention that joint obligation depends only on joint ability. It also depends, I suggest, on the obligation's decision-guiding propensities. Humanity's obligation to halt climate change lacks adequate decision-guiding propensities, because the decision-making powers dwell far below the level at which the obligation is posited. (By contrast, in two-person cases, the decision-makers are not so completely swamped by the scale of the obligation-attributed group; in such cases, joint obligations may be more plausible.)

Conclusion

Ability and obligation come apart in their applicability to large-scale non-agentive groups like 'humanity' or 'carbon emitters.' One lesson is that we shouldn't be surprised if the same is true of other agentive and normative concepts. Another lesson is that ability and obligation are contrastive in a hitherto unrecognised way: one is about agential explanation; the other is about decision-guidance. The same can be hypothesised about agential versus normative concepts more generally, though this requires further investigation.

More concretely, humanity is able to end some global ills, but humanity lacks an obligation to end them. The most we can say—and what we should say—is that each agent has an obligation that derives from those ills. Questions remain: is humanity *blameworthy* for climate change? Can humanity *act* on climate change? And what about smaller-scale non-agentive groups, in which the content of a purported group obligation might be able to guide members' decisions? Answers to these questions are left open by my arguments.

References

- Aas, Sean. 2015. "Distributing Collective Obligation." *Journal of Ethics and Social Philosophy* 9(3), 1-22.
- Björnsson, Gunnar. 2014. "Essentially Shared Obligations." *Midwest Studies in Philosophy* 38(1), 103-120.
- Björnsson, Gunnar. 2020. "Group Duties Without Decision-Making Procedures." *Journal of Social Ontology* 6(1).
- Blomberg, Olle. 2020. "What We Ought to Do: The Decisions and Duties of Non-agential Groups." *Journal of Social Ontology* 6(1).
- Bratman, Michael. 2014. *Shared Agency: A Planning Theory of Acting Together*. Oxford University Press.
- Brown, Mark A. 1988. "On the Logic of Ability." *Journal of Philosophical Logic* 17(1), 1-26.
- Collins, Stephanie. 2019. *Group Duties: Their Existence and Their Implications for Individuals*. Oxford University Press.
- Cripps, Elizabeth. 2013. *Climate Change and the Moral Agent: Individual Obligations in an Interdependent World*. Oxford University Press.
- Cross, Charles B. 1986. "'Can' and the Logic of Ability." *Philosophical Studies* 50, 53-64.
- Darwall, Stephen. 2006. *The Second-Person Standpoint: Morality, Respect, and Accountability*.
- Dietz, Alexander. 2016. "What We Together Ought to Do." *Ethics* 125(4): 955-982.
- Elder-Vass, David. 2007. "For Emergence: Refining Archer's Account of Social Structure." *Journal of the Theory of Social Behaviour* 37(1), 25-44.
- Fara, Michael. 2008. "Masked Abilities and Compatibilism." *Mind* 117(468), 843-865.
- Feinberg, Joel. 1968. "Collective Responsibility." *Journal of Philosophy* 65, 674-688.
- French, Peter. 1984. *Collective and Corporate Responsibility*. Columbia University Press.
- Gilbert, Margaret. 1989. *On Social Facts*.

- Hayek, Friedrich A. 1969. *Studies in Philosophy, Politics and Economics*. Touchstone.
- Held, Virginia. 1970. "Can a Random Collective Be Morally Responsible?" *Journal of Philosophy* 67, 471-481.
- Isaacs, Tracy. 2011. *Moral Responsibility in Collective Contexts*. Oxford University Press.
- Jackson, Frank. 1987. "Group Morality." In J.J.C. Smart, Philip Pettit, Richard Sylvan, and Jean Norman (eds), *Metaphysics and Morality: Essays in Honour of J.C.C. Smart*. Blackwell.
- Kant, Immanuel. 1785. *Groundwork of the Metaphysics of Morals*.
- Lawford-Smith, Holly. 2015. "What 'We'?" *Journal of Social Ontology* 1(2), 225-9.
- List, Christian and Kai Spiekermann. 2013. "Methodological Individualism and Holism in Political Science: A Reconciliation." *American Political Science Review* 107(4), 629–643.
- List, Christian and Philip Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press.
- List, Christian. 2014. "Three Kinds of Collective Attitudes." *Erkenntnis* 79(9 supp), 1601-1622.
- Lord, Errol. 2017. "On the Intellectual Conditions for Responsibility: Acting for the Right Reasons, Conceptualization, and Credit." *Philosophy and Phenomenological Research* 94(2), 436.
- Maier, John. 2014. "Abilities." *Stanford Encyclopedia of Philosophy* Spring 2014 edition.
- Maier, John. 2018. "Ability, Modality, and Genericity." *Philosophical Studies* 175(2), 411-428.
- Mandelkern, Matthew, Ginger Schultheis, and David Boylan. 2017. "Agentive Modals." *Philosophical Review* 126(3), 301-343.
- Markovitz, Julia. 2010. "Acting for the Right Reasons." *Philosophical Review* 119(2).
- Mathiesen, Kay. 2006. "The Epistemic Features of Group Belief." *Episteme* 2(3), 161-175.
- May, Larry. 1987. *The Morality of Groups*. University of Notre Dame Press.

- Mele, Alfred R. 2003. "Agents' Abilities." *Nous* 37(3), 447-470.
- Miller, Seumas. 2001. *Social Action: A Teleological Account*. Cambridge University Press.
- Miller, Seumas. 2015. "Joint Political Rights and Obligations." *Phenomenology and Mind* 9, 138-146.
- Miller, Seumas. 2020. "Joint Abilities, Joint Know-How and Collective Knowledge." *Social Epistemology* 34(3), 197-212.
- Núñez, Carlos. 2019. "Intending Recalcitrant Social Ends." *Erkenntnis*.
- Parfit, Derek. 1984. *Reasons and Persons*. Clarendon Press.
- Pettit, Philip. 2007. "Responsibility Incorporated." *Ethics*.
- Pinkert, Felix. 2014. "What We Together Can (be Required to) Do." *Midwest Studies in Philosophy* 38, 187-202
- Putnam, Hilary. 1967. "Psychological Predicates." In W.H. Capitan and D.D. Merrill (eds.), *Art, Mind, and Religion*. University of Pittsburgh Press, 37-48.
- Rovane, Carol. 1998. *Bounds of Agency: An Essay in Revisionary Metaphysics*. Princeton University Press.
- Sawyer, R. K. 2005. *Social Emergence: Societies as Complex Systems*. Cambridge University Press.
- Schwenkenbecher, Anne. 2014. "Joint Moral Obligations." *Midwest Studies in Philosophy* 38, 58-74.
- Schwenkenbecher, Anne. 2020. "The Epistemology of Group Duties: What We Know and What We Ought to Do." *Journal of Social Ontology* 6(1).
- Schwenkenbecher, Anne. 2021. *Getting Our Act Together: A Theory of Collective Moral Obligations*. Routledge.

- Shapiro, Scott. 2014. "Massively Shared Agency." In Manuel Vargas and Gideon Yaffe (eds), *Rational and Social Agency: The Philosophy of Michael Bratman*. Oxford University Press.
- Smith, Adam. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations*. W. Strahan and T. Cadell.
- Southwood, Nicholas and David Wiens. 2016. "'Actual' Does Not Imply 'Feasible'." *Philosophical Studies* 173(11), 3037-3060.
- Stemplowska, Zofia. 2016. "Feasibility: Individual and Collective." *Social Philosophy and Policy* 33, 273-291.
- Strawson, Peter F. 1974. "Freedom and Resentment." Reprinted in P.F. Strawson, *Freedom and Resentment and Other Essays*. Taylor and Francis, 2008.
- Tuomela, Raimo. 2007. *The Philosophy of Sociality: The Shared Point of View*. Oxford University Press.
- van Inwagen, Peter. 1983. *An Essay on Free Will*. Oxford University Press.
- Velleman, David. 1997. "How to Share an Intention." *Philosophy and Phenomenological Research* 57(1), 29-50.
- Vetter, Barbara. 2015. *Potentiality: From Dispositions to Modality*. Oxford University Press.
- Vihvelin, Kadri. 2004. "Free Will Demystified: A dispositional account." *Philosophical Topics* 32(1/2), 427-450.
- Williams, Bernard. 1985. "Moral Luck." In Bernard Williams, *Moral Luck*. Cambridge University Press, 20-40.
- Wright, Richard W. 1985. "Causation in Tort Law." *California Law Review* 73, 1735-1828.
- Wringe, Bill. 2005. 'Needs, Rights, and Collective Obligations', *Royal Institute of Philosophy Supplement*, 57, 187-208.

Wringe, Bill. 2016. 'Collective Obligations: Their Existence, Their Explanatory Power, and Their Supervenience on the Obligations of Individuals.' *European Journal of Philosophy* 24(2).

Wringe, Bill. 2020a. 'Global Obligations, Collective Capacities, and 'Ought Implies Can.''' *Philosophical Studies* 177, 1523-1538.

Wringe, Bill. 2020b. "The Duties of Non-Agential Groups: Some Comments on Stephanie Collins' Group Duties." *Journal of Social Ontology* 6(1).