

Secondary Publication



Schulz, Moritz

So What's My Part? : Collective Duties, Individual Contributions, and Distributive Justice

Date of secondary publication: 05.07.2023

Version of Record (Published Version), Article

Persistent identifier: urn:nbn:de:bvb:473-irb-599690

Primary publication

Schulz, Moritz: So What's My Part? : Collective Duties, Individual Contributions, and Distributive Justice. In: Historical Social Research. 48 (2023), 3, S. 320-349. DOI: 10.12759/hsr.48.2023.36.

Legal Notice

This work is protected by copyright and/or the indication of a licence. You are free to use this work in any way permitted by the copyright and/or the licence that applies to your usage. For other uses, you must obtain permission from the rights-holder(s).

This document is made available under a Creative Commons license.



The license information is available online:

<https://creativecommons.org/licenses/by-nc/4.0/legalcode>

So What's My Part? Collective Duties, Individual Contributions, and Distributive Justice

Moritz A. Schulz *

Abstract: »Und was ist meine Rolle? Kollektive Pflichten, individuelle Beiträge und Verteilungsgerechtigkeit«. Problems in normative ethics paradigmatically concern what it is obligatory or permissible for an individual to do. Yet sometimes, each of us ought to do something *individually* in virtue of what we ought to do *together*. Unfortunately, traversing these two different levels at which a moral obligation can arise – individual and collective – is fraught with difficulties that easily lure us into conclusions muddying our understanding of collective obligations. This paper seeks to clearly lay out a systematic problem central to the relation between collective and individual duties in view of alleviating some such confusion and facilitating more concerted attempts at resolving it: collective duties require individuals to act in order for the collective duty to be fulfilled. Yet typically, a collective duty does not entail any one set of individual duties that would prescribe such contributory actions: the work we need to do in order to fulfil a duty can be allocated differently to the individuals collectively bearing it. Much like in matters of distributive justice, then, deriving individual duties from a collective one requires us to employ a distributive scheme, which raises separate normative concerns that have rarely taken centre stage in the debate so far.

Keywords: Collective duties, collective obligations, distributive justice.

1. Introduction

If there are any moral duties at all, there certainly are ones we bear individually. Plausibly, there are also ones we bear collectively.¹ Consider moral claims such as these:

* Moritz A. Schulz, Otto-Friedrich-Universität Bamberg & Friedrich-Alexander-Universität Erlangen-Nürnberg, Lehrstuhl für Politische Theorie, Otto-Friedrich-Universität Bamberg, 96045 Bamberg, Germany; moritz.schulz@uni-bamberg.de, ORCID iD: 0000-0003-2833-4451.

¹ This paper will be of interest only to those prepared to countenance some type of genuinely collective duty or other. Beyond this minimal presumption and although my examples will lean towards accounts permissive of joint duties, I remain agnostic as to which varieties of apparently collective duties exist in fact and under which conditions they obtain – an issue hotly contested in the collective morality literature, which, for instance, is the subject of Maïke

- (1) It is the duty of white people to eliminate structural racism.
- (2) As our university's student Climate Society, we have a duty to organise a protest pressing for fossil fuel divestment.

This paper aims to clearly lay out a systematic problem facing accounts of collective obligations – be they addressed to a collective agent (as in (2)) or some form of a plurality of agents (as in (1)). One way of framing this problem is by analogy to claims like the following:

- (3) Together, all the guests at the birthday party have a claim to the whole birthday cake.
- (4) All taxpayers have to pay enough taxes to fund our state's public expenditures.

These likewise make a normative claim at the collective level. As with (1) and (2), the problem with (3) and (4) is not that they would be vacuous. After all, it is conceivable that we have to reserve some cake for an absent friend, or that the state has some other sources of revenue – just as it is conceivable that only some white people (perhaps those implicated in reproducing racial prejudice) have such remedial duties, or that it is permissible but not required for the Climate Society to stage a protest. So these claims do tell us *something*. And yet, (3) and (4) may strike us as uninformative given how plainly they do not answer our ultimate practical concerns in the matter: how much cake do I get? And how much tax do I have to pay? Much the same, or so I will argue, goes for (1) and (2): knowing what we ought to do *collectively*, we still need to ask what it is that *any one of us* ought to do in virtue of it. Yet answering this latter question is not trivial for, again, it raises much of the same problem as going from the collective to the individual level in (3) and (4): it poses what is fundamentally a problem of distributive justice, as we must allocate contributory duties to individuals, such that jointly they lead to the collective duty being fulfilled. So in individualising collective duties, we necessarily rely on additional normative presumptions about *how* this distribution ought to be made – what I will dub *distributive schemes*.

Perhaps, this problem seems too trivial to warrant extensive discussion. Yet keeping it squarely and clearly in view matters for at least two reasons: firstly, systematically identifying it allows us to compare more transparently how different existing accounts address the problem (and to note those that do not address it at all). Even though the aim of this paper is not a critique of existing accounts of collective duties, this may lead us to re-evaluate their respective plausibility or make us see some of them as incomplete (as section 4.1 will briefly illustrate). Secondly, stripping the problem down to its bare bones

Albertzart's (2023) contribution to the present issue. By the same token, collective duties in the sense discussed here will certainly include, but will likely not be limited to, duties borne by *non-hierarchical group actors* of the kind outlined in the introduction to this special issue (Gehring and Marx 2023).

alerts us to just how structurally similar it is to the concerns of the firmly separated literature on distributive justice. For the most part, I here will remain agnostic as to whether this is because the two really are instances of one and the same problem or whether they merely look alike in important respects. Either way, or so sections 4.2 and 4.3 will argue, their similarity presents theorists of collective morality with fruitful resources barely tapped thus far.

In both ways, then, this article hopes to provide some orientation in the recently prolific debate on collective obligations: proposing a way of helpfully sorting through existing accounts, and an agenda for further developing them. Yet note that even beyond the concerns of philosophers studying collective duties, pinpointing normative assumptions involved in dealing with them should be of help to explanatory undertakings in empirical research too, insofar as it allows for reconsidering what contestation of collective moral claims may spring from. When members of a group fight over what together they ought to do or take steps to avoid fighting over it, we should keep in view the different problems that may motivate them to act as they do.

For the most part, this paper will be concerned with spelling out the above in less sketchy terms. I am going to proceed in three rough steps, concerned in turn with why collective-level and individual-level obligations differ and why each matter (section 2); what exactly transitioning from the collective level to the individual level entails (section 3); and what the upshot of this is for the debate on collective obligations (section 4).

2. Relations between Collective and Individual Duties

I have just alluded in somewhat intuitive terms to a difference between individual-level and collective-level moral claims. The problem I aim to draw out arises in the transition between these levels. So it is first advisable to get a clear picture of what they consist in (section 2.1), and why we cannot simply leave out one of them and conduct the business of collective morality in terms of the other (sections 2.2 and 2.3). In addressing one side of the latter question, section 2.3 will then also prepare the ground for the problem central to this article: knowing a collective duty is not enough, we need to know about individual contributory duties as well.

2.1 Why Does the Difference between Collective and Individual Duties Matter?

In loosely speaking of moral obligations or duties, I here mean to refer simply to moral claims that designate options in practical deliberation – actions or,

indirectly, outcomes aimed at in acting – as morally obligatory or impermissible² choices. They look, at bottom, like this:³

- (5) It is obligatory for A to ϕ in circumstances C.

Many of the obligations we ordinarily think about are borne by individuals, which we might pick out individually or quantify over. For example:

- (6) It is impermissible for anyone to kill an innocent bystander (in all circumstances).
- (7) It is obligatory for Ahmed to save Ben from drowning in this shallow pond.

For present purposes, I will use the term “collective” broadly to denote any entity apt to bear duties and constituted by agents.⁴ A collective obligation, then, is simply one in which the position of the subject A of such a moral claim is occupied by a collective. Let us, for instance, adopt our earlier examples:

- (1*) It is obligatory for white people to eliminate structural racism.
- (2*) It is obligatory for the Climate Society to organise a protest pressing for fossil fuel divestment.

So far, the gulf between collective and individual duties seems to be purely grammatical. The interesting problem, though, arises not in comparing, say, (2*) to (7), but in how (2*) relates to duties like these:

- (2a) It is obligatory for Aisha to obtain relevant permissions for the protest from the authorities.
- (2b) It is obligatory for Betty to start a social media campaign promoting the protest.
- (2c) It is obligatory for Carlos to paint a large banner to be used in the protest.

Collectives, as I said, are constituted by agents. And when a collective bears a duty, this has moral implications for its members too. The difference we are here concerned with, then, is that between stating what a collective has a duty to do (2*), on the one hand, and what the agents constituting the collective have a duty to do in virtue thereof (2a–2c), on the other. Note already that the

² I will speak of obligations only since ϕ -ing is *impermissible* for A iff it is obligatory for A not to ϕ . It may also be *optional* for A iff it is neither obligatory for A to ϕ nor obligatory for A not to ϕ . But as the latter demands nothing of its subject, it is of no import here.

³ Some authors (prominently, Williams 1985) have contested the assumption that principles of this sort are at the heart of ethics. On the whole, this has remained the rebellion against a majority view that it was back then.

⁴ Unfortunately, a unified terminology is yet to emerge in the field. My use of “collective” follows the lead of Wringer (2016) as the most encompassing category including any type of *group* or *collective agent* – i.e., for instance, “combinations,” “coalitions,” and “collectives” in the terms of Collins (2019), both structured and unstructured groups in the terms of Ritchie (2015), or both “agential” and “non-agential groups” (inter alia: Chant 2015; Haan 2021; Blomberg 2020).

individual duties derived from a collective duty typically do not at all have the same content as the collective one. So the switch from one to another is not merely a grammatical one in the kind of subject of which we predicate a duty, but also in its content. (We will return to this point in sections 2.2 and 3.1.1.) The difference between collective- and individual-level duties matter, then, simply because the two can be quite different even where they provide a moral account of one and the same situation.

Let me add a proviso on terminology: admittedly, the difference at stake would be captured more accurately by speaking of collective and *contributory* duties at a *member* level. For note also that the agents constituting a collective may in turn have collective agents among their ranks, not just individuals (paradigmatically, think of international organisations having states as their members or of a group of shareholders comprising both individual investors and companies).⁵ Framing the issue in terms of a juxtaposition of *collective* and *individual* duty-bearers and duties, which I will by and large stick with here, simply has the merit of aligning with ordinary use as it is suggested by standard cases: when we think about what each of us ought to do in light of climate change, for instance, I assume it is most natural to phrase this as a question concerning our “individual duty” or responsibility⁶ in light of climate change.

2.2 Why Do Collective Duties Matter?

Perhaps this might lead one to think of collective moral claims as, at best, “second-class citizens” within the realm of moral knowledge: an imprecise description we can dispose of once we can ascertain all the individual duties obtaining in the circumstances. This, however, would fail to recognise the role collective duties play for individual duties derived from them: knowing the former functions not only as a handy epistemic steppingstone towards knowing the latter. Rather, the latter obtain *in virtue of* the former – that is, individual contributory duties metaphysically depend on a collective duty in the sense that the fact that a collective duty obtains *grounds* the fact that the corresponding individual duties obtain (Wringe 2016). So collective duties matter because many individual moral duties would not obtain *but for them*.

⁵ In such cases, the distributive problem I am concerned with in this paper will arise first in transitioning from the obligated collective to its members. Where a member is a collective, the contributory duty thus arrived at will in turn be a collective duty, which again needs to be translated into contributory duties on behalf of the members of that collective member – and so on, until we arrive at a specification of contributory duties that features only individual agents as duty-bearers.

⁶ I largely avoid the term “responsibility” here for its well-known ambiguity between forward-looking moral responsibility, which consists in duties, and backward-looking moral responsibility, which is concerned with blameworthiness for wrongdoing; let alone causal responsibility or, in some contexts, liability. While there are, of course, strong links between the first two in particular, they warrant separate discussion (Isaacs 2014; Collins 2019, 7–9).

This distinguishes a duty like (7) from a duty like (2a): Aisha has to liaise with authorities only because she is a member of the student Climate Society, and the Climate Society ought to organise a protest. Her mother, for instance, has no such duty.

Yet while overlooking this point will hardly lead us to mistakenly assign a duty to Aisha's mother that she does not bear, it can easily lead us to wrongly reject duties where the actions they prescribe would be futile or otherwise morally irrelevant outside of a collective pattern of action of which they form part. Consider again the example of climate change: arguably, the lifetime carbon footprint of any one individual makes no difference whatsoever for global climate change. Framed as individually grounded, then, surely we must reject any purported duty to reduce our individual greenhouse gas emissions: emission reduction is purportedly obligatory *because of* its effect on harmful climate change, and if my reduction in emissions has no discernible effect on climate change, then there is nothing to justify a duty for me to reduce my emissions.⁷ (In fact, this is an argument influentially defended by Walter Sinnott-Armstrong [2005].) Yet that is not how claims about such duties should be construed: plausibly, they are in the business of spelling out what *our collective* duty to mitigate climate change amounts to *for each of us individually*. Accordingly, they must be assessed according to the standards applying to contributory duties deriving from a collective one, and not according to those applying to stand-alone individual duties.

So even if we found ourselves with all individual duties in fact obtaining in a situation laid out before us, we would still need to inquire into which collective duties obtain in order to assess the plausibility of the individual duties that derive from them. We could not justify them otherwise.

2.3 Why Do Individual Duties Matter?

Conversely, one may doubt that individual contributory duties matter for collective morality – that is, one might think that upon ascertaining a collective obligation, there is no further practical question to be settled. Doubting this may take two forms: firstly, one may doubt whether there is any need to derive individual duties from a collective one at all. And secondly, one may doubt whether thus deriving them, albeit necessary, raises any moral questions not already settled upon arriving at the collective duty. This section will address the former concern, whereas the latter will be the subject of section 3.1.

⁷ This practically overlaps with but is theoretically distinct from the problem discussed under the rubric of aggregative harm (Parfit 1987, secs 27-9; Kagan 2011; Kahn 2014; Spiekermann 2014). Or so, at any rate, it would seem barring a closer examination of the scope and nature of collective *prohibitions*.

Recall that in the introduction I appealed to our intuition that collective-level distributive claims do not settle our practical concerns in matters of distributive justice and that much the same would go for collective duties. But perhaps you resist this intuition. So why exactly is there a need to individualise collective duties?

As agents, we face practical questions: what is it that we should do in circumstances C? Moral obligations of the form sketched in section 2.1 provide us with (moral) reasons for practical conclusions answering such questions.⁸ In so doing, they fulfil what is plausibly the essential function of moral norms: governing our conduct. To govern our conduct, moral obligations must, as is widely assumed (H. Smith 2012) be action-guiding – specifically, in the following weak sense: they must provide practical conclusions for those who must act in order for the obligation to be fulfilled; such that if these conclusions get to be acted upon successfully, the respective obligation will be fulfilled.⁹

Now, I take it that the fulfilment of collective duties depends on individual actions at least in the weak sense that, necessarily, if the collective duty is fulfilled, some facts about actions of members of the collective at least partially ground the fact that it is fulfilled.¹⁰ Subscribing to this does not bar us from attributing certain actions or beliefs or desires to collective agents over and above their members, or from postulating genuinely joint actions and abilities, or from holding that facts about many things other than actions of individuals might figure in grounding facts about collective or joint actions. It is merely to say that *some individual actions* necessarily go into anything a collective does.

It should be uncontroversial that the performance of a joint ability thus depends on what individuals do individually. The interesting bit is that if I lift the left end of the sofa and you lift the right end, together we can do something more, namely carry the sofa. But clearly, our carrying the sofa together

⁸ Typically, obligations do not directly present practical conclusions because they include conditions of applicability that require subsuming the present situation under those governed by the obligation. Pro tanto obligations further complicate the picture, as they may require the agent to weigh up independently applicable moral reasons against one another. All the same, the following toy version of a practical syllogism is sufficient for illustration: in circumstances C, you must ϕ (*normative premise*). Circumstances C obtain (*empirical premise*). Therefore, you must ϕ (*practical conclusion*). Of course, this is supposed to be an *analysis* of practical reasoning and not a description of what is explicitly on an agent's mind upon every instance of acting.

⁹ Somewhat more precisely, the picture is supposed to be this: look at the grounding conditions of the fact that the collective duty is fulfilled. Among these, pick out all facts concerning actions. The agents performing these actions must be supplied with corresponding practical conclusions in order for the collective duty to be action-guiding. (Readers unfamiliar with such metaphysical jargon will be helped out by Epstein [2015, chap. 5].)

¹⁰ In this formulation, I adopt Epstein's (2015, 107) notion of *metaphysical dependence*. A discussion of how this relates to the more common notion of *supervenience* is supplied there (cf. 2015, 110f). For present purposes, however, I have no stakes in how the purported difference pans out.

depends on each of us doing something that can perfectly well be described as an individual action: lifting one side of a sofa.

The same goes for collective agents. Again, the interesting bit is that there may be actions (or beliefs or desires) that are attributable only to the collective agent and not to its individual members – List and Pettit (2011) prominently aim to vindicate this in their discussion of doctrinal paradoxes.¹¹ Yet even where its decision procedure leads a collective agent to hold beliefs that no member holds individually, it is still true that the collective holds this belief only because of what the members contributed to that procedure. That is, even if there are facts about a collective agent that are not facts about any individual agent, it should still be uncontroversial that the collective facts would not obtain if not for some individual facts or others.

But if members (that is, at least one member) must do something (construed broadly, including omissions) in order for the collective to do something (such as fulfilling a collective duty), then an action-guiding collective obligation must guide members in their individual contributory actions. This is what individual contributory duties do.

To recap, the argument from action-guidingness I propose for the need to derive individual contributory duties from collective duties runs like this:

- (i) Moral obligations must be action-guiding.
- (ii) Action-guiding moral obligations must provide practical conclusions for those who must act in order for them to be fulfilled.
- (iii) The fulfilment of collective duties depends on individual actions.
- (iv) *Therefore (from i–iii),* action-guiding collective obligations must provide practical conclusions for individuals.
- (v) Individual agents can only directly act on individual practical conclusions.
- (vi) *Therefore (from iv, v),* fulfilling a collective obligation requires providing individual-level practical conclusions for one or more members of the obligated collective.

This argument draws on the idea that a world in which the business of collective morality consisted only in specifying collective-level obligations would be a world in which individuals would not know what to do – collective morality would fail to govern their conduct. Other authors have pursued different strategies for supporting the view that collective duties always come with

¹¹ In a doctrinal paradox, a collective decision-making procedure settling the collective agent's beliefs has the effect of producing a collective belief that none of the members would endorse individually. For a simple case, imagine that three members agree that $((P \wedge Q) \wedge R) \leftrightarrow S$, yet due to their respective belief sets $\{P, Q, \neg R\}$, $\{P, \neg Q, R\}$, $\{\neg P, Q, R\}$ each conclude individually that $\neg S$. If their collective decision procedure now is to take proposition-wise majority votes on whether P , Q , and R , they collectively conclude that S . On the role such cases have been taken to play for accounts of collective agency, refer also to the introduction of this special issue (Gehring and Marx 2023).

individual member duties. For instance, Stephanie Collins opts to argue the point like this:

Do collectives' duties imply members' duties? We can answer this question by answering a slightly different one: when a collective fails to do any duty, does this entail that a member has failed to do a duty? If so – and assuming all collective duties are such that the collective can fail to do them – then the collective's (violated) duty implies a member's (violated) duty – thus, a collective's duty implies at least one duty for at least one member. (2019, 182)

In supporting the central premise that the violation of a collective duty always implies the violation of some individual duty, Collins (2019, 182-3) then retreats to showing that no successful counterexample to this claim has been produced as of yet, lending provisional inductive support to it. To this end, she rebuts two types of purported counterexamples discussed in the literature on collective responsibility (Pettit 2007; List and Pettit 2011, chap. 2; as well as Jackson 1987; T. H. Smith 2009).

Arguably, such counterexamples would be harmful to the picture of contributory duties I here rely on, too. This is because I defined contributory duties as duties prescribing actions whose performance would lead to a collective's obligation being fulfilled. Surely, then, the non-fulfilment of the collective duty implies the non-fulfilment of at least one contributory duty. If a collective duty can be violated without any member duty being violated, something must be wrong about this picture. However, I concur with Collins's (2019, 184-91) as well as Hindriks's (2009) conclusions on these cases: for either type, we are able to identify individual duty violations in any case (even though it might be that only the collective is *blameworthy*) or should reject the putative collective obligation in the first place because the situation of members rendered it impossible for the collective obligation to be fulfilled (via appeal to *ought implies can*).

So you might wonder: if my argument relies on Collins's rebuttal of such counterexamples going through and if Collins argues for much the same conclusion as I do, why add my argument on top? The answer is simple: the deductive argument above lends stronger support to the conclusion than Collins's inductive one, which infers a universally quantified claim merely from the observation that no counterexamples to it have been identified as of yet. A deductive argument like the above shows why this is so: it is impossible for such counterexamples to exist. That, at any rate, will follow if one grants (i) and its explication (ii). Those denying that moral obligations should be action-guiding in the sense of (ii) may, however, still accept the main point of this section on Collins's grounds.

This, I hope, covers potential resistance to premises (i) through (iii). Perhaps, though, you remain concerned about (v): one might think that *even if* every collective duty requires individuals to act for it to be fulfilled and in this

sense relies on a set of individual contributory duties, we do not need to spell out these contributory duties for the collective duty to be individually action-guiding: why shouldn't we as individual agents, simply be able to act directly on a collective-level practical conclusion?

Taken at face value, I have no better response to this objection than to point out that it simply denies what seems to me to be an obvious difference in practical standpoints, issuing in a blatant grammatical mismatch. Imagine I asked: "What is it that I ought to do about climate change?" And you responded: "We ought to stop climate change." You plainly answered a different question than the one I was posing. A we-answer, as it were, simply does not go into an I-slot. It may be good to know that we ought to do something collectively, just as it may be good to know that some other person ought to do something. But I can immediately act only on what *I ought to do*.

Of course, it may be that it is extremely easy to derive contributory duties from collective duties; such that we do not consciously perform any derivation upon learning the collective duty and perceive ourselves as directly acting on the collective obligation. Yet even if this were so, our case should still analytically be described as acting on an individual obligation issuing in an individual practical conclusion. The next section will be concerned with whether it is in fact as easy as that.

3. Individualising Collective Duties

So far, I have argued that knowing a duty we are under collectively is not enough – we also need to know what each of us ought to do in virtue thereof. This, of course, would not be a major concern if getting from one to the other would be a straightforward matter, such that once we figured out our collective duty, all moral questions proper would have been settled. Unfortunately, it isn't. Section 3.1 will go through four ways in which it might be straightforward and show how they go amiss. Section 3.2 will explore in more detail what is at the bottom of this: distributive schemes.

3.1 How Might Collective Duties Entail Individual Duties?

3.1.1 Collective Duties Semantically Entail Individual Duties

One way in which a collective duty may be all we need to know despite us needing to know about individual duties too would be that the former *semantically entails* the latter. I can see two options for how this may be so – the first and most straightforward one I will discuss here, the second one will be addressed in the next section. First, then, it might be that in collective-level moral claims, we can simply swap the collective for any of its members like so:

(8) Humanity has a duty to limit global temperature increase to 2° C.

(9) I am a human.

(10) Therefore, I have a duty to limit global temperature increase to 2° C.

It is quite easy to see that something must have gone wrong. A highly implausible individual-level moral claim (violating *ought implies can*) has been derived from a perfectly plausible collective-level moral claim (not violating *ought implies can*). Why bother with a proposal that so obviously gets things wrong? Because only keeping clearly in view that arguments like (8–10) are logically flawed prevents us from accidentally reading them as a *reductio* against its collective moral premise (8). Such readings have at times preoccupied theorists in the debate about collective responsibility. While, as I noted in footnote 6, this article is not directly concerned with backward-looking responsibility because it raises somewhat separate theoretical issues, inferences like the following serve well to illustrate our cause for caution:

(11) Citizens of Nazi Germany are collectively responsible for the Holocaust.

(12) Peter, Anton, and Marie are citizens of Nazi Germany.

(13) Therefore, (Peter is responsible for the Holocaust) and (Anton is responsible for the Holocaust) and (Marie is responsible for the Holocaust).

But just like saying that Peter, Anton, and Marie have a collective duty to ϕ does not say that they each have a duty to ϕ individually, so saying that they are collectively responsible for ψ should not be read as an entailing that each one is responsible for ψ . Even if our account of collective responsibility rendered this the correct eventual verdict, it would constitute a separate substantial claim and not a mere entailment of the former.

Surely, if the inference from (11) and (12) to (13) would be valid, doubting the truth of (11) would be a reasonable response. So what exactly is the problem with inferences like these? After all, an argument like the following would clearly be valid:

(14) Citizens of Nazi Germany have parents.

(15) Peter, Anton, and Marie are citizens of Nazi Germany.

(16) Therefore, (Peter has parents) and (Anton has parents) and (Marie has parents).

Note, however, that (14) predicates nothing of Germans *collectively*. It merely quantifies over all individuals that individually have a further property (being a citizen of Nazi Germany) – that is, it is a shorthand for an unwieldy conjunction of individual-level claims like those in (16). This does not seem to be what (8) or (11) do, though. The moral claims made in these are attributed to *the collectives* properly speaking – not just a range of individuals.

Perhaps the notion of attributing something to a collective as distinct from its members may be suspected to smuggle in a contentious form of ontological holism about groups. It does not. Perhaps, one might think this to be a problem peculiar to collective agents that, maybe in virtue of their internal structure, somehow get to be more than the mere sum of their members. It is not. Consider the following inference, which is structurally analogous to (8–10) and (11–13):

- (17) The set of all even numbers has infinitely many members.
- (18) The number 4 is a member of the set of all even numbers.

- (19) Therefore, the number 4 has infinitely many members.

As far as groups are concerned, construing them as mere sets of their members is typically considered the most reductionist analysis on the market (Epstein 2015, 137). Also, mathematical claims about sets like (17) do not seem to have any ontologically contentious whiff of holism about them. And yet, inferences like (17–19) are clearly invalid. That is to say: predicate-preserving individualisation of predications of sets is not truth-preserving. In the same way, arguments inferring an individual duty like (10) merely from a collective duty like (8) and a premise identifying the individual as a member of the collective (9) are not valid either. Instead, we should construe them as relying on a further suppressed premise detailing *how* the collective-level claim should be individualised. This is what I will label a *distributive scheme*. It is the kind of premise rather obviously missing in an argument like this:

- (3) Together, all the guests at the birthday party have a claim to the whole birthday cake.
- (20) Fernando is a guest at the birthday party.
- (21) There are ten guests at the birthday party in total.

- (22) Therefore, Fernando has a claim to a tenth of the birthday cake.

What is missing is something like this:

- (23) The birthday cake should be divided equally among all claimants, i.e., members of the collective identified in (3).

Now, even if we grant that assumptions functionally akin to (23) are technically required to individualise collective obligations, it might still be that they present no obstacle to deriving individual duties from collective ones because distributive schemes are very easy to come by. I will briefly consider three ways in which this might be so in the next sections before we will take a closer look at distributive schemes in section 3.2.

3.1.2 Collective Duties Semantically Entail Distributive Schemes

The first option amounts to this: perhaps we need to add distributive schemes as an additional premise to represent individualising derivations the way that I did above. Yet what this distributive premise consists in is in turn just

entailed by our best understanding of what collective obligations are and how they work, such that once we know that there is a collective obligation and we know how such obligations work, no further distributive question remains.

This, of course, is a logical possibility. But clearly, the burden of proof would here be placed squarely on the proponent of a notion of collective obligations that has such implications. Given the extent of controversy surrounding our individual duties in many cases of collective obligations in both the philosophical literature on applied issues (such as climate change or global poverty) and extra-philosophical discourse, such an account looks bound to fight an uphill battle against both existing theories and common-sense morality. Uphill battles can be fought. Yet I am not aware of anyone waging this particular one as of yet.

Perhaps one might reply that invoking such controversies overlooks that we can derive *some* individual moral statements uncontroversially, albeit ones of a different kind. Think of how Schwenkenbecher (2021, 19) speaks of an individual's obligation to *we-frame* in their practical deliberations when they recognise a collective course of action to be morally best or of Collins's *responsiveness duties* to "act responsively to the other members with a view to remedying the situation" (2019, 96). Does knowing that a collective is under an obligation not at least allow us to say something to the effect of each member individually being obliged *to see to it that* the collective obligation is fulfilled? This may well be so. Yet any such "meta-level" or procedural duties putatively entailed by the concept of a collective obligation seem to amount simply to saying that we should individualise the collective obligation in the sense that is the subject of this article. "Seeing to the obligation getting fulfilled" amounts to figuring out what needs to be done for it to be fulfilled, and that amounts to figuring out what my (and others') individual contributions ought to be. Such "meta-level" individual moral claims may be true, but they would at any rate leave the problem we are here concerned with untouched.

3.1.3 There Is Only One Possible Distribution

If our concept of collective obligations does not specify the way in which they ought to be individualised, it may still be that while *technically* we need to specify distributive schemes, we *typically* do not need to think about the matter in practice because it is plainly trivial, even though not in the sense of *semantic* entailment. One option for how this might be so is that there really is only one distribution of contributory obligations on the menu.

This may be an idea crossing the mind of someone perusing the literature on collective morality. Scholars illustrating how joint abilities can give rise to collective obligations typically use simple case examples in which actions contributing to the exercise of a distinct joint ability are clearly laid out – such as, for instance, Björnsson's (2014, 103) three lake polluters, Schwenkenbecher's (2021, 30) two hikers lifting a tree, or Collins's (2019, 102)

beach rescuers. Similarly, those parts of the literature influenced by cooperative game theory will often present options on the team's deliberative menu directly in terms of a handy set of individual players' choices – à la “(Cooperate, Cooperate)” and “(Hi, Hi).”

Some authors have gone on to note that there are actual cases involving collective obligations that are far less transparent than the ones we often use as toy examples (albeit legitimately so for the purposes they serve). For instance, Bill Wringe (2016, 486) invokes the shared obligation of parents to provide for the well-being of their child to this effect. In fact, though, there is no need for recourse to any specially opaque class of cases to rebut the present proposal.

First, consider that collective obligations will often be attributed to groups that contain some people who end up having to contribute nothing at all – i.e., they will be *wide joint-necessity* cases in the terms of Schwenkenbecher (2021, 8, 100).¹² Any such cases will obviously allow for multiple distributions of contributory duties: we can shift around who gets to do nothing.

Second, even for *strict joint-necessity* cases, in which “the number of available contributors to a collective outcome equals the number of contributors minimally necessary to produce it” (Schwenkenbecher 2021, 8), cases in which only one distribution of contributory duties is possible should be rare simply in virtue of permutations. Consider again:

- (2) As our university's student Climate Society, we have a duty to organise a protest pressing for fossil fuel divestment.

This collective duty is *prima facie* compatible with multiple sets of contributory duties such as:

- (2.1) = {Aisha has to start a social media campaign;
Betty has to paint a large banner;
Carlos has to contact the relevant authorities; ...}
 - (2.2) = {Aisha has to contact the relevant authorities;
Betty has to start a social media campaign;
Carlos has to paint a large banner; ...}
- Etc.

So even if we know which contributory actions have to be performed by individuals and that everyone has to do something, this still does not tell us what each one needs to do – at least unless there is only one possible pattern of individual actions and the contributory actions therein are all exactly alike, which again is a far-fetched case. In short, there *almost always* will be multiple possible distributions of contributory duties for any one collective duty.

¹² This, at any rate, is to be expected if we follow Collins (2019, 118) in holding that collective duties should *ceteris paribus* be stated generally for reasons of parsimony.

Perhaps there is a looser sense in which saying that only one distribution is possible amounts to saying that among a number of *in principle* possible distributions only one will *in fact* allow for the collective duty to be discharged, or that considering any alternatives to one particular distribution will itself thwart the discharge of the duty. Anne Schwenkenbecher (2021, 103) enlists such cases when, in passing, she discounts distributive questions as a secondary concern within the study of collective morality:

[I]t may be overall permissible to withdraw (or decrease) one's contribution to a collective farewell present for a colleague in case other colleagues are not contributing their fair share. But this move seems less permissible where someone's life is jeopardised by my refusal to shoulder an unfair share.

She certainly has a point here. Yet it is not immediately clear that cases in which questioning a distributive scheme would be morally impermissible on grounds of urgency or feasibility should serve as our paradigm cases of collective obligations. Consider the broad range of familiar concerns that lend themselves to asking a "What is my part in ...?" question: from pressing large-scale challenges such as climate change, global justice, structural discrimination, or populism and polarisation to more local and mundane ones such as advocating for a curriculum reform, keeping one's commitment to writing a paper together, or getting a family's household chores done. Quite a lot of cases in which we know that *we* ought to do something seem to be cases in which the distributive question arises prominently or is amenable to contestation: we do ask it and we are known to fight over it.

3.1.4 Distributive Schemes Are Obvious

For one final stab, recall the example of a matter of distributive justice I have given above:

- (3) Together, all the guests at the birthday party have a claim to the whole birthday cake.
 - (20) Fernando is a guest at the birthday party.
 - (21) There are ten guests at the birthday party in total.
 - (23) The birthday cake should be divided equally among all claimants, i.e., members of the collective identified in (3).
-
- (22) Therefore, Fernando has a claim to a tenth of the birthday cake.

In cases like these, many people will just presume the egalitarian principle (23) to be the obviously correct distributive scheme in ordinary circumstances. Under such a firmly established and uncontested assumption, we may take (3), (20), and (21) to directly imply (22) and just proceed to cut the cake into ten equal pieces after having counted the number of guests. Likewise, it may be that while different distributive schemes for individualising a collective obligation are a real option, it is simply *obvious* which one is the

correct one, such that no interesting business for moral philosophy arises with regard to distributive schemes.

Yet firstly, such cases would primarily explain why distributive schemes may not readily come to our attention in dealing with collective obligations (much like how we might take ourselves to be acting directly on a collective practical conclusion, as I discussed in section 2.3). It is not clear, however, why this should have any immediate bearing on the best analytical representation of our underlying practical reasoning. Secondly, the very fact that this normative premise would be taken to be obviously true without any need for further reflection may, in tried and tested Socratic fashion, arouse our suspicion. Many a valuable contribution of normative ethics has consisted in scrutinising precisely such unquestioned ethical beliefs. Lastly, even though it remains a logical possibility that no interesting questions about the matter remain, widespread disagreement about the individualisation of collective duties again seems to fly in the face of this idea.

3.2 Distributive Schemes

So far, I have tried to show that collective obligations come with a particular moral problem: we need to derive individual duties from them, yet doing so is no straightforward matter. Rather, it requires us to deploy additional normative assumptions: distributive schemes. The present section will take a closer look at how these work. The upshot will be, first, that this problem we have now seen to arise is structurally similar to that arising in transitioning from collective- to individual-level claims in distributive justice (sections 3.2.2–3.2.4) and, second, that it really does constitute a normative issue for all manner of collectives, even though in some it will not surface as a practical question precisely in virtue of the nature of the collective (section 3.2.5).

3.2.1 Distributive Schemes Matter for Theorising

Before we get to this, though, let me add a quick note so the relevance of our subject does not get misjudged. The fact that my argument for the relevance of contributory duties in section 2.3 has focussed on action-guidingness may lead to the impression that a concern with distributive schemes is a purely *practical* one: it is what I need to resolve in order to know what I ought to do in virtue of a duty I bear collectively. And sure, we practically need to know about distributive schemes and our contributory individual obligations for largely the same reasons as and to the same extent that we need to know about our collective obligations. However, addressing distributive schemes is a matter of *theoretical* concern, too.

Generally and trivially, it is good philosophical practice to make as transparent as possible what any given theoretical claim one asserts actually entails. Understanding what a claim means is, after all, not least a matter of

what can be inferred from it. More specifically, though, the implications of collective obligations for individual obligations seem to be particularly relevant since our – despite the controversy inevitably surrounding any claim about philosophical methodology – best and most widely practised method in assessing normative claims in ethics is one version or another of reflective equilibrium (Tersman 2018; Cath 2016; Knight 2017). At its heart, this method involves testing the verdicts of moral principles against our considered judgments in particular cases. Yet to assess a principle stating a collective obligation in this manner, it is vital to understand the full range of its normative implications in a given case, not just the fact that a particular collective will therein bear a corresponding obligation. It is, after all, widely accepted in our practices of studying potential counterexamples to principles in individual morality that *any* counterintuitive implication of a principle counts as pro tanto reason speaking against it, even if remote or unintended.

3.2.2 Distributive Schemes and Distributive Justice

With this caveat in place, let us take a closer look at distributive schemes. In section 3.1.1, I introduced the idea of a distributive scheme by analogy to a problem of distributive justice – specifically, an argument inferring an individual distributive claim from a collective one. This led us to identify something like the following as a missing normative premise:

(23) The birthday cake should be divided equally among all claimants.

So far, my claim that things are the same in individualising collective duties entails only this much: arguments deriving individual from collective duties too must include an additional normative premise functioning as a transformation rule between statements of either kind. True, this would be a structural similarity between collective duties and distributive justice. But it would be a minimal one. There are many uses for transformation rules of some form or other – in the broadest sense, any rule of inference constitutes one. However, the similarities do not stop there. Rather, the next two sections will aim to show that both types of missing premises really are *distributive* schemes.

The basic reason for this is the following: like a pool of distributable resources, that which is prescribed by a collective duty is distributable, too. Given a possible pattern of actions fulfilling the collective duty, the distribution of the individual actions that constitute it is a zero-sum game. Note that this distinguishes the issue of collective obligations from that of (backward-looking) collective responsibility. Plausibly, blame does not work in the same way (e.g., Hess 2018, 147-8; May 1992; Kutz 2000). There is a finite amount of work involved in, say, a specific plot for blackmailing someone. Experienced blackmailer Debbie could carry it out on her own. When she evenly shares the workload with Emma, Debbie needs to do only half of the work. But this

does not mean that in partnering with Emma, Debbie can avoid half of the blame she would have deserved had she done it alone.

3.2.3 Birthday Cake

So let us walk through a classic problem of distributive justice: the birthday cake from (3). What is involved in distributing the cake? What information do we need? And what do we do with it? First, we need to know which individuals have the standing to raise a claim in distributing the cake – that is to say, for whom it is that we need to check what they should get. In our case, (3) in conjunction with (21) specified this to be the ten guests at the birthday party. Equally obviously, we need to know what is to be distributed. This too is settled in the collective claim (3) – it is the cake.

Next, we need what is essentially a function mapping sets of morally relevant properties of potential claimants onto sets of properties that are sufficient conditions¹³ for an allocated share of the distribuendum to be just. For a strict egalitarian principle like (23), there is only one such assignment: the only relevant property of a given claimant is their being a claimant, and the just share of cake for someone possessing this property is a share of cake that is $1/n$ of the whole cake in size (where n is the number of claimants picked out by the former property; in this case, all ten). Things might get more complicated on other principles, though. We could, for instance, be prioritarians about cake slices and say that the hungrier a guest is, the bigger their slice should be. Or we could be humble sufficientarians and say that as long as everyone gets a slice at all, all is well. We might as well have a cake lottery, where a share is a given claimant's just share if it is the one allotted to them through a randomised procedure. You get the point.

We also need to know what distributions of the distribuendum are possible. That is, we need to know all possible sets of n shares of the whole cake, where such shares diverge from an intuitive picture of a slice of cake in two respects: individual shares can be null, such that in the extreme one share is a 100% of the cake and the other nine are shares of 0% each. And shares are individuated only by all just-making properties. Perhaps candles have been put on top of our cake in a fashionably irregular manner, such that we could divide the cake into ten equally sized pieces in such a way that one piece has two candles on it or in such a way that all pieces have at most one candle on them. If only size matters for justice, these will count as the same possible set of shares of the cake, for there is no need to proliferate distributions that do not differ relevantly.

¹³ Talk of a set of properties is a bit ambiguous here: there may both be multiple properties (say, A and B) that *jointly* pick out a just (J) share (y) and multiple ways in which a share can be just (say, A and B, or C) for a given type of claimant (x with property M), providing *individually* sufficient conditions. Perhaps this is clearer represented in a somewhat cumbersome conditional: $\forall x(M(x) \rightarrow \forall y(((A(y) \wedge B(y)) \vee C(y)) \rightarrow J(y)))$.

Finally, we will have to determine the claims set of n sets of just-making properties of a share for each of our claimants.¹⁴ In our case, it is ten times the property of being a slice a tenth the size of the whole cake. We then need to find a possible distribution that contains shares that match the properties specified in the claims set¹⁵ and allocate to each claimant a share according to their claim. In our case, this will be the distribution that divides the cake into ten equally sized pieces.

3.2.4 Student Protest

Now let us see what things look like for a collective obligation – say, the duty of the student Climate Society to organise a protest (2). If we want to determine what each of the Society’s members ought to do in virtue of their collective duty, what is it that we need to know? First, we need to know who the obligated collective consists of – that is, for which individuals we need to find out what their individual contributory duties are. Let us assume that, as such things go, the Society currently consists only of Aisha, Betty, and Carlos.

In deriving individual duties from a collective duty, we also assign something to each individual. On a formal level, both the collective and the individual statements in the case of the birthday cake assigned claims. Likewise, both the collective and the individual statements in the present case assign duties, that is to say, they designate it as obligatory for something to be done. As we saw in section 3.1.1, though, that which gets assigned to the collective is not identical to what gets assigned to the individuals – just as distributing the cake is not a matter of assigning *the cake* to anyone, but of assigning *shares* of the cake.

Now, it seems to be one thing to split up a cake into a number of slices, but what is it to split up a collective duty or a collective action into individual ones? It will be helpful to take a step back from the confusing image of cutting a duty into slices and recall our present example. What we need to assign to the individual members in order to individualise the Society’s collective duty to organise a protest are individual duties whose fulfilment will *make it the case that* the collective obligation to organise the protest has been fulfilled. They are shares of the collective duty in the sense that the fulfilment of each of the members’ contributory duties partially grounds the fulfilment of the collective duty. Let us simplify things a bit and again assume that for the protest to be organised, permissions need to be obtained, a social media campaign needs to be set up, and a large banner needs to be painted. If these things get done, then the protest will have been organised.

¹⁴ Or the set of all such sets if there are multiple individually sufficient conditions for a just share, see fn. 13.

¹⁵ If we get unlucky, we may additionally need a higher-order principle settling which distribution is to be chosen if no distribution is completely just, i.e., if there is no set of shares that contains a just share for every claimant (or, yet more unluckily, several orders of these).

Next, we need to know how the assigning – the matching of members to contributory duties – is to be done. In order to settle this, we again need a way of picking out members and a way of picking out a contributory duty that would be justly theirs, so that we can map types of members onto types of duties. Technically, there are no limitations on how this picking out is achieved as long as we attain a morally justified distributive scheme by means of it. However, heeding the latter proviso will eventually require us to specify the distributive principle in such a way that the way in which a just share is picked out also *explains why* a share thus picked out is just for that kind of person – that is, if the identifying properties also are the right-making ones.

We already know that the set of members from which the individual-picking is to be performed contains Aisha, Betty, and Carlos. But we also need to know what the set which the duty-picking is to be performed from contains. Above, I proposed an intuitive decomposition into three contributory tasks (permits, social media, banner). Yet when we consider such tasks, it is obvious that they may well be carved up further insofar as each consists of a number of steps that could in principle be performed by different people. If, for instance, one of them fell ill during the run-up to the protest, one of the others could pick up their task from where they got to and complete it for them. The composite steps which in such a scenario one could not complete but would necessarily have to start over if they had not been completed are typically quite small indeed.¹⁶

This being so, a great many more divisions of labour are conceivable than that of the three intuitive tasks I have described above. There are, as it were, many “action packages” that can be drawn together from the many things that can be done by members in order to collectively do their duty. So the individual duty-picking must be done from the set of all sets of *n* such packages, where each package is a set of individual duties composed such that the fulfilment of all duties across all packages makes it the case that the collective duty is fulfilled. Again, some packages may be empty and they are individuated only in the ways in which they get picked out by the distributive principle.

Then we need to see which shares our distributive principle assigns to our actual group members based on the properties they in fact have, and from this are given a set of specifications of just individual duties. Compare this against the possible distributions and assign to each an individual duty matching the criteria for a just contribution of theirs from amongst a distribution containing a just such contribution for everyone and you are done already.

Let me add one last caveat: above, I kept loosely speaking of what we *need to know* in order to individualise the respective collective-level claim. This is

¹⁶ Cf. the debate on basic (Danto 1963, 1965) or primitive (Davidson 2001 [1971]) or simple (Martin 1972) actions.

certainly true if our intent is to lay out a full theoretical account of the matter at hand. But do we need to know all that from the practical standpoint of one member concerned with their individual duties? Not necessarily. If we know the distributive scheme, we can sometimes identify a just share without knowing much about the larger pattern of contributory actions at all. More generally, it is clear that the realm of possible distributions will often be vast and we need to explore it only so far until we have found one distribution that matches our claims set. Often, we will also be able to directly read off that relevant distribution from the specification of the claims set – as in our cake example, where just shares are picked out by being a tenth of the size of the whole cake. Yet this may not always be the case. We might in principle specify a just share as one that its recipient agrees to or as one of the size that would result from using an ancient cake cutter worshipped by an obscure sect of pâtissiers.

3.2.5 Organisational Structures and Distributive Schemes

There is another rather prominent way in which the complicated issues we have just spelt out may not arise: they may simply have been settled already. Consider the following case:¹⁷

- (24) *Lifeboat*: On a calm day, a person is spotted some 300 metres off the coast next to a capsized dinghy, struggling to keep their head above the water. Luckily, there is a lifeboat nearby ready to be launched. At least three people are needed to man it and safely rescue the drowning person. Now, consider two alternative scenarios:
- (a) In the first, there is only a number of bystanders around, and no way to alert the lifeboat crew in time. As it happens, they are good swimmers and some of them know how to operate this type of boat.
 - (b) In the second, the lifeboat crew happens to be present and ready to launch.

In either case, a collective is able and presumably obligated to save the drowning person. Several individuals forming either collective need to perform individual actions in order to perform their collective rescue: one steers the boat, two pull the swimmer aboard, etc. Yet by contrast to the random bystanders, the lifeboat crew is unlikely to face a practical question of individualisation in carrying out their collective duty: most likely, each crew member will take their position on board and they will head off with each one performing their established role in the team.

Such pre-existing distributions of roles and responsibilities have recurrently been taken as a hallmark of at least certain collective agents. While some accounts of collective agents require an internal structure only in the weaker sense of any established decision-making procedure (e.g., List and

¹⁷ Collins (2019, 102) uses a case similar to (24a).

Pettit 2011; Collins 2019), it is often observed that many collective agents feature organisational structures involving role distributions. Herlinde Pauer-Studer's (2014) constitutivist account of group agency, for instance, focuses on these; and Katherine Ritchie (2013, 2015, 2020) defends an account of collective agents as structured wholes,¹⁸ in part echoing work by Sally Haslanger (2003, 2012) on positions within social structures with regard to gender and race.¹⁹ Ritchie (2020, 411) observes:

Organized social groups have structures. A group's structure captures its functional organization. For instance, a baseball team's structure captures the functional roles of the catcher, pitcher, outfielders, etc. Relations might include *calling the pitch*, *pitching to*, *returns the ball to*, and so on. Nodes in an organized social group's structure carry requirements on who can occupy them.

In a similar vein, Peter French's (1984) account of corporate responsibility centres on what he terms Corporate Internal Decision (CID) Structures, which comprise a "corporate-decision recognition rule" and an "organisational or responsibility flowchart":

The organization chart of a corporation distinguishes players and clarifies their rank and the interwoven lines of responsibility within the corporation. An organizational chart tells us, for example, that anyone holding the title "Executive vice-president for finance administration" stands in a certain relationship to anyone holding the title "director of internal audit" and to anyone holding the title "treasurer," etc. In effect it expresses, or maps, the interdependent and dependent relationships, line and staff, that are involved in determinations of corporate decisions and actions. (French 1984, 42-3)

Regardless of whether we want to draw the conceptual boundaries of collective agency by invoking such internal structures, it is clear that where they are in place, they will typically have implications for the individualisation of collective obligations: organisational structures of the kind Ritchie or French describe can constrain or indeed fully entail distributive schemes for some, many, or all obligations the collective might be under.

Still, this does not mean that distributive questions in individualising collective duties would not arise in structured collectives. Firstly, even though organisational structures often succeed in settling distributive schemes across a range of actions the collective may take to, vagueness or overlap in the

¹⁸ Where structures are taken to be "complexes, networks, or 'latticeworks' of relations," which "can be represented as (although they are not identical to) graphs composed of nodes and edges. Nodes represent positions or places that can be occupied by objects. Edges represent relations that hold between nodes (or node-occupiers). [...] Nodes are defined in terms of both (a) relations to other nodes and (b) (possibly null) additional requirements on occupiers" (Ritchie 2020, 405).

¹⁹ Ásta (2018, chap. 6) defends an account of social identities (which can constitute non-agential groups such as "Austrians" or "liberals") that relies on a similar notion of occupying a "location on a social map" (a node related to others in particular ways).

specification of responsibilities is likewise commonplace. (Consider, for instance, how responsibility for a cross-cutting policy issue may be assigned to different cabinet ministers.) Secondly and more fundamentally, to the extent that the structure indeed settles the distributive scheme, this simply means that the distributive question gets raised and settled at the point of determining the collective's organisational structure and assigning members to the different nodes in it.²⁰

4. Distributive Schemes and the Collective Morality Literature

Up to this point, I have argued that questions in collective morality can be answered on two levels (section 2) and that traversing them poses a problem (section 3.1), which we have taken a closer look at in section 3.2. The part of my promise I have yet to deliver on is to show that noting all this amounts to more than intellectual stamp collecting. So what is the upshot for the debate on collective obligations? Section 4.1 will point out how keeping the problem of distributive schemes clearly in view can provide orientation in the existing literature, whereas sections 4.2 and 4.3 will highlight two ways in which the structural similarity with matters of distributive justice discussed in 3.2.2–3.2.4 could allow us to develop it further.

4.1 Taking Stock of Proposals

The first thing noting a problem allows you to do is, obviously, seeing whether people have a solution on offer. Of course, not everyone writing on collective morality can be expected to cover each theoretical problem arising within that domain. But if this paper got things roughly right about the relevance of distributive schemes,²¹ then we should expect at least those scholars aiming to paint a rather comprehensive picture of collective obligations to address it. Still, some otherwise very informative accounts have surprisingly little to say about it. Consider, for instance, two recent books on collective obligations.

Stephanie Collins (2019, chap. 7) devotes an entire chapter to membership duties (individual contributory duties in my terms), discussing first whether collective duties always imply such individual duties (see above 2.3) and then what exactly the former entail in terms of the latter – acknowledging that “[w]e need to ask what a member’s ‘part’ or ‘bit’ is” (Collins 2019, 193, similarly 176). She arrives at the following conclusion:

²⁰ These two steps obviously need not coincide. When there is a vacancy for a specified job role within a company, the structure is (largely) settled and what is to be determined is which (new) member gets allocated to this position within it.

²¹ Recall especially sections 2.3, 3.2.1, and 3.2.5.

If a collective has a duty to see to it that X, then

1. Each member has a duty to use their role, if possible and as appropriate, to put inputs into the collective's decision-making procedure with a view to the procedure's distributing roles to members in a way that: if enough members used their roles with a view to seeing to it that X, then that would be sufficient for X in a high proportion of likely futures. These are "X-sufficient" roles.
2. If X-sufficient roles are distributed, then each member has a duty to use their role, if possible and as appropriate, with a view to seeing to it that X. (Collins 2019, 198)

Note that clause (1) calls for establishing *any* distribution of roles sufficient for fulfilling the collective duty, and (2) consequently demands members to act within *those* roles in order to fulfil the collective duty in the way that this distribution of roles foresees. But this merely amounts to a blanket duty endorsing any set of action packages (as I called it in 2.3.4) that the collective comes to endorse. How does it come to endorse one over the other? Are they all morally on a par?

Note that Collins is aware of the fact that there will typically be multiple possible sets of contributory duties: she criticises Bill Wringer's (2016, 488) proposal that members have a duty to perform their part in a pattern of actions fulfilling the collective duty if other members are performing or likely to perform their parts in it on the grounds that this could create contradictory duties for a member who in one pattern of actions ought to ϕ but in another possible pattern ought not to ϕ , if both happened to be equally likely to be acted upon (Collins 2019, 195). She prescribes the determination of *one* distribution of contributory roles in clause (1) above to resolve this rather special problem. Yet all the while, the issue of *which one* ought to be picked (and to thereby become obligatory for members qua clause (2)) simply does not come up.

Anne Schwenkenbecher's *Getting Our Act Together* likewise includes a chapter on "What Collective Obligations Mean for Individual Agents" (2021, chap. 5), in the course of which she comes to raise the question of "how the burden of individual contributory duties should be distributed among group members" (2021, 102). Arguably, this acknowledges what I have here framed as the problem of distributive schemes. Yet Schwenkenbecher immediately goes on to discount it, adding that "the question of burden-distribution, while morally relevant, is subordinate here" because "an unjust distribution of burdens need not undermine a collective obligation." This may be so, yet it still does not tell the individuals involved what they are required to do (cf. section 2.3). She adds a brief reference to Iris Marion Young (2004, 2011), but the remainder of the chapter is devoted to other concerns.

These two publications are noteworthy for present purposes because they explicitly discuss contributory duties at all. Arguably, much of the literature on collective obligations is dialectically pitted against sceptics about the

existence of collective duties and thus primarily concerned with exploring whether and how duties obtain at the collective level. Also, early papers during the notable revival of this debate over the past decade have – as, perhaps, is typical of nascent fields – often attempted to cover a broad range of systematic problems within the realm of one article, at times developing an account of collective agency or joint abilities hand in hand with one of collective obligations or responsibility or both (e.g., Björnsson 2014; Jansen 2014; Pinkert 2014; Mathiesen 2006; Collins 2013). Understandably, such ambition requires sacrificing some theoretical nuance or other. At any rate, then, I do not mean to engage in a blame game in pointing out blank spots – doing so is simply a matter of highlighting where we can do even better going forward.

Note also that even though express treatment of the problem of distributive schemes has often eluded the collective morality literature, we may still piece together substantial proposals implied in discussing nearby issues or in fact occasionally draw more explicit remarks from other accounts. In so doing, having a firm grasp on the problem stands to be helpful in a second way: even though, unfortunately, a full-scale review of this literature is beyond the scope of this paper, it should be clear that asking *which* distributive schemes are in turn endorsed by different accounts can be of use in charting the logical space of possible options and thus prepare the ground for a more targeted and systematic discussion of them.

4.2 Theory Transfer

It is in exploring this space of theoretical options that the structural similarity between distributive schemes in matters of collective duties and matters of distributive justice could prove particularly fruitful. In the context of *discovery*, we may simply tap into the rich distributive justice literature for inspiration: after all, it provides us with a great many distributive principles that we can plug into accounts of collective morality and see what we get – even if nothing that has been said about those principles over there would inform their suitability for our purposes. This might not sound like much, yet when tracing how the debate on distributive justice has panned out since Rawls (1971), we can notice recurrently that seemingly intuitive notions like “equality” or “sufficiency” require a fair bit of unpacking that subsequently enriches our conceptual apparatus. It might not be our worst stab to try some of the fruits of these lengthy discussions.

However, what we can learn from that debate will also extend into the context of *justification*, even if those theories can only be transferred and so do not at heart constitute theories of the same phenomenon. This is simply because any systematic debate also teaches us about logical relations between alternative positions within it and the dialectical moves made possible by them, which can carry over into our target domain of collective obligations if

they do not spring specifically from the theories' application to distributable resources. We may, for instance, learn that when someone proposes that what centrally matters with collective obligations is that each member contributes *enough*, we can ask them how they reconcile what Casal (2007) has termed the negative and the positive thesis implied by such thresholds, and our interlocutor might reply that we could just split them apart (Huseby 2010).

4.3 Theory Unification

By contrast to a structural transfer of theories from one domain to another, a focus on distributive schemes may also allow us to unify separate theories by subsuming the phenomena they target under a broader concept and thus recasting the previously stand-alone theories as particular applications of a more general theory. Two ways in which this may play out in our case come to mind. Firstly, we might ambitiously explore unifying distributive questions concerning goods and obligations altogether, such that the two issues I have here depicted as relevantly similar would be cast as simply instances of one and the same fundamental problem. Individualising collective duties would arguably be much easier if we could cast them in terms of a common morally relevant (right-making) currency in which we can compare different contributory tasks – to wit, some interpretation of the *burdens and benefits* that are imbued by assigning them as duties. Any candidate for such an abstract currency is, however, likely not to be far removed from what matters morally about resources, and some degree of translatability is indeed widely presupposed by our established practices of materially compensating others for their contributions to collective actions. Such unification would also allow us to widen our perspective from a given collective obligation at hand to the broader field of moral demands addressed to a given agent in virtue of independently obtaining collective obligations.

On a humbler yet theoretically significant level, we may leverage a potential for unification in the relation between normative and applied theories of collective morality. In contrast to the literature on collective duties and responsibility in general, distributive schemes have more often seen explicit discussion in debates around applied issues, where the practical questions outlined in this paper often surface at centre stage. Consider the relatively extant debate on individual duties in countering climate change, the analogous debate about structural injustice, duties of resistance against unjust policies and institutions, revisionists in just war theory discussing the duties and responsibility of individual soldiers, individual duties of aid in view of global poverty, or individual epistemic duties of citizens within epistemic theories of democracy. In some of these debates, explicit discussion of distributive principles relating to the burdens of collective obligations has gained a prominent place – such as in the set of principles of climate justice discussed by Caney (2010,

2012) following Shue (1999), which incidentally have been taken up under a different guise in the debate about which states specifically have a duty to humanitarian intervention when such intervention is morally required (Hjorthen 2020).

A closer analysis of distributive schemes and a normative defence of certain particular principles for the individualisation of different types of collective duties (if not for all) should lend us a framework theory that can cover different applied issues and highlight their normative structural similarities, explain in a unified theoretical vocabulary the way in which certain facts about these specific cases matter morally and rid them of other contingent features that might distort our normative analysis of them.

5. Conclusion

With plenty more work in sight, the time has come to conclude. The main message of this paper has been that collective obligations structurally resemble collective distributive claims in that both require the deployment of a separate distributive scheme to satisfy the central practical interest we take in either type of normative claim. While the individualisation of distributive claims trivially is the point and purpose of discussing distributive justice, I have argued it is likewise required in discussing collective obligations. This is because the fulfilment of collective obligations depends on actions of individuals, and for such obligations to be action-guiding, we must tell individuals what their part consists in. Doing so, however, raises a further normative question that is not typically settled upon ascertaining the collective obligation itself. I have gone on to provide a closer analysis of what is involved in individualisation in either case, underlining the structural similarity of the operation for either type of claim. Having this distributive question clearly in view helps us identify more readily what existent accounts have to offer in response to it – and exploiting said structural similarities between the two types of collective-level claims should allow us to recruit insights from the rich literature on distributive justice. Both should further our efforts to better understand precisely how distributive schemes play out in collective morality and, importantly, which substantive schemes we ought to endorse.

References

Albertzart, Maike. 2023. Being Jointly Obligated: A Reductive Account. *Historical Social Research* 48 (3): 296-319. doi: [10.12759/hsr.48.2023.35](https://doi.org/10.12759/hsr.48.2023.35).

- Ásta. 2018. *Categories We Live By: The Construction of Sex, Gender, Race, and Other Social Categories*. Oxford: Oxford University Press. doi: [10.1093/oso/9780190256791.001.0001](https://doi.org/10.1093/oso/9780190256791.001.0001).
- Björnsson, Gunnar. 2014. Essentially Shared Obligations. *Midwest Studies In Philosophy* 38 (1): 103-20. doi: [10.1111/misp.12019](https://doi.org/10.1111/misp.12019).
- Blomberg, Olle. 2020. What We Ought to Do: The Decisions and Duties of Non-Agential Groups. *Journal of Social Ontology* 6 (1): 101-16. doi: [10.1515/jso-2020-0025](https://doi.org/10.1515/jso-2020-0025).
- Caney, Simon. 2010. Climate Change and the Duties of the Advantaged. *Critical Review of International Social and Political Philosophy* 13 (1): 203-28. doi: [10.1080/13698230903326331](https://doi.org/10.1080/13698230903326331).
- Caney, Simon. 2012. Just Emissions. *Philosophy & Public Affairs* 40 (4): 255-300. doi: [10.1111/papa.12005](https://doi.org/10.1111/papa.12005).
- Casal, Paula. 2007. Why Sufficiency Is Not Enough. *Ethics* 117 (2): 296-326. doi: [10.1086/510692](https://doi.org/10.1086/510692).
- Cath, Yuri. 2016. Reflective Equilibrium. In *Oxford Handbook of Philosophical Methodology*, ed. Herman Cappelen, Tamar Szabó Gendler and John Hawthorne. Oxford: Oxford University Press. doi: [10.1093/oxfordhb/9780199668779.013.32](https://doi.org/10.1093/oxfordhb/9780199668779.013.32).
- Chant, Sara Rachel. 2015. Collective Responsibility in a Hollywood Standoff. *Thought: A Journal of Philosophy* 4 (2): 83-92. doi: [10.1002/tht3.161](https://doi.org/10.1002/tht3.161).
- Collins, Stephanie. 2013. Collectives' Duties and Collectivization Duties. *Australasian Journal of Philosophy* 91 (2): 231-48. doi: [10.1080/00048402.2012.717533](https://doi.org/10.1080/00048402.2012.717533).
- Collins, Stephanie. 2019. *Group Duties: Their Existence and Their Implications for Individuals*. Oxford: Oxford University Press. doi: [10.1093/oso/9780198840275.001.0001](https://doi.org/10.1093/oso/9780198840275.001.0001).
- Danto, Arthur C. 1963. What We Can Do. *Journal of Philosophy* 60 (15): 435-45. doi: [10.2307/2023429](https://doi.org/10.2307/2023429).
- Danto, Arthur C. 1965. Basic Actions. *American Philosophical Quarterly* 2 (2): 141-8.
- Davidson, Donald. 2001 [1971]. Agency. In *Essays on Actions and Events*, 2nd ed., 43-62. Oxford: Oxford University Press. doi: [10.1093/0199246270.003.0003](https://doi.org/10.1093/0199246270.003.0003).
- Epstein, Brian. 2015. *The Ant Trap: Rebuilding the Foundations of the Social Sciences*. Oxford Studies in Philosophy of Science. Oxford: Oxford University Press. doi: [10.1093/acprof:oso/9780199381104.001.0001](https://doi.org/10.1093/acprof:oso/9780199381104.001.0001).
- French, Peter A. 1984. *Collective and Corporate Responsibility*. New York: Columbia University Press. doi: [10.7312/fren90672](https://doi.org/10.7312/fren90672).
- Gehring, Thomas, and Johannes Marx. 2023. The Emergence and Effects of Non-hierarchical Collective Agency. *Historical Social Research* 48 (3): 7-39. doi: [10.12759/hsr.48.2023.23](https://doi.org/10.12759/hsr.48.2023.23).
- Haan, Niels de. 2021. On the Relation Between Collective Responsibility and Collective Duties. *Philosophy* 96 (1): 99-131. doi: [10.1017/S0031819120000364](https://doi.org/10.1017/S0031819120000364).
- Haslanger, Sally. 2003. Social Construction: The 'Debunking' Project. In *Socializing Metaphysics*, ed. Frederick F. Schmitt, 301-25. Lanham: Rowman & Littlefield.
- Haslanger, Sally. 2012. *Resisting Reality: Social Construction and Social Critique*. Oxford: Oxford University Press. doi: [10.1093/acprof:oso/9780199892631.001.0001](https://doi.org/10.1093/acprof:oso/9780199892631.001.0001).

- Hess, Kendy M. 2018. Who's Responsible? (It's Complicated.) Assigning Blame in the Wake of the Financial Crisis. *Midwest Studies In Philosophy* 42 (1): 133-55. doi: [10.1111/misp.12087](https://doi.org/10.1111/misp.12087).
- Hindriks, Frank. 2009. Corporate Responsibility and Judgment Aggregation. *Economics and Philosophy* 25 (2): 161-77. doi: [10.1017/s0266267109990034](https://doi.org/10.1017/s0266267109990034).
- Hjorthen, Fredrik D. 2020. Who Should Pay for Humanitarian Intervention? *European Journal of Political Theory* 19 (3): 334-53. doi: [10.1177/1474885117697468](https://doi.org/10.1177/1474885117697468).
- Huseby, Robert. 2010. Sufficiency: Restated and Defended. *Journal of Political Philosophy* 18 (2): 178-97. doi: [10.1111/j.1467-9760.2009.00338.x](https://doi.org/10.1111/j.1467-9760.2009.00338.x).
- Isaacs, Tracy. 2014. Collective Responsibility and Collective Obligation. *Midwest Studies In Philosophy* 38 (1): 40-57. doi: [10.1111/misp.12015](https://doi.org/10.1111/misp.12015).
- Jackson, Frank. 1987. Group Morality. In *Metaphysics and Morality: Essays in Honour of J.J.C. Smart*, ed. Philip Pettit, Richard Sylvan, and Jean Norman, 91-110. Oxford: Blackwell.
- Jansen, Ludger. 2014. A Plural Subject Approach to the Responsibilities of Groups and Institutions. *Midwest Studies In Philosophy* 38 (1): 91-102. doi: [10.1111/misp.12018](https://doi.org/10.1111/misp.12018).
- Kagan, Shelly. 2011. Do I Make a Difference? *Philosophy & Public Affairs* 39 (2): 105-41. doi: [10.1111/j.1088-4963.2011.01203.x](https://doi.org/10.1111/j.1088-4963.2011.01203.x).
- Kahn, Elizabeth. 2014. The Tragedy of the Commons as an Essentially Aggregative Harm. *Journal of Applied Philosophy* 31 (3): 223-36. doi: [10.1111/japp.12057](https://doi.org/10.1111/japp.12057).
- Knight, Carl. 2017. Reflective Equilibrium. In *Methods in Analytical Political Theory*, ed. Adrian Blau, 46-64. Cambridge: Cambridge University Press. doi: [10.1017/9781316162576.005](https://doi.org/10.1017/9781316162576.005).
- Kutz, Christopher. 2000. *Complicity: Ethics and Law for a Collective Age*. Cambridge Studies in Philosophy and Law. Cambridge: Cambridge University Press. doi: [10.1017/CBO9780511663758](https://doi.org/10.1017/CBO9780511663758).
- List, Christian, and Philip Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press. doi: [10.1093/acprof:oso/9780199591565.001.0001](https://doi.org/10.1093/acprof:oso/9780199591565.001.0001).
- Martin, Jane R. 1972. Basic Actions and Simple Actions. *American Philosophical Quarterly* 9 (1): 59-68.
- Mathiesen, Kay. 2006. We're All in This Together: Responsibility of Collective Agents and Their Members. *Midwest Studies In Philosophy* 30 (1): 240-55. doi: [10.1111/j.1475-4975.2006.00137.x](https://doi.org/10.1111/j.1475-4975.2006.00137.x).
- May, Larry. 1992. *Sharing Responsibility*. Chicago: University of Chicago Press.
- Parfit, Derek. 1987. *Reasons and Persons*. 3rd corrected ed. Oxford: Clarendon Press.
- Pauer-Studer, Herlinde. 2014. A Constitutive Account of Group Agency. *Erkenntnis* 79 (9): 1623-39. doi: [10.1007/s10670-014-9632-y](https://doi.org/10.1007/s10670-014-9632-y).
- Pettit, Philip. 2007. Responsibility Incorporated. *Ethics* 117 (2): 171-201. doi: [10.1086/510695](https://doi.org/10.1086/510695).
- Pinkert, Felix. 2014. What We Together Can (Be Required to) Do. *Midwest Studies In Philosophy* 38 (1): 187-202. doi: [10/ggwnhz](https://doi.org/10/ggwnhz).
- Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: Belknap Press.
- Ritchie, Katherine. 2013. What Are Groups? *Philosophical Studies* 166 (2): 257-72. doi: [10.1007/s11098-012-0030-5](https://doi.org/10.1007/s11098-012-0030-5).
- Ritchie, Katherine. 2015. The Metaphysics of Social Groups. *Philosophy Compass* 10 (5): 310-21. doi: [10.1111/phc3.12213](https://doi.org/10.1111/phc3.12213).

- Ritchie, Katherine. 2020. Social Structures and the Ontology of Social Groups. *Philosophy and Phenomenological Research* 100 (2): 402-24. doi: [10.1111/phpr.12555](https://doi.org/10.1111/phpr.12555).
- Schwenkenbecher, Anne. 2021. *Getting Our Act Together: A Theory of Collective Moral Obligations*. Abingdon: Routledge. doi: [10.4324/9781003097136](https://doi.org/10.4324/9781003097136).
- Shue, Henry. 1999. Global Environment and International Inequality. *International Affairs* 75 (3): 531-45. doi: [10.1111/1468-2346.00092](https://doi.org/10.1111/1468-2346.00092).
- Sinnott-Armstrong, Walter. 2005. It's Not My Fault: Global Warming and Individual Moral Obligations. In *Perspectives on Climate Change: Science, Economics, Politics, Ethics*, ed. Walter Sinnott-Armstrong and Richard B. Howarth, 285-307. *Advances in the Economics of Environmental Resources* 5. Bingley: Emerald. doi: [10.1016/S1569-3740\(05\)05013-3](https://doi.org/10.1016/S1569-3740(05)05013-3).
- Smith, Holly. 2012. Using Moral Principles to Guide Decisions. *Philosophical Issues* 22: 369-86. doi: [10.1111/j.1533-6077.2012.00235.x](https://doi.org/10.1111/j.1533-6077.2012.00235.x).
- Smith, Thomas H. 2009. Non-Distributive Blameworthiness. *Proceedings of the Aristotelian Society* 109: 31-60. doi: [10.1111/j.1467-9264.2009.00257.x](https://doi.org/10.1111/j.1467-9264.2009.00257.x).
- Spiekermann, Kai. 2014. Small Impacts and Imperceptible Effects: Causing Harm with Others. *Midwest Studies In Philosophy* 38 (1): 75-90. doi: [10/ggwnhs](https://doi.org/10/ggwnhs).
- Tersman, Folke. 2018. Recent Work on Reflective Equilibrium and Method in Ethics. *Philosophy Compass* 13 (6): e12493. doi: [10.1111/phc3.12493](https://doi.org/10.1111/phc3.12493).
- Williams, Bernard. 1985. *Ethics and the Limits of Philosophy*. London: Fontana Press.
- Wringe, Bill. 2016. Collective Obligations: Their Existence, Their Explanatory Power, and Their Supervenience on the Obligations of Individuals. *European Journal of Philosophy* 24 (2): 472-97. doi: [10.1111/ejop.12076](https://doi.org/10.1111/ejop.12076).
- Young, Iris Marion. 2004. Responsibility and Global Labor Justice. *Journal of Political Philosophy* 12 (4): 365-88. doi: [10.1111/j.1467-9760.2004.00205.x](https://doi.org/10.1111/j.1467-9760.2004.00205.x).
- Young, Iris Marion. 2011. *Responsibility for Justice*, ed. Martha Nussbaum. Oxford: Oxford University Press. doi: [10.1093/acprof:oso/9780195392388.001.0001](https://doi.org/10.1093/acprof:oso/9780195392388.001.0001).