

Group Duties: Response to Critics

To be published in *Journal of Social Ontology*

Abstract This is a response to critical commentaries by Gunnar Björnsson, Olle Blomberg, Anne Schwenkenbecher, and Bill Wringe regarding my monograph *Group Duties: Their Existence and Their Implications for Individuals* (Oxford University Press, 2019).

Key Words: Group duties, group agency, group responsibility

I am grateful to each of the commentators for their detailed and generous engagement with *Group Duties*. I have learned a lot from reading, and thinking about, their essays, as well as from our meeting at *Social Ontology 2019* in Tampere.

1. Reply to Anne Schwenkenbecher

Schwenkenbecher questions the usefulness of the ontological questions at the heart of *Group Duties*. As she rightly notes, much of *Group Duties* is concerned with which groups can have duties. Roughly, Schwenkenbecher asks: why care? She suggests that the most important questions are the distinctively *epistemic* questions that arise when many agents must act to produce some morally valuable outcome. Specifically, Schwenkenbecher encourages us to ask the following epistemic questions: “What does the problem look like to the agents involved? What can agents be obligated to do given what they (can) know and how they would reasonably perceive the situation?” (this volume, p. **)

The problem is that these questions cannot be answered until we have answered a prior ontological question: who are the “agents” in a given scenario? Does this include only each individual agent? (This seems to be Schwenkenbecher’s implication.) Or does “the agents” also

include some group agent(s)? That is: in addition to individual agents asking what *they each* should do given what they can know and perceive, perhaps *the group* should also ask what *it* should do given what it can know and perceive. *Groups Duties* affirms this possibility, if there is a group united under a rationally-operated decision-making procedure that can attend to moral considerations; and denies this possibility, if there is not such a group. Thus, the ontology matters for the following question: *which entities* should ask themselves Schwenkenbecher's epistemic questions?

I do not deny that the epistemic issues matter, and, as Schwenkenbecher points out, much of the individual-level analysis in *Group Duties* uses premises about what individuals can reasonably know, believe, expect, and so on. Yet these individual-level epistemic premises derive from the group-level ontology—where that ontology is not merely a matter of how humans communicate (as Schwenkenbecher suggests).

For example, take Schwenkenbecher's suggestion that "agents should we-reason when the collective option is their best (moral) bet given the evidence." (Schwenkenbecher 2019, 163)¹ When is the collective option their best moral bet? *Group Duties* argues that it is their best bet when they together constitute a moral coalition that is oriented around a relevant goal. To use Schwenkenbecher's trail-ride example: if we assume the trail riders constitute a coalition rather than a collective (I'll relax this below), then, from the moment of Jane's fall, when

¹ I was unable to read Schwenkenbecher (2019) before *Group Duties* was due with the press in late 2018. I agree with much of Schwenkenbecher's (2019) suggestions for members of non-collective groups—except for the idea that there is a singular duty that these members hold 'jointly', that they should 'we-reason' (rather than 'coalition-reason', as I argue in Chapter Five), or that convincing others to join in a collective action cannot be the upshot of I-reasoning (indeed, this is often the upshot of I-reasoning amongst members of combinations (Collins 2019, 150)).

Schwenkenbecher and Alice exchange concerned glances or otherwise discontinue the normal swing of the trail ride, then they are commonly known to (i) hold the goal of assisting Jane, (ii) be disposed to work with the other to realize that goal.

That is, swiftly after Jane's fall, Schwenkenbecher and Alice transform into a moral coalition with the goal of helping Jane, not just a non-moral coalition with the goal of trail-riding. At that moment—when the goal that unites them as a coalition changes, with a glance and a halt—then, as Schwenkenbecher puts it, “pattern-based reasons outweigh individual-based reasons.” (2019, 163). In this way, the ontological question (‘what kind of group do Schwenkenbecher and Alice constitute, in *Group Duties*’ taxonomy?’) helps to answer Schwenkenbecher’s practical question (‘how should Schwenkenbecher and Alice each reason?’).

This might seem trivial: why entertain the possibility that Schwenkenbecher could have reasoned on the assumption that Alice would continue to prioritise the trail ride, after Jane's fall? It seems ridiculous to imagine Alice wouldn't have wanted to help Jane, or that Schwenkenbecher could have reasonably believed this. In this small-scale case, it is indeed ridiculous. But there are similarly-structured, larger-scale cases where it is not ridiculous. These are cases where a moral goal has become salient (consider climate change or poverty), but many people have (and their fellow coalition-members should presume they have) continued to pursue the non-moral goal (such as ‘living a comfortable life’). In those kinds of cases, each agent should reason on the presumption that her fellows will continue to pursue their non-moral goal. This requires that the agent pursues the second-best-scenario pattern (where coordinating patterns are counted as one type of ‘pattern’). That is: if your non-moral coalition does not transform into a moral coalition, then you “should do what is best given the assumption that the other member will continue to pursue the original non-moral coalition goal.” (Schwenkenbecher, this volume, p. **) In this way, the ontological distinction between moral and non-moral coalitions has practical implications.

Schwenkenbecher rightly notes that coalition-reasoning is not we-reasoning as most people understand it. It is reasoning on the presumption that others will do their part in some group pattern. It is not reasoning from the group's perspective. This is why I coined the term 'coalition-reasoning.' Importantly, though, coalition-reasoning is not mere I-reasoning. I-reasoning does not allow the aforementioned presumption—not even in the 'pro-group' version of I-reasoning, discussed by Hakli, Miller, and Tuomela (2010)). I called coalition-reasoning a 'variant' of we-reasoning, but perhaps this was confusing and I should rather have said that coalition-reasoning is a new category in between I-reasoning and we-reasoning.

Regarding non-moral coalitions in a group rescue case, Schwenkenbecher asks: "if all members of a non-moral coalition are morally conscientious and act on their coordination duties [that is, if they act presuming that the others will continue to pursue the non-moral goal], do they all stop being a non-moral coalition and therefore become a combination? Does this imply that members of combinations engage in coalition reasoning?" (Schwenkenbecher, this volume, p. **) The answers to these questions are 'yes' and 'no', respectively. The temporal sequencing helps to explain why.

An individual in a non-moral coalition should coalition-reason: she should presume the others will pursue the shared non-moral goal, even while she aims at the moral goal. Her duty to reason in this way arises from, and exists while, she and the others are part of a coalition. But by performing this duty, the individual transforms the group into a combination—because the individual no longer holds the group's non-moral goal. Thus, the duty to coalition-reason is held and performed by members of a coalition, even though, if the coalition is non-moral, this very process of reasoning will transform the group into a combination. If each and every member of the non-moral coalition engages in this reasoning, then we will have a kind of 'pluralistic ignorance' case: everyone will believe everyone else holds a certain non-moral goal (such as continuing a trail ride), but in fact no one holds it (instead, everyone holds a moral

goal, such as helping Jane). They don't hold the latter goal as a coalition, since no one knows that anyone else has the goal of helping Jane. Instead, we now have a combination, where that combination was formed via each member performing the coalition-reasoning duties that they acquired back when the group was a non-moral coalition. The ontology of moral coalitions, non-moral coalitions, and combinations gives us the conceptual resources to categorise what is going on at each temporal stage of this process.

With all that said, as Schwenkenbecher rightly notes in her commentary, the trail-riders are perhaps best categorised as a collective, not a (moral or non-moral) coalition. As I say in the book, “[i]t’s not always easy to categorize real-world groups under the three concepts [of combination, coalition, and collective]. As with many concepts, there are likely to be vague or borderline cases.” (Collins 2019, 10) From the vignette given in Schwenkenbecher’s commentary, I’m not entirely sure whether the riders are united under a rationally-operated group-level decision-making procedure such as conversation-based consensus.

Notice, though, that what matters for whether the members are *united under* a procedure is not just whether they can come to an agreement via talking, or how they communicate, as Schwenkenbecher’s commentary suggests (and which, she implies, would be somewhat counterintuitive as a morally significant difference). Rather, being ‘united under’ a procedure requires something more than mere verbal communication. It requires that the individuals meet three conditions: (1) each is committed to abide by the procedure, (2) each has inputs into the procedure, and (3) each can ‘act as the group’ when they enact the results of the procedure (Collins 2019, 12-14). That is, constituting a collective is not something that “boils down to communicating some kind of plan to others.” (Schwenkenbecher, this volume, p. **).

To return to Schwenkenbecher’s reasonable request that the ontology pays dividends: why does it matter whether the trail riders are a coalition or a collective? One answer is simply

metaphysical: it matters for the list of duties and agents that there are in the world; there is one more duty and one more agent if they are a collective than if they are a coalition.

But this, I suspect, will not satisfy Schwenkenbecher, who wants the group-level ontology to make a difference *for what actions the riders should individually perform*. On that point, the difference is this. If the riders constitute a collective (whether or not that collective pre-dates the accident), then their individual duties are to *use their role in the collective* with a view to helping Jane. Let's assume their decision-making procedure is conversation-based consensus. Then, their mode of decision-making is to reach a consensus via conversation—not to vote, or form committees, or whatever. So their individual obligation is to engage in conversation, with a view to consensus. This is exactly what was done by Schwenkenbecher and Alice in the vignette.

Their decision-making procedure could have been different—it could have been based on gestures, rather than words, for example; or it could have required voting—in which case silent decision-making could have amounted to group decision-making. It all depends on what the procedure is, with regard to which the riders meet conditions (1), (2), and (3) above. And assuming the trail-riders do constitute a collective by each meeting conditions (1), (2), and (3), then the group's duty comes into existence as Jane comes off the horse: the group's duty does not come into existence only when Schwenkenbecher initiates the group's decision on what to do. (This is on the assumption that the riders' being united under a procedure pre-dates the use of that procedure to respond to the specific bucking incident, since the procedure has been used to select trail routes, speeds, and so on.)

By contrast, if they are a coalition—not meeting (1)-(3) with regard to any group-level decision-making procedure—then there is no procedure to guide and structure their individual-level duties. They should instead each plump for their part of the optimal pattern of group

actions for helping Jane, once they have transformed from a non-moral coalition to a moral coalition in the way described above. In this way, the difference between the duties of collectives' members and coalitions' members is not that the former involve communication while the latter do not. A decision-making procedure might need almost no communication (consider secret ballots); and one's part in the optimal coalition pattern might require that many words be exchanged. Rather, the difference is whether the duties require engagement with a group-level decision-making procedure.

What about the other people in Schwenkenbecher's vignette: the paramedics and the cyclist? Are they members of a coalition or a collective, together with Schwenkenbecher and Alice? Here it matters that goals of coalitions can be characterised in a more or less coarse-grained way (Collins 2019, 143-144)—so if the members' goals differ only slightly, then they are a coalition united around a generally-described goal (that is, a goal that elides the slight difference), but not a coalition united around the more specifically-described goals. From the time they were made aware of the accident, the paramedics and cyclist became members of the moral coalition (alongside the collective constituted by Schwenkenbecher and Alice), whose members each held the goal 'helping Jane.' This meant they should each reason on the presumption that the others would do their part towards this goal (where, of course, the first step of that 'part' might require communication about what part each will perform; this doesn't automatically imply a decision-making procedure).

Things would have been different if all those involved together constituted a collective. In that case, their individual duties would have related to putting inputs into the group's decision-making procedure and using their designated roles in certain ways. But I suspect—again, from the brief vignette in Schwenkenbecher's commentary—that conditions (1), (2), and (3) for being united under a rationally-operated group-level decision-making procedure were not met by all these agents (unlike amongst the three riders themselves, or amongst the paramedics

themselves). Thus, the trail-ride example is a good one for demonstrating why collectivisation is often unnecessary; responsiveness amongst members of a moral coalition is often good enough.

Let's assume the trail-riders are a collective. In that case, Schwenkenbecher suggests, my account of this group barely differs from authors who argue that, in cases with non-agent groups, "group members must perform those actions that form part of (or are constitutive of) a collectively available action or outcome." (Schwenkenbecher, this volume, citing herself, Wringer, Pinkert, and Aas.) I think there remain at least two important differences between my account and these ones. First, the structure of members' duties in collectives differs from the quote just given: as Chapter Seven of *Group Duties* argues, the duty is to use one's role to put inputs into the collective's procedure (see esp. Collins 2019, 198). This isn't just a matter of performing one's part in a collectively available outcome. Second, in my view, groups can have duties only if they are agents. This is not so in the views of Schwenkenbecher, Wringer, Pinkert, or Aas. Thus, we reach substantively different conclusions (on the content of group' members' duties) and for different reasons (regarding whether or not the group must be an agent to bear a duty).

2. Reply to Olle Blomberg

The foregoing brings us naturally to Olle Blomberg's commentary. Blomberg describes (without endorsing) a strategy for opposing non-collective groups' duties, which is different from the strategy I pursue in *Group Duties*. The core premise in Blomberg's alternative strategy is this: in order for me to have a duty to ϕ , my exercising my ability to ϕ must depend in some way on my decisions—including if ϕ is an automatic or habitual response, in which case, my "duty to do what now requires the automatic response can be traced back to past opportunities she had to decide whether or not to shape and cultivate her behaviour so that she would acquire

the automatic response behaviour.” (this volume, p. **) A second premise is that non-collective groups cannot make any decisions. From this (roughly), we get that non-collective groups cannot have any duties—because duty-behaviour must be traceable to decisions.²

I hesitate to use Blomberg’s suggested argument because I suspect we can have duties to perform actions we’ve never had the opportunity to make any decisions about, not even of the character-cultivating kind. The duty against murder might be like this. In any case, I do not want to assume there are no such duties. This is why I declined to use this premise. Instead, I relied on (what I view as) a weaker, and therefore less controversial, core premise: that in order for me to bear a duty to ϕ , I must have the ability to make decisions that are sensitive to *some duty or other* (not necessarily decisions that are related to ϕ itself).

Why should we believe my weaker core premise? As Blomberg rightly notes, my answer is an inference to the best explanation: I claim that my core premise is a good explanation of why dogs and babies lack duties, while most adult humans have them. Dogs and babies cannot make decisions that are sensitive to any duties; most adult humans can make decisions that are sensitive to at least some duties. Of course, an inference to the best explanation is never conclusive.

Inferences to the best explanation are inconclusive because we can argue forever about which explanation is the best one. On this, Blomberg asks: “Why isn’t it the presence/absence of responsiveness to moral reasons that best explains why the babies and dogs don’t have duties while normal adult human beings do?” (this volume, p. **) I find this question confusing. I

² Blomberg wonders if this suggested argument would force me to accept non-collective groups’ duties, via an analogy to humans’ temporal parts and groups’ human parts. But my use of that analogy arose in my discussion of groups’ abilities, which, I argue, requires a thinner kind of ‘glue’ than groups’ duties; I wouldn’t use the same analogy in a duty context.

argue that any duty-bearing entity must have the ability to *fulfil* at least one duty (Collins 2019, 87). Fulfilling a duty, in my view, requires that the performance of the duty “has a psychological explanation that accords with its moral justification” (Collins 2019, 86). In plainer terms: an entity’s fulfilling a duty requires that the entity has made decisions because of duties or because of the justifications that ground duties. Duties are grounded in (or perhaps just are) weighty moral reasons. Thus, I think the book endorses the explanation suggested in Blomberg’s question: a creature can bear a duty only if it can be responsive to moral reasons.

Now, Blomberg may reply that mere responsiveness-to-moral-reasons does not require making any decisions. But in that case, I don’t think responsiveness-to-moral-reasons separates me from a trained dog. Such dogs are trained to be responsive to moral reasons; they act responsively to those reasons in order to receive treats or approval. But trained dogs don’t make decisions that have a psychological explanation that accords with their moral justification. So duty-fulfillment separates trained dogs from me (and other duty-bearers); Blomberg’s decision-less interpretation of responsiveness-to-moral-reasons does not. In sum, it still seems to me that the ability to make decisions (that have a psychological explanation that accords with their moral justification) is the best explanation of why duties fall on the creatures they do, in non-group cases.

But what about group cases? Blomberg provides an argument for non-collective groups’ duties, building from the notion of we-reasoning. As he notes, Chapter Five of *Group Duties* offers a reductive account of we-reasoning. (I call my account ‘coalition-reasoning’, to mark its difference from non-reductive we-reasoning and to signal the groups in which such reasoning is appropriate.) Blomberg proposes to see we-reasoning differently: it requires a “perspective shift”, in which an individual “picks the outcome that is best for the group to bring about, and proceeds to ask herself how she should do her part of the group action will bring that outcome about.” (this volume, p. **)

He suggests that, when each of a random group of beachgoers is we-reasoning in a collective rescue case, the statement ““You [second-person plural] ought to rescue the swimmer’ ... would arguably be a fitting input into each beachgoer’s *first-person plural* future-directed decision-making” (this volume, p. **, square brackets and italics original). At face value, this suggestion troubles me, because a rag-tag group of beachgoers isn’t structured so as to be able to make any decisions: there is no plural decision-making into which the statement could be an input. But what Blomberg means is that the duty could be an input into “each member’s deliberation that not only concerns what the group ought to do, but that is carried out from the group’s point of view.” (this volume, p. **) There is only individual decision-making. So far, so good. But it seems I could engage in deliberation about what my dog ought to do, that’s carried out from the dog’s point of view. The dog’s ‘duty’ could be an ‘input into my decision-making’, in this sense. But that is insufficient to show that *the dog* bears the duty. Likewise, I suggest, the idea that a (purported) group duty could be an input into an individual’s decision-making is insufficient to show that *the group* bears the duty.

Blomberg writes that if the random beachgoers successfully pull off the rescue, then “the beachgoers’ joint performance has a psychological explanation that accords with its moral justification.” (this volume, p. **) This is true in a distributive sense: each beachgoer’s performance was explained by their beliefs and preferences regarding the horror of drowning, for example. But there is no psychological explanation at the group level, because (as I think Blomberg agrees) the group has no beliefs or desires. So there are no beliefs or desires that can serve as a *group-level* psychological explanation of the group-level behaviour. The collection of individual-level psychological explanations demonstrates only that each individual in the collection can be a duty-bearer, not that the collection as a whole can. As Blomberg notes, the members of a non-collective group lack the “memories, expectations and intentions” that unite an individual’s temporal parts. (this volume, p. **) So we should not draw an analogy between

Blomberg's long-term duties (to help a man bobbing in the water) and a non-collective group's duties. (As I mentioned above, I think the analogy is more appropriate when discussing abilities, where the interconnections between an entity's parts can be looser than they must be for duties.)

I'm also sceptical that we need the non-distributive moral duty (held by the group or plurality) "in order to explain how it can be that each [individual] ... has an individual moral duty to perform rescue actions." (this volume, p. **) I discuss this generally on pp. 29-33 of *Group Duties*, where I argue that the value of an outcome, combined with the non-distributive ability (not duty) of the group or plurality, can satisfactorily explain each individual's duty to do their part. But Blomberg is talking about a specific kind of case, namely, one structured along *Hi-Lo* lines. He questions how it is that two players could end up settling on "Hi" unless they are engaged in we-reasoning. He writes: "it is not clear why the fact that there is a commonly known shared goal would provide *any* evidence, however small, that the others will do their part of the optimal pattern of actions." (this volume, p. **)

But the commonly known shared goal is not the only premise in coalition-reasoning, which is the reasoning I recommend in this kind of case. Instead, a crucial premise in coalition-reasoning is "It is a common belief [...] between you and me that me doing X and you doing Y is the best way of patterning our respective actions so as to realize [the outcome]" (Collins 2019, 139, also quoted by Blomberg). This premise makes coalition-reasoning different from garden-variety I-reasoning. This premise places coalition-reasoning somewhere between I-reasoning and we-reasoning in Blomberg's sense. Perhaps Blomberg would find this premise "a brute psychological feature" and therefore "unsatisfactory" as a justification for choosing Hi. (This is what Blomberg says of Parfit's response to the puzzle (this volume, p. **).) But I don't see why the premise I propose is any more a "brute psychological feature" than is the perspective shift described by advocates of we-reasoning. Thus, I still don't think we need to appeal to a

group-level duty (or to we-reasoning) to explain what individuals should do in Hi-Lo cases; instead, if each individual has the presumption that the others know Hi is the best way to realise their (respective) goal, this does the same work.

3. Reply to Bill Wringe

We-reasoning also features in Bill Wringe's analysis of my account of collectives. Wringe correctly observes that, in my view, "a coalition whose members are each committed to a policy of we-reasoning does not constitute a group agent." (this volume, p. **) But he thinks the following does: "a group [that] contains a single member who is willing to subordinate their personal decision to what they take to be majority opinion (or even only to unanimous opposition)." (this volume, p. **)

From this description, I'm not yet able to judge whether such a group would count as a collective under my framework. We need to know not just that one individual will subordinate their individual-level decisions to majority opinion, but also that those whose opinions are factored into that majority are also members, that is, that (1) each is committed to abide by the procedure, (2) each has inputs into the procedure, and (3) each can act as the group. If I spontaneously decide to have my personal decisions dictated by what a majority of Anne, Olle, Bill, and Gunnar want me to do, then that doesn't entail that the five of us together constitute a collective. In order for us to constitute a collective, Anne, Olle, Bill, and Gunnar also need to meet conditions (1), (2), and (3), where the procedure might be something like 'using majority vote amongst Anne, Olle, Bill, and Gunnar to decide what Stephanie will do.' So my account of collectives is more restrictive than Wringe's commentary implies (though I agree my account is quite permissive, and intentionally so).

Yet Wringe is correct about his other inference: that one can cease to be a member by ceasing to be committed to abide by the collective's decision-making procedure. In this case, one would

cease to have membership duties (that is, individual-level duties that are entailed by the fact the individual is a member of a collective with a duty). Wringe worries that “a duty which ceases to be a duty if one refuses to perform it isn’t much of a duty.” (this volume, p. **) This is a fair concern. Of course, people often retain membership in collectives, even though that membership entails duties—because the membership also brings benefits. So I’m not sure how often Wringe’s concern plays out on the ground. And notice that—as Wringe suggests—such an individual may still have all sorts of duties: duties to re-join the group and thereby re-acquire membership duties, duties to act upon the group from the outside to get it to do its duty, duties to act independently with a view to the outcome over the which the group has a duty. None of these are, strictly speaking, membership duties. They are bog-standard individual duties. But they will go a long way towards getting the individual onto the relevant hooks. Wringe thinks justifying such duties will be a “non-trivial, though perhaps not hopeless task.” (this volume, p. **) I think it’s a task for which we can draw on centuries of ethical theorising that grounds individual duties in facts about consequences, universalizability, or character.

The notion of character brings us to the second half of Wringe’s commentary. Wringe accurately summarises my view that (a) an entity must perform a duty for the right reasons (‘fulfil’ it) in order to have moral worth, and (b) an entity must be able to fulfil at least one duty in order to bear any duty. In my view, fulfilling a duty requires making decisions—which, I claim, non-collective groups can’t do.

Wringe observes that virtue ethicists might take issue with (a), if they valorise duties that are performed via habituation. But I think this is true only if the decisions that lead to duty-fulfilment must be temporally proximate to the performance. The habituated virtuous agent *does* perform duties for the right reasons, by making decisions. It’s just that those decisions were made when she decided to pursue virtuous habituation via moral education, the emulation

of virtuous agents, and so on (this was suggested in Blomberg's commentary and will reappear in Björnsson's).

Wringe notes a "striking disanalogy" between individuals' decisions and collectives' decisions: individuals typically don't make decisions via procedures. (this volume, p. **) This disanalogy arises because individuals (unlike collectives) are not constituted by a plurality of agents at any one time (Collins 2019, 45). If individuals were thus constituted, then I suspect procedures would feature in (rational) individuals' decisions in much the way they feature in collectives' decisions. So I don't think this analogy does anything to undermine (a).

What about (b)? This was 'an entity must be able to fulfil at least one duty in order to bear any duty.' Like Blomberg, Wringe observes that I claim (b) distinguishes most adult humans from dogs, tables, and young children. Like Blomberg, Wringe suggests an alternative distinguishing feature: most adult humans can "can perform duties in a non-fluky and relatively reliable manner", while dogs, tables, and young children cannot. (this volume, p. **) The problem with Wringe's alternative is that it is false: dogs, tables, and young children *can* perform duties in a non-fluky and relatively reliable manner, bearing in mind that 'performance' is just a matter of behaviour. Tables can be placed where they reliably won't harm anyone; dogs can be trained to reliably be gentle; young children can be bribed into reliably sharing with their siblings. Thus, as in my response to Blomberg, I remain convinced that (b) distinguishes duty-bearers from non-duty-bearers.

Wringe poses another problem for (b): it seems to imply that psychopaths are not duty-bearers. Wringe raises two problems with this result. First, it renders unwarranted our "reactions" to psychopaths who breach (apparent) duties. I assume these reactions are reactive attitudes. I'm no expert on psychopaths. I'm not sure psychopaths can't fulfil duties, in my sense. However, following David Shoemaker (2011; 2015), I think reactive attitudes are a multi-faceted affair.

Shoemaker argues that at least *some* of these attitudes are not apt when directed at psychopaths: attitudes of indignation or resentment. These attitudes are the only ones I meant to subsume under ‘moral worth,’ which I characterised as praiseworthiness and blameworthiness. (This raises the question: are *other* reactive attitudes apt for non-collective groups—perhaps, reactive attitudes that do not presuppose that the target is a duty-bearer? I haven’t argued for it, but I think the answer is ‘yes’.) Second, Wringe claims “we think it’s permissible to coerce psychopaths in some ways, but not in others.” (this volume, p. **) I don’t know which ways Wringe has in mind. Perhaps the idea is ‘we may coerce psychopaths in ways we may not coerce young children.’ If this is true, I would guess that’s well-explained by psychopaths advanced rational capacities compared to young children.

Also on coercion, Wringe asks about group agents that are incapable of attending to moral considerations (and therefore are not collectives, strictly speaking): “May such (collective) agents permissibly be coerced?” (this volume, p. **) No answer to this is implied by *Group Duties*, which has nothing to say about what makes an entity worthy of moral consideration or moral rights. A ‘yes’ answer arises only if you think (i) all and only entities that have the capacity to set ends are worthy of moral consideration, and (ii) an entity’s being worthy of moral consideration implies it’s impermissible to coerce the entity. I don’t believe (i) or (ii). I think animals are worthy of moral consideration. I think it’s permissible to coerce me (for example, to coerce me into performing important duties). None of that is argued for in *Group Duties*. In any case, my answer to Wringe’s question (quite separately from the book’s argument) is ‘of course!’

In his final section, Wringe turns to a different argument for non-collective groups’ duties. This ‘argument from phenomenology’ was given by Wringe (2016) and is discussed in *Group Duties* (pp. 36-38). There, I argue that individuals’ phenomenology in group rescue cases can be understood in individualistic terms. Wringe worries that my discussion of the

phenomenology is third-personal. To allay any doubt: when I have confronted medical emergencies, my phenomenology has been one of asking what we *each* should do (just as Schwenkenbecher concluded: ‘I should tend to Jane; Alice should mind the horses’ (Schwenkenbecher, this volume, p. **)), not what we *together* should do. This holds even if only two out of three present people are needed: in that case, ‘we ought to do something’ is true, distributively: A ought to act; B ought to act; C ought to stay out of the way. (Or if it’s not yet clear who should stay out of the way, then each has a duty to act until enough people are acting, at which point the remaining inactive person has a duty to stay out of the way).

Wringe argues that this distributive analysis cannot account for the backwards-looking phenomenology, which is permeated by the thought “we messed up.” (this volume, p. **) As I say in *Group Duties* (2019, pp. 8-9), I don’t think facts about duty are perfectly parallel to facts about backward-looking responsibility. And as I mentioned above, some reactive attitudes may indeed be warranted towards non-collective groups. This does not, I think, imply that such groups have duties. (Again, the reactive attitudes in my ‘moral worth’ argument (praise and blame) are not the only reactive attitudes there are; I think these attitudes entail that the target is able to make decisions for the right moral reasons; other reactive attitudes may not entail this.)

4. Reply to Gunnar Björnsson

Like Blomberg and Wringe, Björnsson focuses on my claim that non-collective groups cannot bear duties. Björnsson develops two examples intended to show that one can perform a duty, with a psychological explanation that accords with the duty’s moral justification, even though one does not make a decision. As he carefully explains, if these examples succeed then they undermine my overall argument that non-collective groups cannot bear duties.

I'm inclined to take different approaches to each of Björnsson's two examples. As he suggests, his example *Noticing and Wondering* is a case where the 'performance' at issue (namely, noticing and wondering) is not obviously the kind of performance over which there can be a duty. But let's suppose there is a duty over this performance *and* let's suppose the agent, Fara, deserves moral praise for that performance. (It's important that these are two separate suppositions, and Björnsson's point is focused on the second; on my view, the first supposition requires only that Fara can make at least one duty-grounded decision, not that Fara can make a duty-grounded decision that relates to the duty in question.) I'd suggest that the moral praise depends upon the fact that Fara made prior decisions (around disposition-cultivation and so on) that relate to the noticing and wondering. I mean this in an ordinary sense of 'decision.'

What if Fara notices and wonders without *ever* having made any decisions related to the goodness of being the kind of person who notices and wonders? Then we may admire Fara without giving her moral praise. (This is what Shoemaker (2015) calls 'attributability' without 'accountability.' As mentioned above, reactive attitudes are numerous and varied.) In this case, Fara bears a duty and she performs the duty. She simply doesn't fulfil it. Thus, an entity's performance's having a psychological explanation that accords with its moral justification implies that the entity has made *some* relevant decision—though that decision may have been in the distant past and not directly about the performance in question. To clarify how this meshes with my reply to Blomberg: although moral praise for ϕ entails a ϕ -related decision, bearing a duty to ϕ does not. Duties can be performed without decisions and, therefore, without praise.

Björnsson's other example—*Refraining*—is one where Fara refrains from taking a wallet on the seat next to her, without considering the possibility of taking the wallet. Björnsson suggests that Fara's refraining must have a psychological explanation that accords with its moral

justification. Fara even supererogates by finding the owner of the wallet. Yet, Björnsson says, Fara has not *decided* to refrain. Thus, we have moral praise without a relevant decision.

For this example, I want to distinguish two performances: refraining from taking the wallet versus finding its owner. The latter came along with decisions. I suspect these decisions affect our assessment of Fara's praiseworthiness for refraining. So let's focus just on the refraining, not the finding. I would deal with the refraining as I dealt with the noticing and wondering: either there was no relevant decision (*ever*, not even related to being the kind of person who would refrain in situations like this), in which case there is no praise (though there might be admiration and other attitudes); or there was a relevant decision, in which case there is praise. Again, if there wasn't a relevant decision, this doesn't mean 'refrain from stealing' cannot be the content of Fara's duty. I don't require that each duty can be performed via a decision; merely that each duty-bearer is able to perform at least one duty via decision.

Perhaps the refraining and finding are too intertwined to be separated in this way. Perhaps Fara refrains from taking the wallet *by way of* finding its owner. But on that interpretation, the refraining does involve a decision (the decision to find the owner)—again, in an ordinary sense of 'decision.' So this is no counterexample to the notion that praiseworthiness—and, by stipulation, duty-fulfilment—implies decisions.

The upshot is the following. Although my notion of decision is thin, it is not quite as thin as Björnsson suggests it must be to accommodate *Noticing and Wondering* and *Refraining*. Because my notion of decision is not so very thin, I also deny that there is a group-level decision in his example *Rescue*. Thus, I don't view these examples as providing reason to reject the idea that a duty-bearing entity is a decision-capable entity.

Let's suppose I'm right about this. Björnsson then makes a further suggestion: that groups without decision-making procedures *are* decision-capable entities. He rightly reports that I

argue non-collective groups have abilities. I just deny that those abilities include the ability to make group-level decisions (or, for that matter, to produce the outcome that the group makes a decision, unless that outcome includes the group transforming into a collective that then makes a decision). Björnsson finds this incredible: if a bunch of beachgoers is able to rescue a drowning swimmer (via a series of separable individually-made decisions), then surely the group is able to make a decision to rescue that swimmer.

Why do I deny that non-collective groups have the ability to make decisions? The answer does not lie in the idea that the relevant ability is the ability to make an 'immediate' decision, where 'immediate' is understood as 'instantaneous' or 'quick.' I introduced 'immediate decision' to contrast a group with a standing decision-making procedure against a group that does not yet have a decision-making procedure but that might come to have one. The latter group has only the 'mediate' ability to make decisions: an ability mediated by its transforming into a collective. So the word 'immediate' was maybe an unfortunate choice on my part: a better word might be 'standing' or 'preexisting.' My claim was that the *standing* or *preexisting* ability to make decisions must inhere in any entity, if that entity can appropriately be attributed duties. A group with a mediate (non-standing, non-preexisting) ability to make decisions can be attributed duties only once that ability becomes immediate, which requires that the members are unified under a decision-making procedure.

Now, Björnsson might reply that a random group *does* have the standing or preexisting ability to make decisions (such as the decision to respond to a sudden proximate emergency). At least, he might say, if such a group has the ability to produce an outcome (which I grant), then surely it has the ability to produce the outcome in which it has made a decision. I don't think this is correct. We must consider what the concept of 'decision' does for us. Decisions govern the subsequent actions of the entity that made the decision. Decisions create intentions. Decisions are part of the web of mental states (beliefs, preferences, etc) of the entity that has made the

decision. Decisions play a certain role in the mental life the entity that has made the decision. But non-collective groups lack the actions, intentions, and other mental states that make up such a mental life. Unlike a decision, the production of an outcome (such as the outcome ‘the swimmer no longer drowning’) does not have this role to play in an entity’s mental life. Outcomes can be produced by the weather. Decisions are tied to an entity with a mental life. This is why non-collective groups are able to produce outcomes but are not able to make decisions.

Björnsson considers a different route to the conclusion that collectives are the only groups that can bear duties. Like Blomberg, he suggests that reasons-responsiveness might be the key to duty-bearing. Unlike Blomberg, Björnsson suggests that perhaps collectives are the only group capable of reasons-responsiveness. I worry that groups other than collectives are capable of reasons-responsiveness, in the same way that dogs, tables, and children are capable of satisfying the conditions on duty-bearing that were suggested above by Blomberg and by Wringer.

In my view, attributions of duties to non-collective groups should be reinterpreted as attributions of coordination duties held by each member of the group. Importantly, this reinterpretation means that we reject the intentions of the person who attributed a duty to the group. If we engage in this reinterpretation and reject the intentions of the attributor, then our reinterpretation is no tension with the idea that a group-level duty requires a group-level decision-making procedure. So there is no need to resolve this possible tension.

That said, it is worth commenting on Björnsson’s proposed resolution of the tension between the attributor’s intentions and my suggested reinterpretation. He suggests that a non-collective group bears a duty when it’s able to produce an important outcome (at least, able to produce that outcome given that all members perform relevant coordination duties). This allows us to

hold onto the attributor's intentions (to attribute a duty to *the group*), alongside the individual-level coordination duties proposed in my reinterpretation. Björnsson suggests that such group-level duties might be derived from 'basic' duties, where basic duties can be held only by agents. I would prefer not to uphold the attributor's intentions at the cost of philosophical defensibility. I do not view duties of non-collective groups as philosophically defensible. So I would reject this proposal, independently of Björnsson's reasons for rejecting it.

Björnsson gives two final examples: *Beyond Reach* and *Ferry Ride* (this volume, p. **). In *Beyond Reach*, my account says you have "a duty to do what [you] can (within reasonable cost) to make it reasonable for the others to believe" that, roughly, you'll do your part if they'll do theirs (Collins 2019, 117). Björnsson finds this objective duty implausible. I suspect this rides on how we interpret 'reasonable cost.' I gave no interpretation of this in the book. If the others are utterly recalcitrant, then perhaps the only reasonable cost (including opportunity cost) of you signalling your conditional willingness is *zero*. That is, if there is any cost to your signalling, then the 'reasonable cost' rider is not met, and my account is silent on whether you have any duty. But I am happy to say that you have this objective duty if there is literally no cost to you signalling. There is value, I think, in the others knowing that you are willing, even if they are completely unwilling. It means they know that the failure is theirs alone.

As for *Ferry Ride*, my account implies coordination duties, where those duties require that the individuals affect one another's environment with a view to producing the outcome of the ferry not capsizing. As above, if the individuals do not make any decisions related to this, then I deny that they have *fulfilled* these coordination duties—though they do bear those duties, and they perform them by sitting quietly after a hard day's work. Here, "with a view to" was an unfortunate choice of words on my part. It suggests a conscious choice is required. This reflects my focus on duties to act, rather than duties to omit. A better formulation would have required

agents to affect one another's environment "so as to" produce the outcome of the ferry not capsizing. In such cases, our duties can require no thought at all.

To close, I would again like to thank the commentators (and this journal) for giving me this opportunity to clarify and supplement aspects of the book. I am certain my comments here will not convince everyone and I look forward to ongoing discussion.

References

- Collins, Stephanie. 2019. *Group Duties: Their Existence and Their Implications for Individuals*. Oxford University Press.
- Hakli, Raul, Kaarlo Miller, and Raimo Tuomela. 2010. "Two kinds of we-reasoning." *Economics and Philosophy* 26(3), 291-320.
- Schwenkenbecher, Anne. 2019. "Collective moral obligations: 'we-reasoning' and the perspective of the deliberating agent." *The Monist* 102(2), 151-171.
- Shoemaker, David. 2011. "Attributability, Answerability, and Accountability: Toward a Wider Theory of Moral Responsibility." *Ethics* 121(3), 602-632.
- Shoemaker, David. 2015. *Responsibility from the Margins*. Oxford University Press.
- Wringe, Bill. 2016. "Collective Obligations: Their Existence, Their Explanatory Power, and Their Supervenience on the Obligations of Individuals." *European Journal of Philosophy* 24, 472-497.