

Group duties without decision-making procedures

Though short, Stephanie Collins' (2019) *Group Duties* combines vigor with breadth, offering interesting new arguments while bringing together numerous issues regarding group duties that have hitherto been treated separately. My commentary focuses on two particularly original and central theses of the book, named here for easy reference:

DUTY REQUIRES PROCEDURE: For a group to bear duties, it must be united under a group-level decision-making procedure. (Ch. 2–3)

REINTERPRETATION: Attributions of duties to ϕ to groups without decision-making procedures should be understood as attributions of “coordination duties” to each member of the group, duties to take steps responsive to the others with a view to the group's ϕ -ing or express willingness to do so. (Ch. 4)

I will argue that, at least in its current form, Collins' defense of these theses fails.

Start with some terminology. Collins distinguishes *collectives* from *coalitions* and *combinations* of agents. Collectives—henceforth “Collectives” with a capital “C”—are constituted by agents that are united under a rationally operated group-level decision-making procedure that it responsive to moral considerations. Coalitions are constituted by agents that share a goal but lack such a uniting procedure. Combinations, finally, are constituted by members that lack both shared goal and uniting procedure. A *decision-making procedure* “takes in beliefs and preferences, and processes them to produce decisions”. Members are *united* under a group level-procedure if (i) they are at least tacitly presumptively committed to abide by it, (ii) the procedure is operationally distinct from the personal decision-making procedure of any one member, and (iii) the enactment of the decisions requires actions on parts of the members. This might sound complex and advanced, and so it is in the case of vast and intricately organized entities like the UK. But Collectives also include small groups of people that have come together to do something on the initiative of one member who is distributing tasks to the others, or a small group of friends coordinating some of their actions through conversation-based consensus, a procedure that members have committed to by participation in the conversation.

Collins' basic argument for DUTY REQUIRES PROCEDURE appeals to the following two claims, again named here for easy reference:

DUTY REQUIRES ABILITY: For a group to bear duties, it must be able to make decisions.

ABILITY REQUIRES PROCEDURE: For a group to be able to make decisions, it must be united under a group-level decision-making procedure.

I'll return to Collins' reasons for accepting ABILITY REQUIRES PROCEDURE. My focus will first be on DUTY REQUIRES ABILITY, for which Collins offers the following

MORAL WORTH ARGUMENT (pp. 86–7)

- (1) If an entity fulfils a duty [in contrast to merely acting in accordance with it], that entity acquires moral worth [i.e. merits moral praise].
- (2) If an entity acquires moral worth, it has performed in a way that has a psychological explanation that accords with the performance's moral justification.

- (3) If an entity's performance has a psychological explanation that accords with its moral justification, then the entity has made a decision.

Therefore,

- (4) When an entity fulfils a duty, that entity has made a decision.
 (5) For an entity to bear a duty, the entity must have the ability to fulfil a duty.

Therefore,

- (6) If an entity can bear duties, then that entity has the ability to make a decision.

My primary focus will be on premise (3). Premise (1) is meant to be true by definition (p. 88). Something like premise (2) is part of sophisticated quality of will accounts of praiseworthiness (e.g. Arpaly and Schroeder 2014), and I have myself defended a version of it: the object of praise must be explained in a normal way by the agent's concern for the relevant values (e.g. Björnsson 2017; Björnsson and Persson 2012). Though premises (1) and (2) both seem fine, I don't see that much is gained by introducing considerations of moral worth. For the following, which would do the job of (1) and (2), itself looks like a tautology:

- (1*) If an entity fulfils a duty [in contrast to merely acting in accordance with it], it has performed in a way that has a psychological explanation that accords with the performance's moral justification.

In either case, the strength of Collins' case rests on premises (3) and (5). The latter seems plausible indeed if we understand fulfilling a duty as performing what is one's duty as the normal upshot of responsiveness to the reasons grounding the duty. Something like it follows from the idea that duties are what moral considerations demand that we do and the idea that moral demands apply only to those who have the *ability* to respond to the demands. As it stands, the premise is even weaker than that, as it does not even require that it be true of each duty that its bearer is capable of fulfilling it: all that is required is the general ability. In the end, the plausibility of (5) will depend on how we understand the relevant kind of ability. This issue becomes pressing when we ask whether ABILITY REQUIRES PROCEDURE: whether a group's ability to decide, in whatever sense is required for fulfilling a duty, requires that it is already united under a group-level decision-making procedure. I will return to that issue after discussing premise (3).

Must an entity have made a decision when its performance has a psychological explanation that accords with its moral justification?

The problem with relying on premise (3), I will argue, is this: To be uncontroversially acceptable, (3) must be understood in a way that undermines the overall argument for DUTY REQUIRES PROCEDURE. The following two cases begin to bring out the problem:

Noticing and Wondering: In spite of being deeply involved in a crucial step of her pet project, Fara notices that Jakob is in distress, and begins wondering how she can help.

Refraining: Though destitute for no fault of her own, Fara never even considers keeping the wallet a stranger forgets on the seat next to her. Instead she goes beyond the call of duty in making sure that it gets back to its owner.

Consider three of Fara's "performances": noticing Jakob's need, wondering how to help, and refraining from taking the contents of the wallet. Let us assume, as seems plausible, that these performances have psychological explanations that accord with their moral justifications: they are the normal upshots of Fara's extraordinary concern with the values at stake. Even so, it seems false that Fara must have *made a decision* to notice the distress or wonder about how to help, or decided not to take the money, in any ordinary sense of "making a decision".

There are four ways to defend (3) in light of cases like these.

The first, and to my mind least plausible, is to argue that, in spite of appearances, Fara did make the decisions in question, in an ordinary sense of "making a decision".

The second is to deny that Fara's noticing, wondering, and refraining qualify as *performances*. This has some plausibility with respect to noticing and wondering, if we assume that performances, in the relevant sense, must be possible contents of one's obligations: though I think that we can have obligations to notice and wonder about things, these are not uncontroversial or paradigmatic examples of obligations. But it is implausible for Fara's not keeping the wallet: not keeping someone else's wallet is a paradigm case of an obligation.

The third is to deny that Fara's performances have psychological explanations that accord with their moral justifications. At least on the face of it, that seems implausible. (It also seems to be in direct tension with (2), as Fara seems to deserve moral praise for her performances. That might not be a problem if, as I have suggested, (1) and (2) can be replaced by (1*). But *Refraining* also puts pressure on (1*), as it seems that Fara fulfilled her duty not to take the stranger's wallet, and didn't just act in accordance with it.)

The fourth, finally, is to understand "made a decision" in (3) in a weak enough way to make it plausible that Fara did make decisions leading him to notice Jakob's distress, wonder whether he could do something about it, or refrain from taking the wallet. Some of what Collins says might suggest such a very weak reading: her notion of "decision" is explicitly a "thin" one, allowing for decisions that are "sub-personal, offline, tacit, or unconscious" (p. 90). One might think that in this thin sense, Fara's performances could be the upshot of sub-personal decisions to direct cognitive resources in the right direction. However, in discussing what does and does not involve decisions, Collins says:

"if I am walking with my friend—engrossed in philosophical conversation—and she almost steps in front of an oncoming bus, I might jerk out my hand to stop her. This is a performance that accords with duty: my behaviour matches the behaviour that a duty demands. But I have not made any decision." (90)

It is not clear to me why this action should not be seen as the upshot of a sub-personal decision guided by a tacit understanding of the value at stake. This suggests that Collins intends something stronger by "sub-personal" or "tacit" decisions. But whatever Collins' intentions, what is important for the overall argument is this: If we understand "making a decision" in a fairly thick sense, (3) seems implausible. If we instead understand it in a sense thin enough to apply in *Noticing and Wondering* and *Refraining*, it is unclear what ground we have for accepting ABILITY REQUIRES PROCEDURE. Why assume that a group's ability to make a decision in such a weak sense requires that it is united by a decision-making procedure? Consider:

Rescue: A swimmer needs to be rescued. In spite of the offshore current and the risks involved, one of the strangers on the beach starts making her way towards the swimmer and others

follow suit, coordinating with their immediate neighbors and spontaneously creating a human chain. Together they manage to save the swimmer and getting everyone safely back on the beach. But they never committed to a uniting group-level decision-making procedure.¹

In *Rescue*, the group effort was guided by a series of sub-group decisions tracking the justification of the rescue, namely the group's ability to rescue the swimmer and the importance of doing so. On a *very* thin understanding of what it is to make and act on decisions, then, it seems that this group did exactly that.²

In light of all this, I suspect that (3) will either be given a reading on which it is implausible or controversial, or be interpreted in a way that undermines other premises in the argument for DUTY REQUIRES PROCEDURE. At least it seems that more needs to be said to support the argument.

Does a group's ability to make decisions require a decision-making procedure?

Suppose that, in spite of the worries raised about premise (3), we could establish DUTY REQUIRES ABILITY: that for an entity to have duties, it must have the ability to make decisions. The question would then be why we should think that a group's ability to make decisions requires that it is united under a group-level decision-making procedure, i.e. why we should accept ABILITY REQUIRES PROCEDURE.

To understand the latter claim, we need to understand the relevant notion of "ability". The problem is that there is at least *some* sense in which some strangers on the beach are able to decide to rescue a swimmer under favorable circumstances. Indeed, we might think that Collins has explained how easy it is for them to do this: all it takes is that one of them suggests a way to get organized and for the others to agree (pp. 108–109). Clearly they can do this! Since ABILITY IMPLIES PROCEDURE seems to fly in the face of this plausible claim, Collins needs to spell out a sense of "ability" given which the group plausibly needs to be already united under a rationally operated decision-making procedure in order to have such an ability. Moreover, she needs to explain why having an ability to fulfil a duty or make a decision *in that sense of ability* is required for bearing obligations.

In earlier writing, Collins had denied that non-agential groups have abilities. With a notion of ability strong enough to support that claim, any ability to decide would of course be ruled out. As noted, however, we naturally attribute abilities to groups that lack decision-making procedures, and Collins needs a notion of ability for non-agential groups for what she wants to say about the obligations of members of such groups. In the book, she thus goes to some lengths defending the following notion (here: "Collins-ability"):

COLLINS-ABILITY: A non-agential group is able to produce an outcome X at a time t just in case (1) each member has an individual ability at t to perform actions that contribute to X; and (2) given that enough members exercise the abilities in (1) at t, each will do their contributory part of a pattern of behaviours that will robustly secure X in the absence of defeaters. (P. 71)³

¹ Cf. Collins' more elaborate *Responsive Rescue* (102).

² It also seems that the group in *Rescue* deserves praise for rescuing the swimmer, putting pressure on the idea that praise requires a psychological explanation that involves making a decision in anything but a very thin sense.

³ I take it that "t" signifies the time at which the ability is had, not the time at which the outcome is produced or the time of the outcome.

It is clear how the group of people on the beach is Collins-able to produce the rescue of the swimmer, but wouldn't be in different scenarios. The problem is that it is equally clear how the group is Collins-able to produce a group decision to rescue the swimmer that is then acted on. Admittedly, a group's having the Collins-ability at t to *produce the outcome that the group makes a decision* might not be logically equivalent to the group's having the ability at t to *make a decision*. But we still need a principled well-motivated reason for saying that a group's Collins-ability to produce and act on a group decision isn't sufficient for bearing duties, whereas Collectives have what it takes.

I am not entirely clear about Collins reasoning at this point, but her discussion in other parts of the book suggests that she might have either of two reasons in mind.

First, she might think that (i) duty-bearing requires the ability to *immediately* make a decision and that (ii) this ability can only be had by groups that already have decision-procedures in place (cf. what she says on pp. 46–50). But both parts of this suggestion are problematic. In many cases where we fulfil duties, we first notice the duty, then take it into account in deliberation, and then decide to fulfil it. At least as we normally understand matters, the duty precedes our recognition of it, and our decision to act on it comes even later. Given that the normal fulfilment of duties does not involve the immediate decision to fulfil it, it is unclear why duty-bearing should require such an immediate ability. Moreover, if such an ability is required, this would not plausibly draw the line between groups that are united under decision-making procedures and those that are not. After all, the enactment of decision-making procedures can be very time-consuming. Some parliamentary procedures extend over years, involving elaborate rules for how individuals can submit proposals to be decided upon and requiring proposals to be confirmed before and after general elections, and even the mundane procedure of conversational consensus can be frustratingly protracted. In comparison, the route by which the strangers on the beach might reach a decision to rescue the swimmer might be mercifully swift and, it seems, no less immediate.

Second, Collins might think that decision-making procedures provide a particularly *robust* way of securing the relevant outcome. In general, an appeal to robustness seems relevant for distinctions between groups that have and do not have abilities to decide. Whether we like Collins particular way of incorporating a robustness requirement in her analysis of abilities, the ability to produce an outcome that is relevant for obligations does seem to require more than some greater than zero chance that one would do so. Moreover, the suggestion that decision-making procedures provide robustness seems sensible: when a group is united under a decision-making procedure, recall, the members are defeasibly committed to follow the procedure and doing their parts in enacting the decisions. It would also be in line with Collins' argument that duties of Collectives cannot be reduced to duties of their members because Collectives have mechanisms for "reliably producing multilateralism among members" (176–9). But like the ability to immediately decide, robustness does not seem unique to Collectives. When a need for group action to achieve a goal is clear enough to all involved and there are no obstacles to coordinating actions accordingly, robustness should not be an issue even in the absence of prior commitments. Conversely, a decision-making procedure might be ill-suited to produce plans that robustly secure outcomes: think of procedures for forming governments in some countries, or commitments to conversational consensus in groups with suitably convergent interests. So an appeal to robustness does not seem to draw the line where Collins wants it. Some other reason is needed.

Might reasons-responsiveness require decision-making procedures?

In light of the preceding discussion, I remain unconvinced that there is a way of understanding abilities to make decisions that makes it plausible both that having a duty requires such an ability and that only groups unified under decision-making procedures have this ability. I have also suggested that premise (3) is problematic given a reading that doesn't undermine other parts of Collins' argument for DUTY REQUIRES PROCEDURE.

Perhaps, though, other arguments for that conclusion might be more successful. One sort of argument that I find comparatively promising starts from the premise that

REASONS-RESPONSIVENESS: An entity can only have an obligation if it is capable of responding coherently to the reasons grounding the obligation so as to discharge it.

Recall *Rescue*. Though each agent in that case responds to the group's ability to rescue the swimmer and the importance that they do so, we might ask if *the group of agents* is responding to those reasons, as opposed to mere parts or constituents of the group. Perhaps, in order for the responses of members to count as responses *of the group*, it must be the role of the individuals, qua members and not merely as tools of the group, to guide the group by these responses. And this, one might try to argue, requires a mutual commitment with respect to how members are to behave in response to the reasons-guided inputs of other members: something like a unifying decision-making procedure, as Collins understand these.⁴

Group duties without REINTERPRETATION

Suppose that the MORAL WORTH ARGUMENT or an appeal to REASONS-RESPONSIVENESS made it plausible that DUTY REQUIRES PROCEDURE: that only groups united under decision-making procedures bear duties. What, then, should we think about frequent attributions of duties to other groups, as when we naturally think that it is the duty of the strangers on the beach to save the swimmer, or the duty of rich countries to prevent catastrophic climate change? As already mentioned, Collins thinks that,

REINTERPRETATION: Attributions of duties to ϕ to groups without decision-making procedures should be understood as attributions of "coordination duties" to each member of the group, duties to take steps responsive to the others with a view to the group's ϕ -ing or express willingness to do so.⁵

Suppose that this is how the attributions should be understood. Then these attributions are correct when and only when the members in question actually have the relevant coordination duties. And then it seems hard to deny that the groups have duties, *in the sense intended by these attributions*, under exactly those conditions. But duties *of the sort intended in these attributions* are held by groups that lack group-level decision-making procedures. The REINTERPRETATION suggestion thus seems to be in fairly direct tension with DUTY REQUIRES PROCEDURE.

⁴ Even if the only groups that bear duties are Collectives, it doesn't follow that all Collectives bear duties. One might think that many such groups lack the sort of persistent independent existence and interests that make them parties of potential social contracts that contractualist accounts of obligations require, or lack the sort of constitutive aim of rational self-determination that constitutivist accounts require (cf. Björnsson Björnsson 2020a).

⁵ Collins employs the following notion of *responsiveness*: An "agent, A, is responsive to another, B, just in case A acts upon B with a view to B responding to the reasons or duties that (A believes) B holds." (p 98).

I think that the right resolution of this tension begins by noticing that group duties, so understood, are derived from individual duties. This opens for the possibility that *basic* moral duties are borne only by suitably unified moral agents, whereas derived moral duties can be borne by groups constituted by such moral agents even when these groups lack the requisite unity (Björnsson 2020a).

A different response would be to deny that the attributions are well understood as attributions of any collective properties at all. If the individual coordination duties were entirely individual affairs, that would be a plausible response. But they are not, as Collins understands them. I'll ride roughshod over the many intricate details of her full account (pp. 116–7), but three things in particular are worth noting. The first is that coordination duties of members depend on whether the group is able to produce the outcome that it ϕ s, where this ability is a property of the group, not necessarily of any member. The second is that the ability would be exercised if all the members of the group were to do their coordination duties. The third is that coordination duties depend on the importance of the outcome in question, an outcome that the group is best placed to produce. Given that all these conditions involve the group, the conditions under which the members of a group have the relevant coordination duties are themselves irreducibly collective. Simplifying quite a bit, the conditions would be:

REINTERPRETATION DUTY: A group without a decision-making procedure has a duty to ϕ if and only if the group's ϕ -ing is important and the group would exercise its ability to ϕ should the members of the group discharge the coordination duties held based on the importance of the group's ϕ -ing and its ability to do so.

So understood, group duties would not only obtain in virtue of irreducibly collective facts, but also in virtue of facts that are normatively significant. For it matters that something important would be produced as a matter of exercised ability if a group of individual agents just lived up to what can be morally demanded of them, without requiring that they go beyond the call of these basic demands.

I think that REINTERPRETATION DUTY picks out an irreducibly collective and normatively significant relationship, and take it to be compatible with the idea that DUTY REQUIRES PROCEDURE holds for basic duties. But REINTERPRETATION DUTY fails to capture what many of us have in mind when we attribute duties to groups without decision-making procedures. I'll mention two problems.

The first problem with REINTERPRETATION DUTY is that it seems to get negative group obligations wrong. Consider:

Ferry ride: A small ferry would capsize if many of its passengers moved laterally in sync. On this ride, as on most, almost all passengers are sitting quietly after a hard day's work. (Björnsson 2020)

Given that the passengers could easily capsize the ferry if they all wanted to and given that it would be very bad if they did, they plausibly have a duty not to. REINTERPRETATION DUTY fails to explain this: because there is no danger in the offing, no passenger has any obligation to take steps responsive to the others with a view to the group's not capsizing the ferry or to signal their willingness to do so.

The second problem concerns cases where individual attempts at coordination would be futile because the other members of the group would refuse to join. Consider:

Beyond Reach: Two things are true: (1) Kim, Jo, and yourself can rescue a drowning swimmer if, but only if, you all join forces, but neither of you has any evidence for this. (2) Kim and Jo would refuse to help no matter what you would do, but you have no evidence for this.

It is natural to think that in *Beyond Reach*, the group has an objective duty to rescue the swimmer, grounded in its ability to do so; the unwillingness of Kim and Jo does not seem to undermine that duty. (The duty would be objective rather than evidence-relative, as the group lack reasons to think that they have the ability. But this is as it should be, as Collins interest is with objective duties.) For REINTERPRETATION DUTY to yield that group duty, each group member must have corresponding coordination duties: duties to take steps responsive to the others with a view to rescuing the swimmer or, in case this is pointless unless the others do likewise, to signal their willingness to do so (pp. 116–7). But it is implausible that you have any such duty in *Beyond Reach*. Because the exercise of the group’s ability was beyond your reach, you are not only off the hook in the sense that your failure didn’t reflect badly on you. You are off the hook in that there was no significant objective ground for signaling, as indicated by the fact that if you later became aware of the two truths of *Beyond Reach*, it would be odd to regret not having signaled. Contrast this with a case where, contrary to your evidence, your signaling could have contributed to a successful saving attempt, bringing about coordination with the others. Now regret would seem appropriate. (I follow Collins in taking regret to indicate failure to discharge an objective obligation.)

I have suggested that *Beyond Reach* is a problem for REINTERPRETATION DUTY, because it seems that (a) the group has an objective duty to rescue the swimmer while (b) you lack the coordination duty required for the group duty. But (b) is equally a problem for Collins’ account of coordination duties. For according to that account, such duties are not predicated on members’ individual abilities to significantly contribute to the outcome given dispositions of other members, but only on the group’s ability to bring about an important outcome. In *Beyond Reach*, then, Collins would say that you have a duty to signal your willingness to take responsive steps with a view to rescuing the swimmer. But this, again, seems implausible.

Because of problems highlighted by *Ferry Ride* and *Beyond Reach*, I think that we should reject REINTERPRETATION DUTY and, with it, REINTERPRETATION. Interestingly, though, REINTERPRETATION DUTY is structurally very similar to some existing accounts of shared duties or obligations. Here is my own favorite version, which avoids the problems mentioned:

OBLIGATION: For a group to have as its moral obligation to ϕ is for its ϕ -ing to be (i) morally important and (ii) ensured, in a normal way, if members care as can be morally demanded of them (Björnsson 2014; 2020a; b; cf. Aas 2015; Schwenkenbecher 2019).⁶

Unlike REINTERPRETATION DUTY, this proposal is meant to be relevant for attributions of duties to both pluralities and individual agents, where the latter are conceived of as groups of one. If it can capture a wider range of attributions of obligations, this is itself a strike in its favor. And unlike REINTERPRETATION DUTY, which involves demands that members take steps responsive to others with a view to the outcome or expresses responsiveness, OBLIGATION takes group duties to depend on demands that individuals *care* about certain things in certain ways. To care about something, in the relevant sense, is to be disposed to notice factors that affect how well it goes with the object of

⁶ Note that obligation talk is multiply ambiguous. This particular formulation is meant to capture one species of subjective obligations. For related ways of understanding other species, see Björnsson 2020b.

care, and to invest resources—cognitive, behavioral, social, material—based on such information into making it go better. Under the right circumstances, caring about various values thus means acting in certain ways, ways that can ensure outcomes and might include the sort of responsiveness that coordination duties call for. But because OBLIGATION doesn't predicate group duties on individual duties to *do* anything, it can easily account for the negative obligation in *Ferry Ride*: the passengers are plausibly required to care to a certain extent about the values at stake, and if they did, this would ensure that they didn't capsize the ferry.⁷ And because demands to care about the values involved are not predicated on the group's actual abilities, it avoids the sort of problems that I raised for Collins' account of coordination duties. But like REINTERPRETATION DUTY, it captures an irreducibly collective and normatively significant relation: it matters that something important would be ensured in a normal way if a group of individual agents just lived up to what can be morally demanded of them, without requiring that they go beyond the call of these basic demands.

As is customary, my comments here have been critical: I have argued that without further supplementation, the MORAL WORTH ARGUMENT fails to establish that DUTY REQUIRES ABILITY. I have also suggested that we need to hear more to accept that ABILITY REQUIRES PROCEDURE, but that there might be other reasons to think that DUTY REQUIRES PROCEDURE holds for basic moral demands. Finally, I have suggested that there are reasons to reject REINTERPRETATION, and briefly explained why a structurally similar account looks more promising. But I want to end where I started, by stressing that this is a vigorous and very useful contribution to the field, and one that I have both learned from and enjoyed engaging with. I expect more from further engagement.

Bibliography

- Aas, Sean 2015: 'Distributing Collective Obligation'. *Journal of Ethics and Social Philosophy*, 9, pp. 1–23.
- Arpaly, Nomy and Schroeder, Timothy 2014: *In Praise of Desire*. Oxford University Press.
- Björnsson, Gunnar 2014: 'Essentially Shared Obligations'. *Midwest Studies in Philosophy*, 38, pp. 103–20.
- Björnsson, Gunnar 2017: 'Explaining Away Epistemic Skepticism About Culpability'. In *Oxford Studies in Agency and Responsibility*. Shoemaker, David (ed) Oxford University Press pp. 141–64.
- Björnsson, Gunnar 2020a: 'Collective Responsibility and Collective Obligations without Collective Agents'. In *The Routledge Handbook of Collective Responsibility*. Bazargan-Forward, Saba and Tollefsen, Deborah (eds) New York: Routledge.
- Björnsson, Gunnar 2020b: 'Individual and Shared Obligations: In Defense of the Activist's Perspective'. In *Philosophy and Climate Change*. Budolfson, Mark, McPherson, Tristram and Plunkett, David (eds) Oxford University Press.
- Björnsson, Gunnar and Persson, Karl 2012: 'The Explanatory Component of Moral Responsibility'. *Noûs*, 46, pp. 326–54.
- Collins, Stephanie 2019: *Group Duties: Their Existence and Their Implications for Individuals*. Oxford University Press.
- Schwenkenbecher, Anne 2019: 'Collective Moral Obligations: 'We-Reasoning' and the Perspective of the Deliberating Agent'. *The Monist*, 102.

⁷ For X to *ensure* Y is for Y to obtain in all relevant possibilities and to do so because of X in at least one of these possibilities.