

***Group Duties: Their Existence and Their Implications for Individuals***, by Stephanie Collins. Oxford: Oxford University Press, 2019. Pp. vii + 218.

Sometimes one should start to judge a book by its cover – when the cover is as apt as the one for Stephanie Collins’s book *Group Duties: Their Existence and Their Implications for Individuals*. In a large aerial shot of a coastline the waves of the sea are crashing onto a beach. Dotted around the beach are swimmers, beach walkers, lifeguards, and sunbathers. Most of these individuals keep a respectful distance to each other, minding their own business. But, looking at the scene more closely, we realize that this could change quickly. Where the water meets the beach, yellow lifeguard surfboards lie ready for use. If a swimmer gets into trouble and begins to drown, a group of individuals will have to act together to save the stricken aquatic athlete, as no single individual might alone be able to achieve that goal. But when is it correct to say that there is a group on which such a duty falls? This is the question Stephanie Collins aims to answer in her clear and tightly argued essay.

In chapter 1, Collins presents us with a tripartite model of groups. Groups can be *combinations*, *coalitions*, or *collectives*. *Combinations* are random sets of people, e.g. a group of people sitting on the beach who don’t know each other. *Coalitions* are sets of people with a shared goal and a disposition to react responsively to each other in support of that goal. For example, a group of beach goers may have the joint goal of setting up a beach volleyball field together. As one sets up one post, the other sets up another in the right distance, a third starts rigging up the net depending on the progress of the first two, and so on.

*Collectives* are the most complex form of groups. In a first gloss, Collins says that a ‘collective is constituted by agents that are united under a rationally operated group-level decision-making procedure that can attend to moral considerations’ (p. 11). This characterization prepares the ground for a more complex set of ‘typically’ necessary conditions stated in chapter 6 (‘typically’ because Collins is willing to concede that there are some conceivable collectives ‘constituted by agents that are quite different from human agents’, p. 156-7.). Stated in slightly simplified fashion, a collective has at least (i) a *procedure* for making decisions that tends to produce rational choices responding to moral considerations; (ii) *group goals* determined by the procedure; (iii) *roles* for the individuals constituting the collective assigned by the procedure; (iv) such that each member has some *influence* over group decisions; (v) and the members have a *commitment to comply* with group decisions, and this is common knowledge amongst them (cf. Pettit and List 2011).

The central claim of Collins's book is that collectives can have group duties because they are agents. By contrast, combinations and coalitions cannot have any duties *as group* because they are not agents. To support this claim, Collins reviews the arguments in favour of duties for combinations and coalitions and finds that they lead, at best, to a stalemate between proponents and opponents (chapter 2). In chapter 3, Collins presents and defends the 'Decision Argument' against group duties for combinations and coalitions: that to bear a duty, a group has to be an agent. Chapter 4 argues in favour of a duty for individuals to coordinate when coordinated action can promote moral value. Importantly for Collins, such coordination duties are not group duties. Chapter 5 considers the divergence between combinations and coalitions in coordination duties. Because members of coalitions share the same goal, more directly action-guiding coordination duties apply to them, while the duties for members of combinations are often more loosely specified. Chapter 6 characterizes collectives in more depth and then sets out why the duties of collectives are not reducible to individual duties. Notwithstanding this last point, the duties of collectives also entail individual membership duties, as chapter 7 explains.

In this review I zoom in on three points in Collins's analysis that I find particularly important or controversial: her analysis of the abilities of combinations and coalitions, the arguments for the irreducibility of group abilities and duties, and the conditions under which individuals coordination duties arise.

## **1. Group Ability**

Do non-agent groups have abilities? The answer to that question matters for Collins's project because ought-implies-can type arguments suggest that ability is necessary for duty. The argument in favour of assigning abilities to combinations and coalitions turns on the question whether such group abilities can be reduced to individual abilities, or whether a non-reductive account is preferable for reasons of explanatory power, parsimony, or theoretical elegance. To tackle that question, we first need an analysis of group ability. After assessing some possible approaches, Collins proposes this definition:

Ability for non-agential groups: 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)

Put roughly, group ability to X requires, first, a decision to contribute by a sufficient number of group members and, second, contributions that together result in a successful pattern, such that the pattern brings about X. In a slogan, then, what is needed is sufficient *participation* in suitable *choreography*.

I agree with the thrust of this definition but suggest it can be sharpened by thinking of the interaction of agents as a game. One such game is the El Farol Bar Problem (Arthur 1994): 100 players must decide simultaneously (and without communication) whether they go to the El Farol Bar tonight. The bar is small and therefore attendees will have a great party if 60 or fewer players go and an unpleasant party if more than 60 go. Unfortunately, the way to the bar is too long to turn around once one has arrived, so that one cannot change one's mind after the initial decision. Each player prefers attending a great party over staying at home over attending an unpleasant party. Is the group able to ensure that there will be a great party at the El Farol Bar tonight?

Clearly, each player has the ability to perform actions that contribute to having a great party, namely either staying at home or going to the party. The problem is with the pattern of actions, and whether a successful pattern can be brought about robustly. In other words, what matters for 'solving' the El Farol Bar Problem is whether there is adequate choreography guiding the individuals (Pinkert 2014).

Collins's definition of non-agent group ability suggests that the group is able if and only if each player is able to do their part of a pattern that robustly secures great parties. In our example, there are many candidate patterns, namely all possible arrangements such that a subset of 60 or fewer players go to El Farol while the rest stays at home. However, the fact that there are many suitable patterns leaves open the question whether the group is able to *robustly* produce successful patterns.

Here is one uncharitable reading of Collins. She might be suggesting that non-agent groups are able to X if and only if the group is able to reliably act in a pattern that causes X, and then actually does so. However, defining group ability in terms of the ability to act in a pattern would be circular. Fortunately, there is a more charitable reading of Collins's definition: a group is able to perform X if and only if, given enough members perform contributory actions, they will robustly act in patterns that succeed in bringing about X.

Read in this way, the definition is not circular. However, one might complain that this definition is not as illuminating as one would have hoped because it does not tell us under which circumstances non-agent groups will reliably act in suitable patterns. In the El Farol Bar Problem, for example, the definition does not inform us whether the individuals will overcome their coordination problem or whether they will end up in an overcrowded bar. Fortunately, the analysis of coordination problems can enlighten us further (e.g., Gintis 2000). For example, whether the El Farol Bar Problem can lead to great parties depends on whether there is a suitable, reliable mechanism of choreography. Looking at a one-shot game El Farol Problem, there isn't such a reliable mechanism. It is therefore at best a matter of luck whether the group ends up having a great party for 60, so it is not a group ability.

In a one-shot El Farol Bar Problem individuals act simultaneously and without communication. They can therefore not be responsive to each other; they cannot adapt their strategy in light of what other individuals do. And to make matters worse, their individual interests create a collective action problem, as the attempt to satisfy everyone's preferences leads to a suboptimal outcome. The situation is different, however, when the game is sequential and the strategies of players become contingent on what other players do. Consider, for example, a situation where someone gets trapped under a car while changing a tire. A group of bystanders (a combination) might be able to help by lifting the car and moving it off the trapped person. Is the group able to do so? That depends on whether enough individuals have the right individual strategies: strategies that ensure the right pattern of action come about when the different individual strategies are played out interactively. In the car-lifting example, a useful strategy might look like this:

(i) help if not enough people are already helping, and if there are, remain on standby to help if others fail but make sure you are not in the way; (ii) when helping, position yourself at the car so that you don't inhibit others and make sure the lifting power is approximately symmetrically distributed; (iii) start lifting and encourage others to do so; (iv) adjust position and lifting power to keep the car balanced; (v) pull the car slowly in the direction of the nearest safe location; (vi) set the car down if the nearest safe location has been reached.

If all individuals separately follow a strategy like this, it is highly likely that a suitable number of people will lift (not too few, not too many), that they lift in suitable positions, and move the car in a good location to help the victim. Such a group is able to lift a car off a person trapped underneath.

The example reveals an interesting aspect of group ability: whether the group can follow a reliable choreography leading to success depends on the individual strategies and their interaction. This, in turn, depends on the structure of the interaction: whether it is sequential, whether there is time for learning, whether players can communicate, and what players know. The robustness of success required to establish ability is therefore a property of the set of individual strategies and the game form describing the rules of interaction. Therefore, when Collins refers to the robustness of the pattern ('each will do their contributory part of a pattern of behaviours that will robustly secure X'), she does not mean that all individuals need to robustly act in one maximally specified way. Rather, the pattern can be 'multiply realized' (p. 73). Perhaps a clearer way of putting this would be: what needs to be robust is the production of outcome X with the interaction of individual strategies, given the game form. The detailed patterns of actions might not be robust. What is robust is the production of the outcome, and what makes this production of outcome robust is the ability of different individuals to respond in suitable ways to the contingent behaviour of others, as well as other contingent events.

A plausible refinement of Collins's definition of non-agent group ability could make use of thinking in terms of the strategies that agents use, given the game form G (the rules of the interaction set out in terms of available actions for the individuals, possibly incorporating 'moves of nature', i.e. non-agential events, e.g., Osborne 2004). Suppose that for each individual there are non-responsive strategies which will not support the production of X. There is also exactly one responsive strategy which tends to lead to X, provided that enough others are responsive in the right way. (This restriction to exactly one responsive strategy rules out a coordination problem that were to arise if there were several, mutually incompatible ways to be responsive; I put this complication to the side.) Given these assumptions, we can then state a revised definition:

*Ability for non-agential groups\**: A non-agential group is able to produce an outcome X in a game of form G at time  $t$  just in case (1) each member has an individual ability to perform their responsive strategy at  $t$ ; and (2) if enough members enact their responsive strategies, these strategies will together robustly secure X in G.

The important improvement provided by this definition is that it links the property of robustness explicitly to the tuple of all individual strategies *together*. This shows that group ability does not only depend on enough people being responsive but also on them being *responsive to each other* in the right way. Ability is therefore a property of the system

constituted by the individuals and their strategies – and only the system, not the individuals can reliably perform the group action.

Groups of well-meaning, responsive people often fail to bring about desired outcomes because they do not manage to coordinate their actions; the El Farol Bar Problem is one stylized example. In a similar vein, the group trying to lift the car might also fail if the responsive individuals are too eager to help and inhibit each other's actions. If the individual strategies are such that the car lifters are not likely to coordinate their actions in an effective way, they are – as a group – not able to perform the rescue action.

The definition also foreshadows Collins's view about the reduction of abilities. Ability for non-agent groups is not straightforwardly reducible to individual ability because success of a tuple of individual strategies depends on how the strategies mesh and interact in suitable ways so that individual deviations or unanticipated changes do not derail the overall performance of the group. I will now turn Collins's unusual stance on the non-reducibility of group ability in more detail.

## **2. Are Group Abilities and Duties Reducible?**

The starting point for Collins's argument against group duties for combinations and coalitions is the

Agency Principle. Groups that are not agents cannot bear duties. (p. 60)

To be an agent, a group has to be able to make decisions (p. 44). And since the only groups that can make decisions are collectives (by definitional fiat), only collectives (but not combinations nor coalitions) can have duties.

Interestingly, Collins rejects a different line of argument against group duties: that combinations or coalitions are unable to perform the required action. As we have seen above, Collins maintains that non-agent groups can have abilities. She therefore rejects the 'Ability Argument' against group duties for combinations and coalitions and does *not* think that ought-implies-can (or, more precisely, duty-implies-ability) rules out group duties for coalitions and combinations (see Wringe 2019, with further references).

This is an interesting combination of claims. It can be summarized like this:

		Irreducible Abilities for combinations/coalitions?	
		no	yes
Irreducible Group Duties for combinations/coalitions?	no	Reductivists	Ability holists (Collins)
	yes	Duty holists	Ability and Duty holists

A reductivist with regard to combinations and coalitions denies that non-agent groups have irreducible abilities or irreducible duties. Instead, all such apparent group properties can be re-described in terms of individual abilities and duties. The reductionist position appeals with its ontological parsimony but struggles to explain intuitions that some non-agent groups appear to have abilities and perhaps duties.

An abilities and duties holist maintains that coalitions and combinations have irreducible abilities and duties. This thoroughly holist view can easily explain the aforementioned intuitions but is challenged to explain how groups without agency can have duties of their own.

A duty holist thinks that groups can have irreducible duties but no abilities. The duty holist needs to explain how non-agent groups can have duties. In addition, there is an ought-implies-can related challenge: surely a duty must require the ability to perform the action demanded by the duty – but the duty holist believes that groups do not have group-level abilities. However, the duty holist might retort that groups have abilities in a reductive sense: they are simply nothing but the aggregation of individual abilities. And since the aggregate of individual abilities may suffice to perform the required action, this concern can be defused.

Finally, there is ability holism, the position Collins argues for. She defends the view that combinations and collectives have irreducible abilities but she also thinks that they do not have any group duties. What are we to make of her view?

One interesting aspect of Collins's view is that she does not appeal to ought-implies-can-style arguments against group duties. However, the reason for her rejection of these sorts of arguments is unusual: rather than insisting that the aggregate of individual duties suffices to establish that groups 'can' perform the required actions (the standard way to defuse ought-implies can concerns about group duties), Collins takes a holistic position with regard to abilities, as we have seen above. The problem with combinations and collectives is not that they are unable to perform, Collins insists, the problem is that they are not able to make decisions and therefore lack agency. And since, according to the Agency Principle, only agents can bear duties, combinations and collectives do not.

Being an ability holist and a duty-reductionist about non-agent groups, Collins has to deal with the asymmetry between abilities and duties, an asymmetry that comes with a cost. In particular, it is a challenge for ability holists like Collins to explain how duties can reduce to individuals while the corresponding abilities reside with groups. Presumably, Collins assumes that the duties of individuals correspond with individual abilities, while combinations and collectives may have additional group abilities that are never 'called upon' by a corresponding group duty. An example might help to see the challenge. Suppose a group of strangers in a train carriage can overcome a knife-wielding attacker if and only if they act together (Held 1970). This combination therefore has the ability to overcome the attacker, and, following Collins's ability-holist view, they have this ability as a group. But the combination does not have a duty to overcome the attacker, Collins insists, as the combination lacks agency. Instead, the individuals have duties to coordinate, which involves being positively responsive to the efforts of others, or, if necessary, take appropriate actions to make the formation of a collective more likely (see chapter 4).

An upshot of this view is that each individual train traveller does not have a duty to overcome the knifeman (see p. 106) because each individual – on their own – is not able to do so. And, as we have seen, neither does the combination have a duty to overpower the knifeman (even though they are able to). This has the noteworthy counter-intuitive implication that *no one* has a duty to overpower the knifeman. This may strike some as implausible. However, Collins would insist that that is the right result: the individuals lack ability, the combination lacks agency, so both candidate duty bearers are ruled out. In Collins's analysis, when we talk about the duties of combinations and coalitions, we are – strictly speaking – making a mistake. We should instead avoid such statements and talk about individual coordination duties.

### **3. When and Why Do Individuals Have Coordination Duties?**

Collins's proposal takes us some way towards explaining why talk about duties borne by combinations and collectives is confused. But I am concerned that Collins's view does not offer a plausible account of when and why individuals have coordination duties. To see this, consider the natural response to the question 'Why do I have a duty to coordinate?' It would be: 'Because the combination/coalition you are part of has a duty to  $\phi$ '. Such a group duty would explain why there is a duty on individuals to do what needs to be done to make the group meet its obligations. But for Collins, this line of response is not available, as she rejects the idea that combinations or collections have duties. Collins has to say, instead, that



individuals have coordination duties whenever coordinated action by a combination or coalition can prevent moral disvalue. Collins often refers to this as 'remedying the situation' (p. 96) and sometimes she calls the state to be obtained 'morally important' (p. 116). The problem with these formulations is that they under-determine the situation in which coordination duties arise and the reasons why they arise, as it remains rather vague what is or is not morally important.

Consider two examples. It is arguably morally important for me to keep my promise to lunch with my colleague Hwa (because promise-keeping is important), but I doubt that that triggers coordination duties on my office neighbours to work together and ensure I show up at that lunch appointment. Similarly, it might be morally important for all children on a beach to have sunscreen lotion applied (because preventing harm is important), but that does not mean that this provides reasons for beach goers to coordinate and create a sunscreen vigilante group to achieve that goal. The point is: there are many morally important states that do not trigger any duty on individuals to coordinate. Nor do they give reasons to individuals in favour of coordination. The question 'Why do I have a duty to coordinate?' cannot be answered by merely pointing to 'morally important' states.

Collins does not appear to address this problem, but there may be a simple fix that is compatible with Collins's framework. Rather than appealing to moral importance or net moral value, one could say the following. A coordination duty arises if and only if: were the coalition or combination to be turned into a collective (with agency), that collective would have a moral duty to bring about the morally valuable state. The answer to the question 'Why do I have a duty to coordinate?' is then 'because if the combination/coalition had agency, that collective would have a duty to act'. This way of triggering coordination duties also avoids the problems with the examples just stated: even a collective of my colleagues would not have a duty to ensure I keep my lunch promises, regardless of whether keeping such promises is morally important. Neither does a collective of beach goers have a duty to form a sunscreen vigilante group, despite the increased moral value it would produce.

The solution proposed here has an interesting modal element: individual coordination duties arise in the actual world  $a$  if (i) the individuals are in a combination or coalition in world  $a$  and (ii) there exists a nearby possible world  $w$  in which the combination from world  $a$  is a group agent and that agent has a duty. This approach seems superficially similar to a duty to incorporate (Collins discusses this duty in relation to Held 1970 and Erskine 2014). However, I agree with Collins that incorporation is not always necessary to meet duties. As we have seen

above, it often suffices for members of combinations or coalitions to become responsive to one another, without creating decision-making process within the group. The upshot is that not all moral challenges require a group agent; sometimes mutually responsive individual behaviour will do. For instance, to save our person trapped under the car the bystanders do not need to form a collective – what is needed is everyone trying to lift the car until enough people try and make it happen.

The point of the solution proposed here is to identify the situations in which coordination duties arise by considering nearby possible worlds in which a collective agent exists. If that hypothetical collective agent has duties, then, I suggest, coordination duties arise to perform actions akin to the actions the hypothetical collective agent would need to perform to discharge the duties.

#### **4. Conclusion**

I have briefly looked at Collins's definition of group abilities, whether group abilities or group duties are reducible, and at the conditions under which individuals have coordination duties. Due to space constraints, I had to put aside other fascinating topics, especially Collins's insightful discussion of the differences between combinations and coalitions and her analysis of the duties of collectives and the membership duties that arise when group agents have duties. Stephanie Collins's deceptively slim book repays careful reading and shows us something important: that the existence of group duties hinges not on whether groups are able to perform required actions – it hinges on whether a group is an agent able to make its own decisions. Collins advocates for an important shift in the debate and deserves a wide readership in the areas of social ontology, ethics, political philosophy, and beyond.\*

Kai Spiekermann

London School of Economics

[k.spiekermann@lse.ac.uk](mailto:k.spiekermann@lse.ac.uk)

---

\* I am grateful to Steph Collins for comments and earlier discussions during a manuscript workshop at the Central European University in Budapest.

- Arthur, W. Brian 1994, 'Inductive Reasoning and Bounded Rationality', in *The American Economic Review*, 84(2), 406–411
- Erskine, Toni 2014, 'Coalitions of the Willing and Responsibilities to Protect: Informal Associations, Enhanced Capacities, and Shared Moral Burdens', in *Ethics & International Affairs*, 28(1), 115–145
- Gintis, Herbert 2000, *Game Theory Evolving* (Princeton: Princeton University Press)
- Held, Virginia 1970, 'Can a Random Collection of Individuals be Morally Responsible?', in *The Journal of Philosophy*, 67(14), 471–481
- List, Christian, & Pettit, Philip 2011, *Group agency: The possibility, design, and status of corporate agents*. (Oxford: Oxford University Press)
- Osborne, Martin J. 2004, *An Introduction to Game Theory* (Oxford: Oxford University Press)
- Pinkert, Felix 2014, 'What We Together Can (Be Required to) Do', in *Midwest Studies in Philosophy*, 38(1), 187–202
- Wringe, Bill 2019, 'Global obligations, collective capacities, and "ought implies can"', in *Philosophical Studies*. Online First