



Book Symposium

Olle Blomberg*

What We Ought to Do: The Decisions and Duties of Non-agential Groups

<https://doi.org/10.1515/jso-2020-0025>

published online July 14, 2020

Abstract: In ordinary discourse, a single duty is often attributed to a plurality of agents. In *Group Duties: Their Existence and Their Implications for Individuals*, Stephanie Collins claims that such attributions involve a “category error”. I critically discuss Collins’ argument for this claim and argue that there is a substantive sense in which non-agential groups can have moral duties. A plurality of agents can have a single duty to bring about an outcome by virtue of a capacity of each to practically reason about what they ought to do together. I also argue that Collins’ attempt to give a reductive account of such “we-reasoning” fails.

Keywords: collective moral obligation; group identification; non-agential groups; Stephanie Collins; team reasoning

1 Introduction

In *Group Duties: Their Existence and Their Implications for Individuals*, Stephanie Collins (2019) argues that only a moral agent can have a moral duty. A moral agent need not be an ordinary individual moral animal. It can be also be group agent. According to Collins, a group agent is a group whose members are united under a group-level decision-making procedure. Examples of such agents—“collectives”, as Collins calls them—include everything from a nation state or a corporation to, say, a group of snowboarding friends who together decide (by reaching a “conversation-based consensus”) which route they should take down the mountain to base camp.

A group of moral agents that is not itself an intentional agent cannot have a moral duty on Collins’ view. (Collins distinguishes between two types of non-agential groups: combinations and coalitions. A *coalition* is a group whose

*Corresponding author: **Olle Blomberg**, Department of Philosophy, Linguistics and Theory of Science, University of Gothenburg, Gothenburg, Sweden, E-mail: olle.blomberg@gmail.com.
<https://orcid.org/0000-0001-7696-7996>

members share a goal and have common knowledge that the goal is shared. A *combination* is simply any group of agents that is neither a coalition nor a collective.) Hence, according to Collins, a group of sunbathing strangers on a beach will normally not constitute a group agent and cannot therefore have a moral duty to save the large and unwieldy man that is drowning nearby. Neither can any or each of them have a moral duty to do so, since none of them is able to save him on their own. If they created a group agent, then that group agent could have a moral duty to save the drowning man. But the bearer of the duty would then not be *them*, but *it*—the group agent itself rather than the coalition of sunbathing strangers who created and sustains it.

Group Duties is a rich, comprehensive and lively treatment of the kind of moral duties we have as members of different kinds of groups, as well as the kind of moral duties that group agents that we partly constitute or interact with have. I have learnt a lot from reading it, and there is much that I agree with. However, my commentary will be a critical engagement which will solely focus on a point of disagreement, namely whether non-agential groups can have moral duties. I will argue, contra Collins, that there is a substantive sense in which non-agential groups can have moral duties. (I will have nothing to say about the moral duties of group agents and of their members.)

Collins is not denying that the sunbathing beachgoers will normally have moral duties related to the fact that the swimmer is in the process of drowning. Nor is she denying that most of us have moral duties related to, say, alleviating the harms of air pollution or climate change, or to stopping the populist extreme right-wing party from winning the election. But, she argues, these are ordinary individual moral duties (assuming, that is, the absence of a group agent). Each individual beachgoer has a duty, roughly, to perform actions that are responsive to what the others do, with an aim of bringing about the morally important outcome that the swimmer doesn't drown. According to Collins, these so-called "responsiveness duties" of the individual group members are sufficient for capturing what is morally required in this kind of situation: "We need not posit a group-level duty to explain or specify [the responsiveness duties]." (2019, p. 96).

I suspect that most philosophers who work on the moral duties of groups would take this to be a revisionary view according to which things aren't quite as they appear. Wringer (2014) would at least take Collins' view to be revisionary in this sense. He argues that it would normally be part of the first-person moral experience of agents such as the beachgoers that they had a group-level moral duty that explained their individual moral duties. Collins argues though, that her view saves the appearances just as much as a non-distributive view such as Wringer's (Collins 2019, p. 36–38). To do this, Collins appeals to various considerations that are supposed to explain why non-distributivists have mistakenly taken the appearances to have a collectivist character (Collins 2019, p. 63–64, 152). When it comes to the character of our first-person

moral experience as members of non-agential groups, I think Collins is right that it is consistent with both individualism and collectivism. However, it is important to note that Collins does not deny that we often talk as if we were attributing duties to groups in a non-distributive fashion. About the beachgoers, for example, we might say: “They ought to rescue the swimmer!” But Collins (2019) argues that we must “charitably reinterpret” (p. 181) such talk as a “mere shorthand for a conjunction of members’ duties” (p. 4). To actually attribute a single duty to a plurality of agents would be a “category error” (Collins 2019, p. 25, 59). The fact that Collins’ account requires this reinterpretation of all such attributions is arguably a theoretical cost that counts somewhat against it. Other things being equal, an account that didn’t systematically require reinterpretation of ordinary discourse would be preferable.

Note that the collective reading of, for example, “they ought to rescue the swimmer” does not entail that each group member has a duty to rescue the swimmer. If Ben and Laura together have an ability and a duty to rescue the large and unwieldy man that is drowning nearby, but Ben knows that Laura is unwilling to do her part, then it would be false that Ben has an ability and a duty to rescue him (depending on the circumstances, Ben could perhaps have a duty to *try* to rescue him though).¹ If there were such an entailment, then the linguistic appearances could be saved so to speak, and the problem of non-agential groups having duties could be partly dissolved. After all, the principle that only an agent can have moral duties would be respected. In a collective rescue case such as that involving the beachgoers, there would be a duty that each of the agents had precisely because they had it together. Again, there is no such entailment though.²

1 Assuming, at least, that duty implies ability. Both Collins and I assume that this is so (Collins 2019, p. 61).

2 Collins follows Kirk Ludwig in thinking that the semantics of singular action sentences in English are such that they include a *sole agency requirement*. As Ludwig puts it: “In English, ‘it was done by x’ implies that x was the sole agent of it” (2016, p. 25; quoted approvingly in Collins 2019, p. 105). As Collins notes in a footnote (p. 105, fn. 6), I have suggested that such a sole agency assumption is just a typical pragmatic implicature, not something that is semantically encoded (see also Blomberg 2019, sect. 4). On this view, it would be true of each beachgoer that he or she saved the drowning man if they did it together, simply because each of them was an agent, even if not the sole agent, of the rescue. Now, if it is true that “Ben rescued the swimmer” when he did it together with Laura, then one might think that it would also be true that he would have both an ability and a duty to rescue the swimmer if he and Laura had a joint ability and a joint duty to rescue the swimmer. Let me note here that I now think that, *if* the sole agency requirement must either invariably be part of the truth conditions of singular action sentences or else invariably not be part of them, then Collins and Ludwig are probably right that the sole agency requirement is semantically encoded (see especially Ludwig 2019, sect. 3). More generally, they are probably right insofar as sentences do have a context-independent logical form that is relatively insulated from pragmatic meaning (that this is so is certainly not a foregone conclusion though, see e. g., Wilson & Sperber 2012).

In the next section, I critically discuss Collins' argument against non-agential groups having moral duties. According to Collins' so-called "Decision Argument", only an entity that can make decisions can have moral duties. I emphasise that this argument, as Collins acknowledges, is inconclusive. Because of this, I will offer a line of reasoning that I suggest Collins could make use of to defend the idea that non-agential groups cannot have moral duties. I consider why Collins does not make the argument that I present. Arguably, she is blocked from using it due to her incomplete use of an analogy between, on the one hand, different temporal stages of an individual agent and, on the other hand, the members of a non-agential group.

Despite the discussion in section 2, I go on in section 3 to sketch an account according to which there is a substantive sense in which non-agential groups can, after all, have moral duties. This sketch appeals to the possibility that a non-agential group's members each have the capacity to engage in so-called "we-reasoning". In her book, Collins tries to do away with we-reasoning by giving a reductive account of what she calls "coalition-reasoning". In section 4, I argue that her account of coalition-reasoning arguably presupposes we-reasoning.

2 Two Decision Arguments

Collins argues that a non-agential group cannot be a duty-bearer because it cannot deliberate and make decisions. A duty, she writes, "functions as a fitting input into an entity's first-person future-directed decision-making about what to do" (Collins 2019, p. 8). I accept this characterisation of the function of duties. Now, since a non-agential group by definition lacks the ability to make decisions, one might think that this is enough to make it a category error to attribute a moral duty to such a group. However, this inference would be too quick.

This is because a moral agent can have a duty to do something that requires an automatic response that isn't—and perhaps cannot be—the direct result of a decision. To redeploy an example that Collins uses to illustrate that one can have duties and act in accordance with them without fulfilling them: Stephanie can arguably have a moral duty to save her friend's life in a situation where this requires her to quickly jerk out her hand to stop the friend from stepping out in front of an oncoming bus (Collins 2019, p. 90). The situation may be such that there is no time for making a decision, or perhaps the behavioural routine is so ingrained that it is simply triggered by the situation. The behaviour would then unfold beyond the direct control of Stephanie's first-person future-directed decision-making about what to do (p. 92). She could nevertheless *have* the duty, as well as *act in accordance with* it, even if she could not *fulfil* it—the latter would require that

she acted on a decision directly. At one point, Collins seems to suggest that Stephanie could have this kind of moral duty even if the requisite behaviour would have to be triggered in a way similar to how someone's patellar reflex is triggered by a sharp tap below the kneecap (*ibid.*). However, Collins would probably rather say—and I think she should rather say—that Stephanie can have a moral duty to behave in a way that requires her skills and habits to be spontaneously activated by the situation she is in. At one point, this behaviour would need to have been under her long-term control though (in a way that the patellar reflex isn't). That is, under some general description, the type of behaviour would need to have been the object of her first-person future-directed decision-making about what to do.

Hence, Stephanie's 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.³ If it couldn't be traced back to such past opportunities, then she would arguably lack the kind of ability that is required for moral duty (see Collins 2019, p. 68–69). Hence, it seems that a moral duty need not be a fitting input into the duty-bearer's decision-making *at the time of action*. But a duty-bearer must, according to the argument I have presented, nevertheless be able to make decisions that are related in the right way to what moral duty demands. Hence, the conclusion that only agential groups can have moral duties would still be secure.

Were Collins to defend her idea that there is an intimate connection between decision-making and moral duty against putative counterexamples involving duties to do things that require automatic behaviour, she could give this argument. Furthermore, this argument would allow her to defend that there is an intimate connection between a fairly robust sense of decision-making (e. g., conscious personal-level decision-making) and moral duty, which could in turn support her thesis that non-agential groups cannot have duties.

Collins' conclusion that non-agential groups cannot have moral duties is instead based on an inference to the best explanation (Collins 2019, p. 93). She argues that the best explanation of the (alleged) fact that “babies, dogs, and tables” don't have moral duties is that they (each individually) cannot decide to do what duty requires based on the reasons that make it their duty. Without the ability to make such a decision with respect to at least one duty, an entity cannot deserve moral approval, commendation or praise. According to Collins, the concept of moral duty is intimately connected to the possibility of such moral desert. Now,

³ The idea that duty might be traced back to earlier opportunities for making a decision is compelling, but not without problems. See Vargas (2005) and Shabo (2015) for problems with the use of the same kind of tracing procedure applied to an individual's backward-looking moral responsibility.

since a non-agential group cannot make decisions, it cannot deserve moral praise or positive reactive attitudes from others, nor can it have duties, according to Collins. As Collins admits, the inference to the best explanation—her “Decision Argument”—does not conclusively show that it is the presence/absence of decision-making ability that best explains why normal adult human beings have moral duties while babies and dogs (let us set the tables aside) do not. It is not conclusive because of the fact that the ability to fulfil a moral duty is a complex ability that not only requires the ability to make decisions but the ability to make decisions that are responsive to moral reasons. 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? Such responsiveness to moral reasons may perhaps be exhibited by a non-agential group. This possibility at least isn’t ruled out by definition.

Now, Collins’ notion of ‘decision’ is very minimal. According to Collins, “any intentional performance that is explained by the beliefs and desires of the performer is a performance for which the performer has made a decision.” (2019, p. 91) And: “Decisions—like the desires and beliefs on which they are based—are very often sub-personal, offline, tacit, or unconscious.” (Collins 2019, p. 90)⁴ In light of this, perhaps Collins could say that an agent who is responsive to reasons in the right way *must* be making and acting on decisions. If the notion of decision is thin enough, then she arguably could, and this would strengthen the case for her proposed inference to the best explanation. However, with such a thin and fuzzy notion of decision, I take it that it is no longer clear that babies, dogs and even a group such as that of the sunbathing strangers cannot make decisions.

So, it is unclear whether Collins’ Decision Argument really delivers the conclusion that she needs. In light of this, why doesn’t Collins rely on a more robust notion of decision-making to then make the different “Decision Argument” that I sketched? According to this argument, a duty-bearer must be able to make decisions that are related in the right way to any moral duty that she bears. This is because the behaviour that is in accordance with the moral duty must at least be under the agent’s long-range control in order for her to be praiseworthy for it. The reason why Collins does not rely on this kind of argument, I suspect, is that she thinks that endorsing such an argument would put pressure on her to accept that some non-agential groups might, after all, be able to make decisions and have

⁴ Given this, it is actually not obvious why we should not interpret Stephanie’s jerking out her hand to stop her friend from being run over by the bus as the immediate result of her beliefs, desires and decisions. At any rate, Collins uses the example to illustrate that one can have duties and act in accordance with them without fulfilling them. In the example, this is supposed to be because the performance isn’t the outcome of a decision made by Stephanie.

duties. In a discussion about how an agent's abilities to do things are tied to the possibility of the agent's trying to do them, Collins makes argumentative use of an analogy between the temporal parts of an individual agent and the members of a group (Collins 2019, p. 70). According to this analogy, the relations between temporal parts of an individual agent is relevantly similar to the relations between the members of a group. Based on this, we might reason as follows: Suppose that Stephanie's highly automatic behaviour of jerking out her hand to stop her friend from being run over by a bus occurred at time t_2 . Suppose furthermore that this behaviour was in accordance with a moral duty she had at that time. Who was the bearer of that moral duty? In light of the analogy, one might think the duty-bearer would have been Stephanie-at- t_2 . The relevant earlier decision that instilled the behaviour in Stephanie-at- t_2 's suite of automatic routines was made by a different temporal part though, namely Stephanie-at- t_1 (I am simplifying here: there would have been many decisions made by different earlier temporal parts). If we identify the decision-maker and the duty-bearer in this way, then the Decision Argument that I have proposed is undermined. After all, it looks like the duty-bearer, Stephanie-at- t_2 , need not be able to make decisions that are related in the right way to what moral duty demands. In principle, this could make room for a group that is a duty-bearer but not a decision-maker.

However, were one to reason in this way, one would be making a mistake. The duty-bearer is not Stephanie-at- t_2 , but Stephanie, the whole agent. Stephanie-at- t_1 and Stephanie-at- t_2 are one and the same agent. Hence, the decision-maker and the duty-bearer are one and the same. Stephanie's life has different temporal stages of course, and these are connected causally as well as psychologically through memories, expectations and intentions, as well as by capacities at each stage for making decisions and evaluating outcomes from a temporally extended perspective. Arguably, it is partly by virtue of these connections and capacities that she, Stephanie, bears the moral duty at t_2 with which her behaviour accords. In contrast, something analogous to this is generally not true with respect to the members of a group and their moral duties. Suppose that Stephanie's automatic jerking out of her hand is completely insulated from her decision-making ability, but I, another agent, can turn it on and adjust it by remote control, so that it is triggered in certain situations. Here, Stephanie, who is a member of the combination consisting of herself and I, would not bear a moral duty at t_2 to jerk out her hand to stop her friend from being run over, despite the fact that I (let us suppose) could decide whether or not her automatic routine would be activated.⁵ Nor would

⁵ This does not mean that I cannot ever in some sense be responsible for another agent's action. When the other agent is acting on my orders, for example, or when we are engaged in joint intentional action, socially extended moral responsibility is arguably possible (Shoemaker 2012).

the combination of Stephanie and I bear such a moral duty at t_2 . Hence, I do not think Collins has any reason not to endorse the argument I have suggested. Relying on it would not push her into having to accept that non-agential groups can have moral duties, and it would allow her to rely on a more robust notion of what a decision is.

Now, in the next section, I will argue that, despite the Decision Argument that I have sketched in this section, non-agential groups can have moral duties after all.

3 We-reasoning and Group Duties

The unorganised beachgoers do not together constitute a group agent with its own decision-making capacities. However, it is plausible that if they were all co-present on the beach and if what was happening with the drowning man was common knowledge between them, then they would all think of themselves as members of that group of co-present beachgoers who were witnessing someone about to drown. They would realise that they were all in the same “moral boat” so to speak, and that it was up to *them* to urgently do something about the situation.

In this situation, one might think that the beachgoers would engage in what Collins and others have called “we-reasoning” or “team reasoning” about what they should do together (see also Hakli, Miller and Tuomela 2010). According to theories of we-reasoning, an individual who faces a social decision situation may ask herself “What should we do?” as well as the more familiar “What should I do?”.⁶ Such we-reasoning results from a perspective shift that occurs when the individual identifies as a member of the group of agents that are present in the situation. Group identification can be prompted by various cues in the environment. According to Bacharach’s (2006) “interdependence hypothesis”, among these cues are the facts that the agents appear to have common interests and a scope for common gain (for discussion, see Hindriks 2012). The idea is then that an individual in a social decision situation who has group-identified considers the options that are available to the group (that is, which outcomes could be brought about by combinations of actions by group members) rather than the options that are available to herself as an individual agent. She then 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.

Whether there is this kind of we-reasoning is controversial, but I find it to be an attractive and plausible idea. The mechanism of group identification has been

⁶ See Gold and Sugden (2007) for a philosophically oriented overview of the team reasoning literature.

proposed by so-called Social Identity Theory in social psychology and has helped explain a host of psychological intra- and intergroup effects (Hogg et al. 2017; Tajfel 1979). In addition, we-reasoning theory offers elegant solutions to some theoretical and empirical problems facing orthodox game theory (Bacharach 2006; Gold and Sugden 2007).

One theoretical problem facing orthodox game theory is to explain how the theory can endorse what intuitively is the obviously rational choice for each player in a game called Hi-Lo. Collins constructs a useful collective rescue case called *Hi-Lo Rescue* that can be used illustrate the problem (2019, p. 127–128). In *Hi-Lo Rescue*, the optimal pattern for the two beachgoers Ben and Laura to stop the swimmer from drowning is that they each perform rescue actions. A good but suboptimal pattern is that they each instead do their part in making a call to the emergency services (one of them has the phone and the other knows the phone number to the emergency services). The worst outcome would result if one of them performs rescue actions and the other tries to do his or her part of phoning the emergency services: the swimmer will then definitely drown. Suppose that, for whatever reason, it is not possible for Ben and Laura to communicate in this situation. All this is common knowledge between them. Intuitively, it seems obvious that each should perform rescue actions, as they then have the best chance of saving the swimmer.

If each of Ben and Laura are only allowed to ask themselves “What should I do?” and to do what is best in light of what they expect that the other will do, then they will not be able to rationally arrive at a determine answer the question. All orthodox game theory instructs each of them to do is to perform rescue actions if the other performs rescue actions, and to do their part of phoning the emergency services if the other does their part of phoning the emergency services. By the standards of orthodox game theory, there is no determinate rational solution, despite it being intuitively obvious what each ought to do. However, if the *Hi-Lo Rescue* situation prompts each of them to identify as members of the group consisting of themselves and the other and each then proceeds to ask themselves “What should we do?” and then “What should I do as part of what we should do?”, then each could rationally choose to do their part of the optimal pattern, the pattern that will bring about the best outcome. Hence, the we-reasoning theory could vindicate and explain our intuition about what is the rational thing to do in *Hi-Lo Rescue* and other similar cases.

Given that this picture of we-reasoning is accurate, the statement “You [second-person plural] ought to rescue the swimmer”, when addressed to the beachgoers in a collective rescue case, would arguably be a fitting input into each beachgoer’s *first-person plural* future-directed decision-making about what they should do, as well as into their first-person singular decision-making about how

they should do their own part of what they should do together. Given that all of the beachgoers have the capacity to identify as group members—as included in one and the same “we”—and given that they are therefore each poised for we-reasoning in the situation at hand, *they* can arguably have a moral duty together to rescue the swimmer. That is, there is a substantive sense in which the duty can figure in 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. There is no need to reinterpret the statement that they ought to rescue the swimmer as a mere shorthand for a conjunction of ordinary individual duties. Furthermore, in *Hi-Lo Rescue*, we need the non-distributive moral duty of which Ben and Laura collectively are the bearer in order to explain how it can be that each of Ben and Laura has an individual moral duty to perform rescue actions.⁷

Note that, on the account I am proposing, group members need not actually group identify and we-reason in order to *have* a moral duty together; it is sufficient that they have the required capacities to do so. If each member does group identify and we-reason, however, and does their parts of pattern of actions that is most likely to bring about the morally best outcome, then it seems that the group has not only acted in accordance with its moral duty, but may also have fulfilled its moral duty. Suppose that each beachgoer desires that the drowning man be saved, each believes that the combination of their rescue actions will result in this desired outcome being brought about, and as a result of group-identification and we-reasoning, each follows through on a decision to do their part of that combination. As a result, the drowning man is saved. Here, the beachgoers’ joint performance has a psychological explanation that accords with its moral justification. The joint performance is required precisely because they have a joint ability to bring about the morally important outcome while no individual have an ability to bring it about on their own. That this joint performance is required in turn explains why each beachgoer is morally required to do their part, on the presupposition that the others do their parts. Furthermore, the primary target for our moral approval, commendation and praise with respect to the rescue of the drowning man would arguably be the group of beachgoers (the beachgoers considered plurally), not each beachgoer considered individually. Individual group members would be praiseworthy for doing their parts in a way that were responsive to the actions of the others, but only the group would deserve praise for rescuing the swimmer. However, the group is not a group agent in Collins’ sense. There is no group-level

⁷ This proposal is based on joint work-in-progress with Björn Petersson. The account we are developing is in many ways similar to Schwenkenbecher’s (2019) team reasoning-based account of collective moral obligations.

decision-making procedure that the group members are united under. Each group member is doing their own we-reasoning.

In response to this suggestion, Collins would probably say that I, and perhaps the beachgoers themselves, have been misled by the experiential character of the emergency situation on the beach. Regarding a type of we-reasoning, she writes that it gives rise to a “feeling of closeness” which might “induce the thought that the group has duties” (Collins 2019, p. 152). According to Collins, this induced thought is mistaken because it involves the thought that the individual who group-identifies conceives of herself as a part of a larger group-level agent. In other words, the “we” that an individual includes herself in is conceived as an additional agent, over and above herself and the other group members. Furthermore, perhaps Collins would add that without such a group-level agent, the group is not really deliberating and making decisions, which is what is needed for the group to have a duty according to her Decision Argument. I agree with Collins that if group identification is conceived in this way in the context of a non-agential group, then it will involve a mistake, error or perhaps a “fiction” (Collins 2019, p. 133–134, 137, 139). She is right, furthermore, that we-reasoning theorists have often conceived of group identification in precisely this way, even when the groups in question are not group agents. However, we can arguably conceive of group-identification differently, so that the group does not have to constitute an additional intentional agent in order to deliberate and make decisions.

To explain what I have in mind, compare the case of the coalition of beachgoers rescuing a swimmer with a case where there is only a single beachgoer. Suppose I live by the beach on a deserted island. The only other living creature nearby is a man bobbing in the water. For reasons that are too complex to go into, the man must stay in the water to survive, but since he gets tired staying afloat, he is at risk of drowning unless I occasionally help him in various ways (for example, by holding him and lifting him above water so he can rest, by giving him something to eat, etc.). In this situation, I arguably have a moral duty to help him stay afloat. This duty to ensure that the bobbing man stays afloat is one that concerns not just myself here and now as I deliberate about what to do in the next moment (cf. Collins 2019, p. 51, 200). Rather, it concerns my present deliberations about what I should do throughout next week, say. At any one time, the duty is a fitting input to my first-person future-directed decision-making about what to do, even though the duty concerns something that I cannot directly do or control at that present time. This is possible even if me-now and me-in-the-future don’t together constitute a larger *additional* agent that itself can deliberate and perform a large action consisting of the part that me-now performs and the part that me-in-the-future performs. Similarly, I want to suggest that the fact that the beachgoers identify as group members and reason about what they should do together need not involve

the mistaken thought that they constitute an additional group-level agent. (Exactly how group identification should best be conceptualised is an interesting and difficult question. I am not able to answer it here unfortunately.⁸) I am thus extending the analogy between temporal parts of an individual agent and the members of a group which Collins draws on – it is not only that the relations between temporal parts are analogous to the relations between group members, the temporally extended whole is also analogous to the group as a whole.⁹ Neither the temporally extended whole nor the group should be conceived of as an additional agent over and above its stages or members.

4 We-reasoning and Coalition-reasoning

I suspect that Collins would resist this suggested account of the moral duties of non-agential groups in part because she thinks that there isn't really such a thing as a *sui generis* we-form of practical reasoning. Instead, Collins provides an account of a form of reasoning that she dubs "coalition-reasoning". Her account is reductive in the sense that coalition-reasoning is actually just a form of I-reasoning where each agent is assured about what the other agent(s) is (are) going to do. Each member simply reasons prompted by the question "What should I do?" The responsiveness duty of each of the two beachgoers in *Hi-Lo Rescue* become something like "I ought to coalition-reason with the aim of the swimmer surviving".¹⁰ The primary input to each beachgoer's practical I-reasoning in this case would thus not be "we ought to rescue the swimmer". No such group duty is needed in order to explain what each individual has a moral duty to do.

Coalition-reasoning is similar to we-reasoning insofar as it is supposed to enable an individual agent to choose the optimal pattern of actions for a group consisting of himself or herself and others, and then conclude that he or she should do his or her own part of that optimal group pattern. According to Collins, coalition-reasoning will enable each of Ben and Laura to settle on the optimal

8 One possibility is that a group-identifying individual thinks of herself and the others as parts of a single "group causal agent" rather than an intentional agent (Pettersson 2007). In a recent paper, I argue that in order to use group identification and we-reasoning to give an account of joint intentional cooperative action, the group identification ought to involve the thought that oneself and the other participants are part of a properly functioning single body in action (Blomberg 2020). This is another possible conception.

9 Natalie Gold (2018) develops this analogy in an account of "intrapersonal team reasoning".

10 In this respect, Collins proposal is somewhat similar to Anne Schwenkenbecher's (2019) account of collective moral obligations. Schwenkenbecher does not think that we-reasoning can be reductively understood as a kind of coalition-reasoning though.

pattern of rescuing the swimmer and do their part of it in *Hi-Lo Rescue*. However, I am sceptical that Collins' account of coalition-reasoning works without a tacit appeal to something like a sui generis form of we-reasoning. She gives the following schema for coalition reasoning:

Coalition-reasoning:

1. It is a reciprocal belief between you and me that each of us has objective O [e. g., that the swimmer is no longer drowning].
2. 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 O.
3. You and I reciprocally believe that each other will do their part in the best way of patterning our respective actions so as to realize O. (From 1, 2).
4. Given that I believe you will do Y, I know what will best realize my objective.
5. Therefore: I will do what will best realize my objective.

(Collins, 2019, p. 139–140)

The crucial transition here is from 1 and 2 to 3. Collins thinks that the mere fact that we have a commonly known shared goal gives each of us evidence that the other will do their part of the optimal pattern rather than their part of a merely good but suboptimal pattern (performing their part of the rescue rather than doing their part of calling the emergency services). Now, once there is such evidence which can ground a justified belief that the other will do their part through “a ‘ratcheting up’ effect” (Collins 2019, p. 144–145), each can just employ normal I-reasoning to maximise expected (moral) utility (in steps 4–5). The initially small evidence provided by the commonly known shared goal is supposed to be ratcheted up because, in a state of common knowledge, each of us will know that the other knows that we have some evidence that they will do their part of the optimal pattern, and this gives them further reason to choose to do their part of the optimal pattern, which gives each of us further evidence that they will do their part of the optimal pattern, and so on and so forth.

However, if each member is limited to employing ordinary I-reasoning—so-called “best reply” reasoning—where practical reasoning is giving an answer to the question of what “I” should do given what I expect others to do, then 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 (the Hi equilibrium) for realizing the goal, as opposed to doing their part of the good but suboptimal pattern of actions (the Lo equilibrium) for realizing that goal. This is the Hi-Lo problem all over again. The commonly known shared goal would only give a member reason to choose Hi if they expect that the others will choose Hi. But the only reason for expecting this would be if the others expected that they themselves would choose Hi. If the others expected that they themselves

would choose Lo, then they should expect that the others would choose Lo as well, and this would give them reason to actually choose Lo. Given the standards of orthodox game theory and ordinary I-reasoning, there is no determinate rational solution to the problem.¹¹

So, as far as I can see, Collins has not shown that the transition from 1 and 2 to 3 in her coalition-reasoning schema does not actually rely on there being a distinct mode of practical reasoning which we can call we-reasoning. Furthermore, insofar as group-identification and we-reasoning provides the best explanation of how people solve Hi-Lo problems in practice, there is good reason to think that each coalition-reasoner must assume that the other group-identifies and engages in such we-reasoning in order to transition rationally from 1 and 2 to 3.

Despite Collins' attempt to do away with (a non-reducible form of) we-reasoning, I think that there is such a form of reasoning and that we often engage in it as a result of group identification. As I have briefly sketched, there is a way of using we-reasoning and group identification to make sense of the idea that the beachgoers, considered as a coalition, can have a moral duty to rescue the swimmer. Furthermore, such a group-level duty is needed to explain and specify the duties that each individual has to do their part of the optimal pattern for bringing about the morally important outcome.

5 Conclusion

It is initially appealing to think that only a moral agent can have a moral duty. However, I have argued that the possibility of group members practically thinking about what they should do together, rather than each just individually thinking about what they individually should do in light of what their expectations about what the others will do, shows that there is a substantive sense in which a group without a group-level decision-making procedure can have a duty. This is a duty that is needed in many cases to explain what the individual group members have moral duties to do. One of the central theses of *Group Duties* is thus false or is at least in need of substantial qualification. After all, there is a

¹¹ Collins (2019, p. 131) discusses Derek Parfit's (1988) idea that the members naturally each choose Hi because the best equilibrium (Hi) is salient. It is unclear whether or not she endorses this though. Anyway, it is not clear why the worst equilibrium (Lo) isn't equally salient. Furthermore, it seems unsatisfactory to appeal to a brute psychological feature such as salience to explain why members would choose Hi. What we want is an explanation that vindicates that this is the rational thing to do.

substantive sense in which a group that is not an intentional agent can have a moral duty.

Acknowledgments: I am greatly indebted to Björn Petersson for discussions about we-reasoning and group duties. I would also like to thank Dave Ward and Gunnar Björnsson for insightful and helpful questions and comments on a draft of this commentary. I am grateful for funding received from the Lund Gothenburg Responsibility Project (PI: Prof. Paul Russell), which is in turn funded by the Swedish Research Council (Vetenskapsrådet), as well as for funding from the Swedish Research Council project grant 421-2014-1025.

References

- Bacharach, M. 2006. *Beyond Individual Choice: Teams and Frames in Game Theory*. Princeton: Princeton University Press.
- Blomberg, O. 2019. "From Simple to Composite Agency: On Kirk Ludwig's from Individual to Plural Agency." *Journal of Social Ontology* 5 (1): 101–24.
- Blomberg, O. 2020. "Intentional Cooperation and Acting as Part of a Single Body." *Mind & Language*. <https://doi.org/10.1111/mila.12274>.
- Collins, S. 2019. *Group Duties: Their Existence and Their Implications for Individuals*. Oxford: Oxford University Press.
- Gold, N. 2018. "Putting Willpower into Decision Theory: The Person as a Team Over Time and Intrapersonal Team Reasoning." In *Self-Control, Decision Theory, and Rationality*, edited by J. L. Bermudez, 218–39. Cambridge: Cambridge University Press.
- Gold, N., and R. Sugden. 2007. "Theories of Team Agency." In *Rationality and Commitment*, edited by F. Peter, and H. B. Schmid, 280–312. Oxford: Oxford University Press.
- Hindriks, F. 2012. "Team Reasoning and Group Identification." *Rationality and Society* 24 (2): 198–220.
- Hogg, M. A., D. Abrams, and M. B. Brewer. 2017. "Social identity: The Role of Self in Group Processes and Intergroup Relations." In *Group Processes & Intergroup Relations* 20 (5): 570–81.
- Ludwig, K. 2016. *From Individual to Plural Agency. Collective Action*, Vol. 1. Oxford: Oxford University Press.
- Ludwig, K. 2019. "Individual and Collective Action: Reply to Blomberg." *Journal of Social Ontology* 5 (1): 125–46.
- Parfit, D. 1988. *What We Together Do*. Manuscript. Available on PhilArchive: <https://philarchive.org/archive/PARWWT-3>.
- Petersson, B. 2007. "Collectivity and Circularity." *Journal of Philosophy* 104 (3): 138–56.
- Schwenkenbecher, A. 2019. "Collective Moral Obligations: "We-Reasoning" and the Perspective of the Deliberating Agent." *The Monist* 102 (2): 151–71.
- Shabo, S. 2015. "More Trouble with Tracing." *Erkenntnis* 80 (5): 987–1011.

- Shoemaker, D. 2012. "Responsibility without Identity." *The Harvard Review of Philosophy* 18 (1): 109–32.
- Tajfel, H., and J. Turner. 1979. "An Integrative Theory of Intergroup Conflict." In *The Social Psychology of Intergroup Relations*, edited by S. Worchel, and W. G Austin, 7–24. Chicago: Nelson-Hall.
- Vargas, M. 2005. "The Trouble with Tracing." *Midwest Studies in Philosophy* 29 (1): 269–91.
- Wilson, D., and D. Sperber. 2012. *Meaning and Relevance*. Cambridge: Cambridge University Press.
- Wringe, B. 2014. "Collective Obligations: Their Existence, their Explanatory Power, and their Supervenience on the Obligations of Individuals." *European Journal of Philosophy* 24 (2): 472–97.