

This is a preprint of an article whose final and definitive form will be published in the Australasian Journal of Philosophy; the Australasian Journal of Philosophy is available online at: <http://www.tandf.co.uk/journals/>.

COLLECTIVES' DUTIES AND COLLECTIVISATION DUTIES

Stephanie Collins

Abstract

Plausibly, only moral agents can bear action-demanding duties. This places constraints on which groups can bear action-demanding duties: only groups with sufficient structure—call them ‘collectives’—have the necessary agency. Moreover, if duties imply ability then moral agents (of both the individual and collectives varieties) can bear duties only over actions they are able to perform. It is thus doubtful that individual agents can bear duties to perform actions that only a collective could perform. This appears to leave us at a loss when assigning duties in circumstances where only a collective could perform some morally desirable action and no collective exists. But, I argue, we are not at a loss. This article outlines a new way of assigning duties over collective acts when there is no collective. Specifically, we should assign *collectivisation duties* to individuals. These are individual duties to take steps towards forming a collective, which then incurs a duty over the action. I give criteria for when individuals have collectivisation duties and discuss the demands these duties place on their bearers.

Key Words

Group agency, collective responsibility, collectivisation.

1. A Problem for Group Duties

When morally pressing circumstances call for synchronised actions by several individuals, we often say the group has a duty: ‘the hikers have a duty to lift the fallen tree off the child’, ‘the pedestrians have a duty to stop the mugging’, ‘the beachgoers have a duty to save the

drowning person'. In each case, suppose the morally desirable outcome can be brought about only if the individuals work together.

How should we understand these locutions? Sometimes, the group as such will have a duty to perform some morally desirable action. This can be so only if the group is an agent—at least if we assume, as I will, that only agents can bear duties to perform actions. This assumption is based on the natural thought that only agents can act. So, if a group acts (or has a duty to act) then the group is an agent. Acting requires having something that plays the role of aims, on which the agent acts. Let us assume that agents form aims by processing reasons for them, that is, by implementing a decision-making procedure.¹ I take duties to be pro tanto, weighty, objective, moral reasons. Where duties demand that an action be done, they demand it be done by an agent that can process *moral* reasons when forming their aims. I will therefore assume that each action-demanding duty belongs to a moral agent: to an entity with a decision-making procedure that can process moral reasons to produce aims on which the entity acts.

We often have more than this in mind when we speak of agents: complex attitudes, phenomenology, conceptual understandings, and so on. Groups are incapable of such full-fledged human agency. But I argue (in §2) that they are capable of at least approximating the minimal conditions for agency, by having group-level decision-making procedures that can take in and process reasons, produce outputs that have the role of aims, and distribute aim-related instructions to members. (Others take similar views of group agency [French 1979, 1984; List and Pettit 2011; Pettit 2010, 2007; Pettit and Schweikard 2006; Wendt 2004: 299].) I label such groups 'collectives'. If a collective's decision-making procedure can process moral reasons, then that collective is a moral agent. When members follow the procedure's instructions and at least one of the collective's aims is thus achieved, the achievement of the aim is an action of the collective.² Collectives can have duties regarding their actions.

Yet suppose the group of hikers is not a collective and so cannot have a duty to perform the tree-lifting action. How, then, should we understand the ordinary locution? This is the problem with which I am primarily concerned. One possibility is that each hiker has a duty to perform the action: each has a duty to lift the tree. After all, each hiker is an agent.

¹ Agents might acquire aims in other ways. This article is concerned only with aims agents form by taking moral reasons into account, which presumably requires a reason-processing method.

² If one rejects collective intentions, the intentional component of the collective's action can be reduced to the members' several intentions to play their part in achieving the collective's aim. See similarly Tuomela [2006: 36-7]. Importantly, though, I will argue that the collective's aim and role distribution are produced by the collective itself.

But this violates a highly plausible variation on ‘ought implies can’: that action-demanding duties belong to only those agents that are able to perform the action the duty demands. It has been stipulated that one hiker, acting alone, is unable to lift the tree. Assuming ‘duty implies ability’, we cannot assign the duty over this action to any of the individual hikers.

We might instead assert that the group of hikers is a moral agent and can bear action-demanding duties. Yet this entails relaxing our concept of agency so much that it loses its distinctiveness. This also gets things backwards: moral agency should constrain duty bearing, not vice versa. Alternatively we could agree that the group lacks agency, and claim that non-agents can bear duties to things other than act, such as bring about or see to outcomes. We would then have to explain how these duties differ from duties to act, such that bearing them does not require agency. Given the tight intuitive connection between having agency and being eligible to bear duties, this looks unpromising. Or we might say that there are simply no action-demanding duties when an outcome requires individuals to work together and no collective is present. Yet denying there are duties related to the tree being lifted, the mugging being stopped, or the swimmer being saved is deeply counterintuitive. These outcomes present themselves as extremely important. It is imperative that there are action-demanding duties in the vicinity.

§2 solves the problem of what the ordinary locutions must mean when there is no collective. I argue that if the hikers are not a collective, then when we say ‘the hikers have a duty to lift the fallen tree off the child’ we must mean that each individual in the non-collective group has an individual duty to take steps to transform the group. This solution allows us to maintain that each action-demanding duty belongs to an agent; that agents cannot have duties they cannot fulfil; that each agent, and no non-agent, has a decision-making procedure; and that there are duties related to morally pressing group actions.³ Taking steps to transform the group requires the individuals to be what I call ‘responsive’ to the others. Sometimes, the conjunction of some individuals’ responsive actions is sufficient for the morally pressing circumstance to be addressed reliably. Then, each individual simply incurs a duty to perform responsive actions with a view to the circumstance’s being addressed. But often responsive actions alone are insufficient. In some such cases, each individual incurs a duty to perform responsive actions *with a view to there being a collective that can reliably address the circumstance*. These are individual duties to collectivise. The newly formed

³ If we reject these premises, other options may unfold. In §2.3, I consider the possibilities that non-agent groups have duties to do things other than act and that duty does not imply ability.

collective agent then incurs a duty to perform a circumstance-addressing action and each individual incurs a duty to perform the role the collective assigns him in that action.

This analysis matters. The specific acts or omissions for which we hold individuals retrospectively responsible, and that morality prospectively demands of them, will likely differ depending on whether their duties are (a) straightforward duties to perform a group action, (b) duties to respond to others with a view to addressing some circumstance, (c) duties to respond to others with a view to forming a collective that can address some circumstance, or (d) duties to play their role in a collective. This article is primarily concerned with (c), though there will be much to say about (a), (b), and (d). Ample attention has been paid in recent years to collective agents, collective action, and collective responsibility [Bratman 1992; Gilbert 2001, 2006; Kutz 2000; List and Pettit 2011; Mellema 2006; Pettit 2007, 2010; Pettit and Schweikard 2006; Sudgen 2000; Tuomela 2006]. Yet surprisingly few have taken up the issue of when and why individuals might severally bear duties to take steps towards *creating* a collective agent.⁴

2. Collectives' Duties

2.1 Responsiveness

Consider six strangers at the beach: one drowning and the others sunbathing. Suppose each sunbather infers that every sunbather desires the swimmer be rescued. The swimmer can be reached only with a motorboat. It will take two people to drag the boat to the water and hold it while a third starts it. The boat will take off straight away, so the fourth and fifth, who will pull the swimmer into the boat, must already be in the boat. Two are required to pull in the swimmer because none are strong enough to do so alone.

⁴ One exception is Held [1970: 480], who comments in passing that

. . . it may well follow that in some cases all the individuals in a random collection are responsible for not acting to transform the collection into an organized group, even though none of these individuals is responsible for not taking the action that ought to have been taken by an organized group in these circumstances.

However, she doesn't explain any further and the rest of her article defends ad-hoc groups as duty-bearers, which I will deny. Lawford-Smith [2012: §2] suggests that 'four strangers who happen to be passing in the stairwell . . . [where a piano has been dropped on a child] might be blameworthy if they do not attempt to *form* the kind of collective that could lift the piano', but worries that 'this is plausibly itself a collective action' and that 'non-existent agents cannot have obligations.' Later, I address these concerns.

Suppose one beachgoer, Ben, starts trying to drag the boat. Another, Laura, sees him and decides to help because she believes it will satisfy each of their desires that the boat move. Two others, Jon and Julie, run to the water, ready to jump in the boat once it is there. They do this because the others are already dragging the boat, otherwise they would drag it. Stan sees what the others are doing, remembers he has great boat-starting skills, and so runs to the water to be ready to start it. Jon and Julie each jump in the boat once it gets there because they can see Ben and Laura are now tired. Ben and Laura each hold one side of the boat because each believes this will help Stan start the boat. Stan starts it. When the boat arrives at the swimmer, Jon and Julie each grab one of the swimmer's arms and pull. The swimmer ends up in the boat.

No beachgoer could have rescued the swimmer alone. Moreover, the complementary nature of the individuals' actions—that they were different types of actions requiring synchronisation—meant they were unlikely to have been distributed successfully if each chose his or her action independently of the others. So each made inferences about the others' aims and used those inferences to help form his or her aims. The rescue was the result of the autonomously formed decisions of separate individuals, each of whom treated the others as features of the environment. If all but one individual's actions were fixed—say, by laws of nature—then we could perhaps say the remaining individual performed the rescue. But assuming the others' actions were not fixed, no individual performed the rescue. Each performed only his or her individual actions—though each did so with the hope and intention that the swimmer would stop drowning as a result.

The beachgoers were responsive. To put it precisely, an individual, A, is responsive to another, B, just in case A's inferences about the inputs or outputs of B's decision-making procedure become inputs to A's decision-making procedure, in virtue of A responding positively to B's perceived interests. That is, responsive individuals are coordinated. They have a 'shared end', in Bratman's [1992] terminology. In the limiting case, it is not commonly known that the end is shared and responsiveness is unilateral. Yet between minimal responsiveness and collective agency, there are many degrees of responsiveness. These include Bratman's [1992: 328] 'shared cooperative activity', in which 'each participating agent [knows] that the other is attempting to be similarly responsive'; Kutz's [2000] acting together through 'participatory intentions' to do one's part; and Tuomela's [2006: 38] joint intentional action, which 'amounts to the group members' jointly intending X

and jointly acting on the basis of their joint intention.’⁵ None of these entail—though none are incompatible with—a group-level decision-making procedure, which I will argue is necessary for collective agency.

There are two interpretations of the responsive beachgoers’ actions. The first holds that only agents can act and that each agent has a decision-making procedure. ‘Only agents can act’ implies not just ‘only an agent can do any component of any action’, but also ‘only an agent can do any action as a whole’. That is, if rescuing the swimmer was an action, then only an agent could have done it.⁶ The only decision-making procedures at the beach were those of each individual, so the group was not an agent. Thus, on this first view, the group did not perform the action of rescuing the swimmer. Neither did any individual perform that action. The rescue, then, was not an action. However, ‘where there is an action which some individuals perform, those individuals make up an agent’ is consistent with ‘where there is an outcome which results from some individuals’ actions, those individuals need not make up an agent’. On the first view, ‘the swimmer is no longer drowning’ was an *outcome* that resulted from the conjunction of five people’s *actions*. While we might colloquially say ‘the beachgoers rescued the swimmer’, the non-drowning outcome was rather a happy consequence of a conjunction of individual actions. There were no more actions than these.

The second interpretation says that the responsive beachgoers, as a group, rescued the swimmer.⁷ If the beachgoers did not make up an agent and the rescue was an action, then this view denies that only agents can act. This is motivated by the intuition that the swimmer was intentionally rescued, and so must have been rescued by something. We must weigh up this intuition against the intuition that only agents can act. As I see it, that the first view can account for the rescue—it was an outcome aimed at by several individuals, but not an action—weakens the pull towards denying that only agents can act. Moreover, one who holds the second view is committed to the existence of innumerable acts, performed by only minimally responsive groups of individuals. If we should be parsimonious in our ontology of actions, we should favour the first view. However, someone who believes the beachgoers rescued the swimmer might claim that they constituted an agent. I will now provide reasons why a group-level decision-making procedure is necessary for a group to have agency and to act as a group.

⁵ Tuomela takes his joint intentional action to be more than just individuals’ sharing an aim and taking each other as features of the environment in acting to achieve that aim. As will become clear in §2.2, however, his ‘joint action’ is weaker than my ‘collective agency’.

⁶ I am grateful to an anonymous referee for explicating these implications.

⁷ The views of Tuomela [2006], Bratman [1992], and Gilbert [2006] are potentially consistent with this interpretation.

2.2 Collectives

Consider again the drowning. Suppose that, instead of being able to respond efficiently, each sunbather is wholly unknowledgeable about rescuing swimmers—except Laura. All see the drowning, but only Laura knows what any of the required individual actions are. Indeed, she knows what *all* the required actions are. While the others look around, concerned but puzzled, Laura starts commanding them to drag the boat, pull the starter cord, and so on. Each tacitly commits to abide by her commands and supposes the others have too. Laura divides the necessary actions among the sunbathers. Each performs his or her action because of Laura's commands. The swimmer is saved.

Here, Laura acted as the group's decision-making procedure. This example shows that creating a collective can be extremely easy. A collective has a group-level decision-making procedure—a process that takes in reasons and produces aims and instructions—that is *operationally distinct* from the procedures held respectively by members. It is operationally distinct in that its method for processing reasons is different from the method of any one member when deciding for herself; and in that its output set is not the straightforward conjunction of individuals' output sets.

As another example of operational distinctness, consider a three-person majoritarian democracy that is deciding whether $(P \& (Q \& R))$. If A believes $(P \& (Q \& \sim R))$, B believes $(P \& (\sim Q \& R))$, and C believes $(\sim P \& (Q \& R))$, then, if the group votes on each of P, Q, and R in turn, the group will believe that $(P \& (Q \& R))$ even though none of the members believe this [List and Pettit 2011: 43–7]. Its method for reaching the output (i.e., voting) is distinct from that of each member, and the output itself is not just the conjunction of the members' outputs. More generally, Arrow [1963] famously showed that, under certain conditions, *any* decision-making procedure for aggregating individual judgements into group judgements (when there are more than two possible group judgments) must treat some individuals or issues as more important than others, or must let the group's view on an issue be determined by the group's view on other issues. The point is that some groups can—indeed, must—make decisions that are not merely the sum of members' decisions.

For my purposes, a 'collective' is a group of individuals that have a shared decision-making procedure. A group has such a procedure if enough members:

- (i) have agreed (tacitly or explicitly) to the use of a mechanism that each intends and expects to generate the following outputs:
 - (a) propositions that have the role of collective aim/s within the procedure's processes,
 - (b) individual roles that are each necessary and jointly sufficient to achieve those aim/s, and
 - (c) a division of roles between members, and
- (ii) have committed to abide by the procedure's outputs, and
- (iii) have an expectation, commonly known to be held by each member, that enough other role-bearers will abide by the procedure, and
- (iv) are able to receive sufficient information about the mechanism's outputs for them to abide by it.

A few clarifications. First, being 'committed' requires that members abide by the procedure's outputs *because* they result from a procedure to which they and (they believe) sufficient others have committed. So members might be mistaken about whether they are in a collective, if they incorrectly assume sufficient others have committed. Second, one might 'abide' without acting, for example if a collective aim is to let ϕ happen, and ϕ will happen if and only if no member acts. Third, (i) does not require all members to know what all the outputs are. Fourth, to meet (iii) members need not know each other's identities or roles. Finally, to meet (iv), intra-collective communication might be only unilateral.

These conditions are sufficient for group agency. But they may not each strictly be necessary. We could imagine, for example, a large group in which the expectations in (iii) are not quite common knowledge, or in which one member abides reliably but is not committed. A defence of these precise conditions is beyond the scope of this article. However, they give a sketch of the requirements a group must meet in order to have agency. A reader who endorses a different account can insert their preferred conditions here and will be able to accept my argument in §3, *mutatis mutandis*.⁸

⁸ My conditions are similar to Pettit and Schweikard's [2006: 33] conditions for 'group agents', in which '[f]irst, the members act jointly to set up [which might minimally be 'each commits to'] certain common goals [a minimal goal might be 'that the group instructs members'] and to set up a procedure for identifying further goals on later occasions. Second, the members act jointly to set up a body of judgments for rationally guiding action in support of those goals, and a procedure for rationally developing those judgments further as occasion demands. And third, they act jointly to identify those who shall act on any occasion in pursuit of the goals...' French's [1979, 212] 'corporate persons' must meet much weaker conditions: '(I) an organizational or responsibility flow

On my account, an individual *performs his role* in a collective act, ϕ , when the following hold:

(i)-(iv), and

(v) he is assigned a role, R, by the procedure. R is necessary for ϕ and defines his duties and authorities in relation to ϕ . Achieving ϕ is a collective aim.

(vi) he acts within R.

A collective act is the achievement, through members acting within their roles, of one of the aim-like outputs for the sake of which roles were distributed. Each role will be necessary for the performance of at least one collective act. Thus, his meeting (vi) is necessary for the collective's achievement of one of its aims, that is, for the collective's performance of an act. We can attribute the act ϕ to the collective because performing ϕ was a distinctively collective aim (it was produced by the collective's procedure), the necessary means for which the collective determined and distributed to its constituent parts (members), and the members played the roles they did towards ϕ because the collective so instructed them.

In §1, I said agents act on aims, which are formed by processing reasons. The only groups that can process reasons and form aims in their own right—distinctly from their members—are collectives, i.e., groups with decision-making procedures. Collectives govern their parts (members) by giving them roles, and act through the conjunction of their parts' role-based acts. If my gloss on 'agent' was correct, collectives are the only group agents.⁹ Of course, collectives' acts are constituted entirely by individuals' acts, and their reasons must, at each decision, first be recognised as reasons by the individuals that insert reasons into the procedure. But the procedure can process reasons and produce instructions for members in ways that enable members to act not just on their own inferences and decisions, but as members of a collective and in a manner sanctioned by a decision-making procedure that is different from his or any other member's procedure—that is, in a manner sanctioned by the

chart that delineates stations and levels within the corporate power structure and (2) corporate decision recognition rule(s) (usually embedded in something called "corporation policy").

⁹ See fn. 2. Some will deny that decision-making procedures are sufficient for agency; or that any groups can have distinct decision-making procedures. (For example, Ludwig [2007] gives an individualist semantics for group actions and intentions.) From here, my argument is directed at those who agree that collectives have minimal agency. Conversely, one might hold a more permissive account than mine of the conditions a group must meet to act, bear duties, or be an agent. For example, one might maintain that groups of responsive individuals can act or are agents. Such a person will still find §3's account of collectivisation duties of interest, since he can view them as duties to form whatever kind of group he believes can bear action-demanding duties.

collective agent. When conditions (i)-(vi) are met and a collective aim is thus achieved, then the collective performed the act of fulfilling the collective aim.

2.3 *Implications for Duties*

On the assumption that only moral agents can bear action-demanding moral duties, the only groups that can bear such duties are collectives, specifically collectives whose decision-making procedures can (a) take moral reasons as inputs; (b) process those reasons to form an aim-like output where the aim is acting as a duty demands; and (c) distribute roles to members such that, if each did her role, the duty-demanded aim-like output would reliably be achieved. When the collective has done (a), (b), and (c), it has gone some way to discharging the collective duty. But each member whose role pertains to that duty also has a duty, namely to perform her role. When all this is done, the collective has fulfilled the duty.

One might have a residual concern, mentioned in §1: perhaps non-collective groups can bear duties even if they are not agents and cannot act. Perhaps these non-agents can have duties to do things other than act, such as bring about or see to outcomes. I have accepted that non-agent groups can bring about outcomes. But what would their *duties* to do so amount to? These could not be duties to have the outcome as an aim and distribute roles to members with a view to achieving that aim, since these groups lack these functions. It is hard to see how these non-agent duties would be more than a conjunction of individual duties to perform various acts. These duties would likely be conditional on others' fulfilment of their individual duties, and they would be justified with reference to the desirable outcome. But they would be individual duties nonetheless. If this is right, we can eliminate the group-level duty. (Someone who insists that non-agent groups can bear duties over non-actions can understand §3's argument as an account of when individuals have duties to form the kind of group that can bear duties to, for example, bring about outcomes. For simplicity, however, I will continue to discuss only action-demanding duties and to assume that collectives are their only group bearers.)

Over which acts, then, might a collective bear duties? Assume, following Vranas [2007], that the action-demanding duties an agent has at time t are constrained by what the agent has the ability and opportunity to do at t . That is, duties imply ability. This assumption is not uncontroversial. Avid opponents of it might say that one unknowledgeable beachgoer has a duty to save the swimmer on his or her own despite being unable to do so, perhaps in

the way someone who promised to meet her friend in thirty minutes has a duty to do so despite being sixty minutes away.¹⁰ Yet it is unclear that this ‘duty’ would really be action-demanding rather than an expression of desirability or similar. In what follows, assume at least some duties imply ability and opportunity, and that saving a swimmer is among them.¹¹ (Consistently with ‘duties imply ability’, one might think individuals can have duties to *try* to perform collective acts. I consider this in §3.1.)

Consider ‘can’ and ‘is able to’ to be synonyms. One might think agent A ‘can’ ϕ if it is *possible* that he ϕ s if he tries. That is, A can ϕ if A has a miniscule (objective) chance of ϕ -ing if he tries. In my view, using this account of ‘can’ to constrain duties would be too harsh on agents: if A can have a duty to ϕ if he has only a 0.0001 chance of succeeding if he tries, then he is very likely to do something wrong—even though his failure is beyond his control. We can partially solve this by saying that A ‘can’ ϕ just in case he has a *sufficiently high* likelihood of ϕ -ing, conditional on trying.¹² As a constraint on collectives’ duties, then, ‘the collective can ϕ ’ means ‘if the collective *tried*—that is, if the decision-making procedure distributed roles to members with respect to ϕ and each member attempted to perform their role—then the collective would have a p probability of ϕ -ing, where p is over some threshold’. Even if a group meets (i)-(vi), the group cannot have a duty to ϕ if it is very likely that, if it tries (i.e., has ϕ as a collective aim and members attempt to perform their role), the collective does not ϕ .

Thus ‘the group has a duty to ϕ ’ implies that the group meets two conditions. First, the Agency Condition: the group is a collective, that is, it has a group-level decision-making procedure so can act; and the procedure can process moral reasons. This is based on the account of collective agency outlined in §2.2 and the assumption that only agents can act.¹³ A surprisingly large number of groups to which theorists assign prospective or retrospective responsibilities—such as those whose members have shared interests,¹⁴ shared attitudes,¹⁵

¹⁰ This counterexample to ‘ought implies can’ comes from Sinnott-Armstrong [1984].

¹¹ Vranas [2007] compellingly defends ‘ought implies can’. Graham [2011] denies ‘ought implies can’ for omission-demanding oughts, but accepts it for some action-demanding oughts, akin to those I discuss. I thank an anonymous referee for questioning this assumption.

¹² This solution is only partial because A may have a high chance of ϕ -ing if he tries, yet still try to ϕ and fail for reasons beyond his control. This analysis of ‘can’ sidesteps much important literature on abilities, but is sufficient for current purposes. Brennan and Southwood [2007] give a similar analysis of ‘ ϕ is *feasible* for A’.

¹³ One who rejects my account of collective agency and action can insert her preferred criteria for being eligible to bear duties.

¹⁴ Feinburg [1968: 687] describes a train robbery in which passengers could save the day by acting together. He concludes that ‘the group was collectively but not distributively at fault’. See similarly Held [1970, p.477].

¹⁵ May and Strikwerda [1994] argue that men are collectively responsible for rape. This is distinguishable from the claim that *individuals* are responsible for harms perpetrated by their attitudinal communities, to the extent they share those attitudes. On the latter, see May [1992].

shared culture,¹⁶ or similar incomes¹⁷—seem not to meet the Agency Condition. Second, the Can Condition: the agent is sufficiently likely to ϕ , given that it tries. This is based on the account of ability just given and the assumption that duties imply ability. The upshot is this: we can assign a duty over an act to a conjunction of individuals *only if* the conjunction constitutes a collective that is sufficiently likely to perform the act, given that the collective has the act as an aim and members try to follow the collective’s decision-making procedure.¹⁸

One might query the ‘given that members try’ addendum. After all, members often won’t try. So we cannot assume that members will try when defining a collective’s ability. However, for individuals, it makes perfect sense that A’s not trying to ϕ has no bearing on whether A can ϕ (in the sense of ‘can’ relevant to duty-bearing). If A is a member of a collective we are perhaps more likely to *excuse* A’s not trying, if A believes others will not try and A’s trying will be ineffective unless others try. But since condition (ii) for collective decision-making procedures says that A has committed to the procedure, the assumption that A tries is natural when we are thinking about whether the collective can ϕ .

When a collective meets the Can Condition, the chance that the collective will successfully ϕ (given trying) is over some threshold. Context determines this threshold. Context includes ϕ ’s importance: if ϕ is more important, then the chance that the collective will ϕ if it tries can be lower. If, for example, ϕ is ‘detering a giant asteroid from hitting Earth’, then a collective might meet the Can Condition, even if it has an (absolutely speaking) low chance of deterring the asteroid if it tries. This may not be the only factor that determines the Can Condition’s likelihood threshold, but it indicates the kinds of considerations involved.

Importantly, sometimes the moral importance of (what looks like) a group action can give an *individual* a duty to perform it through responsiveness to others. Suppose Ben, Jon, Julie, and Stan are oblivious to the drowning and just happen to be taking the boat to where the swimmer is. Suppose Laura is strong enough to pull the swimmer to safety alone. Laura can respond by jumping in the boat and doing the pulling. If this is likely to result in the swimmer’s not drowning, Laura meets the Can Condition for rescuing to swimmer and can incur a duty to do so. Her duty to rescue the swimmer entails that she alone, by treating others as environmental objects and acting responsively to them, meets the Can Condition for doing

¹⁶ Miller [2004] argues that nations (cultural groups distinct from states) can be responsible for outcomes.

¹⁷ Cullity [2004: ch. 4] argues that the affluent ‘collectively’ have a duty to alleviate world poverty.

¹⁸ Of course, we often cannot know a collective’s likelihood of ϕ -ing. Such epistemic limitations are a problem for all moral philosophy. I take it our best theory (whatever that turns out to be) of moral duties under uncertainty can apply here. Presumably, it will refer to the reasonable beliefs that some actual or ideal agent does or should or would have about the collective’s likelihood of ϕ -ing. For discussion, see Smith [2010].

so—which entails that others’ actions are sufficiently fixed. If she succeeds, the action is not a group action, but an individual action. Her responsiveness is likely to fail if she makes an incorrect guess about what others doing, or what she should do in response. Thus, her meeting the Can Condition may require a low likelihood threshold.

Other (what look like) group actions might necessitate responsiveness that is more than minimal—such as mutual responsiveness between several individuals with overlapping, commonly known intentions. Plausibly, if such responsiveness were sufficiently likely to result in the non-drowning outcome, then each individual beachgoer would have a duty to be responsive in this way. Such duties seem to exist in the first beach drowning, where each beachgoer was knowledgeable. There, it was highly likely that the swimmer would cease drowning if each was responsive. But these are not individual duties *to rescue the swimmer*, since then success would mean that each individual performed a distinct act of rescuing the swimmer. Rather, they are duties to act responsively (presumably, conditional on others so acting), with a view to bringing about the non-drowning *outcome*. Each beachgoer’s responsive action meets the Can Condition for ‘through responsiveness, making a contribution to the non-drowning outcome’s being realised’. The contributions in aggregate meet the likelihood threshold for the outcome’s being realised. This gives each individual a duty to be responsive.¹⁹ Similarly, there are individual duties but no collective duty if: (i) each of several individuals would ϕ if they tried, but if more than one tried then all would fail—then, each has an *individual* duty to ϕ until someone else tries to ϕ (at which time the others’ duties are dissolved); or (ii) ϕ requires actions from multiple individuals, but they can reliably perform these actions without needing responsiveness or collectivisation—then, each has an *individual* duty to perform her independent action.

A problem remains: sometimes, even if it is extremely likely that all will be responsive with overlapping, commonly known intentions, responsiveness does not meet the likelihood threshold. Consider the unknowledgeable beachgoers. Plausibly, even if all are responsive, it is very likely that someone will make a tragically incorrect or belated guess about others’ actions or his appropriate response. The non-drowning outcome is insufficiently likely to result from responsiveness to generate individual responsiveness duties. Or suppose all the beachgoers are experienced rescuers, but the rescue’s temporal length, spatial distance, and roles’ complexity, variety, and synchronisation makes the success of even these people’s responsiveness a terribly tall order. Or consider the Australian Volunteer Coast Guard, and

¹⁹ Someone who maintains that the knowledgeable beachgoers performed a group act could presumably assign the rescue duty to the group. I have already provided reasons against this.

ask whether the outcomes it produces would have a hope of resulting from the acts of the best-placed responsive individuals who weren't subject to a group-level decision-making procedure. Sometimes, the Can Condition requires a group-level procedure for distributing instructions and planning for contingencies. And for all actions, a *group-level* duty requires a group that meets the Agency Condition.

The problem for the remainder of the article is that sometimes the likelihood threshold is not met by an existing collective, or by would-be responsive individuals' aggregated actions. This appears to imply there can be no individual or collective duties in these circumstances. I will argue that there is often more to be said.

3. Collectivisation Duties

3.1 Generating Collectivisation Duties

The unknowledgeable beachgoers appear permitted to go about their way: the non-drowning outcome is insufficiently likely to result from responsiveness to generate individual responsiveness duties; and the group cannot bear duties because it is not an agent. But this cannot be right. Someone is drowning. If the beachgoers follow Laura's instructions, the person could be saved. To avoid this result, we need collectivisation duties. These are individual duties to *be responsive to others with a view to there being a group that meets Agency and Can Conditions*. Collectivisation duties thus differ from the duties of the knowledgeable beachgoers, which were to be responsive to others with a view it being the case that the swimmer no longer drowns. Only once the unknowledgeable beachgoers' collectivisation duties have been discharged can there be a duty to rescue the swimmer, held by the newly-formed collective agent that *can* rescue the swimmer.

One might think deriving collectivisation duties is simple: all we need do is build on the proposition that 'A (the individual agent) has a duty to secure the means to discharging his duties'. If A has a duty to ϕ (where ϕ requires a collective), and ψ is the best means to ϕ (where ψ is A's individual responsive collectivising action), then A has a duty to ψ . This appears consistent with 'duty implies ability'. After all, A *can* ϕ in the long-term: A can ψ and then A can ϕ through the collective. Because he can ϕ in this two-step way, he cannot

avoid the duty to ϕ by pointing out that he cannot ϕ at t_1 . At t_1 , he *can* bring it about that he *can* ϕ at t_2 . So, at t_1 , he is eligible for a duty to ϕ .²⁰

However, this almost always runs afoul of ‘duty implies ability’. If ϕ is a collective’s action, then A often cannot ϕ through the collective at t_2 , even if (at t_1) A ψ s and others’ ψ -ing is very likely. This is because the collective’s t_2 actions are not performed *by A*. At t_2 , A can do only his role in the collective ϕ -ing. He cannot ϕ himself. This is perhaps false of Laura, the beachgoers’ dictator. If at t_1 it is very likely that others will collectivise, then perhaps at t_1 Laura has an individual ability to rescue the swimmer: she is able to manipulate others as if they were features of the environment, so that they perform the necessary roles. Describing this as an individual ability is peculiar, perhaps because of the ongoing possibility of others’ defection. Yet suppose we grant this. Then, at t_1 , Laura is eligible for a duty to rescue the swimmer by taking responsive steps to form the group—but Ben, Jon, Julie, and Stan are not. So what can we say about their duties, and more generally about the duties of any non-collectivised group of individuals who do not have a clear dictator?

Consider unknowledgeable Ben. He is unable, at t_1 , to make himself such that he can rescue the swimmer through the collective, since he lacks expert control over the collective’s decision-making procedure. If Ben—as an individual—cannot bring it about that he can ϕ through the collective, then Ben cannot have a duty to ϕ through the collective. If Ben does not have a duty to ϕ through the collective, then he cannot have a duty to take responsive collectivising actions to *form* the collective, because the only reason he would have had this duty was because this would have been a necessary condition for his ϕ -ing (over which he would have had a duty). Ben is thus not going against his duties if he refrains from collectivising.

Some may yet be unconvinced. One might think an agent can have a duty to *try to* perform an action he has a low likelihood of performing if he tries, if the action is important enough. Then each unknowledgeable individual could have a duty to try to ϕ . A duty to try would plausibly amount to a duty to perform those actions that most efficaciously increase the likelihood of one’s ϕ -ing. We then straightforwardly get individual duties to collectivise, if collectivisation is the most efficacious way for individuals to increase the likelihood of their ϕ -ing.²¹

²⁰ In Jensen’s [2009] terms, A is indirectly diachronically able to ϕ : A can ϕ later, provided he performs an enabling action first. See also Goldman [1978].

²¹ I thank Christian Barry for pressing this possibility.

However, a plausible condition for an agent being *able* to try to ϕ is that he believes it is possible that he will ϕ .²² If this condition is correct, then only individuals satisfying it can have a duty to try to ϕ . Of course, each beachgoer might believe ‘it is possible that I will ϕ if I try to ϕ by ψ -ing, since possibly then others will also ψ and then I will ϕ ’. The individual’s belief that *he or she will ϕ* if others ψ is false (perhaps unless she is Laura), because ϕ is a collective’s action. Still, each beachgoer might falsely believe this. Then, each could have a duty to try to rescue the swimmer. But these duties would not be to rescue the swimmer, only to try to do so. My points would still hold that duties cannot be held by non-collective groups and that individuals cannot have duties *to perform*, rather than try to perform, collective acts. The advocate of individual duties to try to ϕ must, then, accept that ‘the beachgoers have a duty to rescue the swimmer’ is false, if we interpret it as attributing a duty to a group. There are only *individual* duties to transform the group—specifically, where responsiveness is insufficient, to transform it into a collective. This is the most coherent way to understand the ordinary locutions with which §1 started.

These collectivisation duties—individual duties to ψ —cannot depend on each individual having a duty to ϕ , and they should not exist when individual responsiveness alone is best.²³ To establish whether these exist in a given context, we must consider the abilities that would be had by a collective that could result from the individuals’ responsiveness. After all, there is no point in positing collectivisation duties if the resulting collective would be unable to ϕ . I propose the following basis for collectivisation duties:

Criteria for Collectivisation Duties (CCD):

If:

(1) ϕ is morally pressing, and

²² I am grateful to Nic Southwood for this suggestion.

²³ As an anonymous referee pressed, the possibility remains of a general solution, generating both responsiveness and collectivisation duties. Any general solution would need to compare the desirability of individual responsive action with the desirability of collective action, if it is to distinguish when we have responsiveness duties from when we have collectivisation duties. To do this, I suspect it would need something like my Criteria for Collectivisation Duties. I am thus sceptical that a general solution could be less complex while illuminating the distinction between duties to be responsive with a view to a certain outcome and duties to be responsive with a view to there being a collective that can produce the outcome (where the outcome justifies the collectivising duties).

- (2) at t_1 , either: no (collective or individual) agent/s have duties, either to φ or to take responsive actions with a view to there being the morally desirable outcome that φ produces;²⁴ or too many agent/s with such duties default, and
- (3) if, at t_1 , A, . . . , N each took responsive steps towards there being a collective-that-can- φ , then, at t_2 , that collective would incur a duty to φ , and
- (4) at t_1 , A, . . . , N are each able to take responsive steps towards there being a collective-that-can- φ at a reasonable expected personal and moral cost, and
- (5) other individuals will not successfully take responsive steps towards there being a collective that will incur a stronger duty to φ ,

then:

- (6) at t_1 , A, . . . , N each have a duty to take responsive steps towards there being a collective-that-can- φ , and
- (7) at t_2 , once a {A, . . . , N} collective-that-can- φ is formed, that collective has a duty to φ , which entails
- (8) at t_3 , once the collective has distributed φ -related roles, each member with a φ -related role has a duty to perform that role.

Condition (3) does much of the work here. Instead of deriving Ben's collectivisation duty from his duty to perform a collective act, we derive it from a counterfactual about what would happen if Ben and a number of others each took responsive, collectivising actions: there would be a duty-bearing agent. Different ethical theories will justify (1) and (3) in different ways. It might be important that a certain outcome obtains, or that these particular individuals do the right thing or live virtuous lives, or that someone's rights are fulfilled. My account is agnostic about which moral considerations make φ pressing, and give a collective a duty to φ .

3.2 Implications for Individuals

²⁴ This outcome could be: that φ occurs.

There remains a problem. If Ben's collectivisation will result in there being a collective-that-can- ϕ only if others collectivise, does Ben have a collectivisation duty if others defect? Here, there seems no point in Ben's collectivising. Perhaps, then, he is off the hook if he reasonably believes others will defect on the duty in (6). However, consider what is entailed by a duty to be responsive to phenomena in one's environment with a view to there being a collective. The duty plausibly entails trying to affect the phenomena so as to make the collective's existence more likely. If the phenomena are agents with collectivisation duties, this will plausibly include convincing them that they ought to collectivise. Moreover, if the phenomena are agents, consider that I have said responsiveness entails responding *positively* to their perceived interest. Suppose Julie defects, but Ben reasonably believes collectivising is in her interest. After all, it is a 'tick' in her moral ledger and, since (4) holds, is not too costly for her. Ben's responding positively to her perceived interests then includes trying to induce her to collectivise. This is all part of Ben's collectivising duty. However, if Julie had done her duty from the beginning then his duty would have been less demanding—so she should shoulder some of this cost.

This is worth making explicit:

If,

(1)– (6), and

(9) a subset of $\{A, \dots, N\}$ defects on the duty in (6), such that the remaining set of individuals ('compliers') do not meet (3) and (4),

then:

(10) each complier has a duty either to induce defectors to comply or to contribute to discharging the collectivisation duties on which defectors have defected, and

(11) each defector has a duty to compensate some compliers, each defector incurring cost up to the amount that defector would have incurred in collectivisation and the subsequent collective's action.

There is more to say about collectivisation duties. First, there are many ϕ s for which individuals meet the CCD. If the CCD produced all-things-considered duties, discharging them would require that individuals ignore other important goods. Importantly, collectivisation duties are *pro tanto*, to be weighed against other moral values.

Second, suppose there are seven beachgoers and exactly five are needed. According to the CCD, each of the seven has a collectivisation duty. Because exactly five are needed, it might appear that the CCD is wrong here. But we can say (consistently with the CCD) that all seven have a collectivisation duty, *until* five have discharged the duty. Then, (2) is false: a collective agent has, and is discharging, a duty to ϕ . The respective collectivisation duties of the remaining two are then dissolved.

Third, consider appraisal of individuals that produce a morally desirable outcome through insufficiently reliable responsiveness. I have said the unknowledgeable beachgoers' respective responsiveness would not meet the likelihood threshold. But suppose they are responsive anyway and chance upon one of the unlikely futures in which their responsiveness succeeds. My account says that their duties were to take responsive steps towards there being a collective that can rescue the swimmer, not to take responsive steps towards the non-drowning outcome. Thus my account says they did not do their duties. This, one might think, is counterintuitive: after all, the morally pressing need is met. Yet I am happy with this result. The unknowledgeable beachgoers did not do their respective duties. While we should be happy that their responsiveness produced this result, we should disapprove of how they achieved it. They each took a gamble. They had a duty not to gamble with the swimmer's life in this way, despite the gamble having paid off.

Of course, this is false when taking responsive steps towards some outcome *is* better than collectivisation. But then collectivisation incurs too much cost for (4) to be met. When mere responsiveness would be better than collectivising, collectivisation incurs too high an expected moral cost. This is true if, for example, collectivising would take too long and responsiveness would be sufficiently fast. Yet the beach example illustrates just how cheap and fast it can be for a collective to result from responsiveness. It took merely each individual's willingness to help impose order on the situation. Collective decision-making procedures might have just one aim; members' agreement and commitment might be non-verbal; a role might just be 'pull the starter cord'. Even in emergencies, then, collectivising might not be prohibitively costly.

Fairness between collectivisers is a fourth implication. Imagine a succession of drowners, where a new group of five is required for each one. If Ben collectivises for the

first, second, and third drowners, Ben might have done his ‘fair share’ and not have further collectivisation duties. But the CCD forces Ben to collectivise in each drowning. This was intentional: in each drowning, Ben has a duty. Yet we can say that his duties in later drownings are weaker based on his prior compliance and others’ are stronger based on their prior non-compliance.

Finally, suppose I could perform one action if I acted alone, but myself and others could produce a collective able to perform that action plus others, or that action more efficiently. Do we have collectivisation duties? Under the CCD, the answer depends on how we describe ϕ : there is the ϕ I could perform on my own, and the (conjunctive or efficiency-indexed) ϕ that the collective could perform. To which do I have a duty to attend? This is an all-things-considered question, on which the CCD is silent. If there is a morally pressing ϕ of any description (including a conjunctive or efficiency-indexed ϕ) for which (1)–(5) are met, then individuals each have *pro tanto* duty to collectivise. Whether this outweighs other individual duties is a separate question.

4. Conclusion

How should we construe ‘the beachgoers have a duty to save the drowning person’? In §2, I argued that it largely depends on whether the beachgoers constitute a collective. Collectives, I claimed, are agents in virtue of their possession of a decision-making procedure, which enables them to act through the conjunction of their members’ actions. If a collective’s procedure can process moral reasons, it meets the Agency Condition and can have duties to act. I explained that, for determining which duties a collective has, we should consider the likelihood that it will perform certain actions if it had them as an aim and if members abided by the decision-making procedure. The actions the collective is sufficiently likely to perform under these two conditions are the actions it is ‘able’ to perform; the actions for which it meets the Can Condition. I used the Agency and Can Conditions to argue that the group of beachgoers has a duty to rescue the drowning person only if it is a *collective* that *can* save the drowning person. Otherwise, when we say the beachgoers have the duty, this could mean that the individual beachgoers each have a duty to be responsive to others with the aim of producing a non-drowning outcome. However, these duties require the other beachgoers’ actions to be sufficiently reliable. In §3 I argued that we need collectivisation duties, since otherwise there are no duties when only a collective could perform some morally desirable

action or produce some morally desirable outcome, and no able collective exists. Collectivisation duties are individual duties to take responsive steps towards there being a collective that has a duty over that action or outcome. I provided a set of conditions for such duties, called the Criteria for Collectivisation Duties, and discussed some of their implications.

The upshot is this. The group of unknowledgeable beachgoers does not have a duty to rescue the swimmer. Rather, each beachgoer has an individual duty to act responsively to others with a view to there being a collective that can rescue the swimmer, if such a group would incur a duty to do so. This matters because the actions that the beachgoers (and non-collectivised individuals generally) each have a duty to perform—and the actions they are blameworthy for not performing—are likely to be different from those they would each have a duty to perform, if each had an individual duty to rescue the swimmer.²⁵

Australian National University

References

- Arrow, Kenneth 1963. *Social Choice and Individual Values*, 2nd ed., New York: Wiley.
- Bratman, Michael 1992. Shared Cooperative Activity, *The Philosophical Review* 101/2: 327–341.
- Brennan, Geoffrey and Nicholas Southwood 2007. Feasibility in Action and Attitude, in *Hommage a Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*, ed. Toni Rønnow-Rasmussen, Björn Petersson, Jonas Josefsson, and Dan Egonsson, URL = < <http://www.fil.lu.se/hommageawlodek/site/papper/Brennan&Southwood.pdf> >.
- Cullity, Garrett 2004. *The Moral Demands of Affluence*, New York: Oxford University Press.
- Feinberg, Joel 1968. Collective Responsibility, *Journal of Philosophy* 65/21: 674–88.
- French, Peter 1979. The Corporation as a Moral Person, *American Philosophical Quarterly* 16/3: 207–15.
- French, Peter 1984. *Collective and Corporate Responsibility*, New York: Columbia University Press.

²⁵ For discussion of this article, I am grateful to audiences at the 2011 Australasian Postgraduate Philosophy Conference (Monash University) and the 2011 meeting of the Australasian Association of Philosophy (University of Otago). For extensive comments on written versions, I thank Bob Goodin, Nic Southwood, Holly Lawford-Smith, Christian Barry, Seth Lazar, Jonathan Farrell, and two anonymous referees.

- Gilbert, Margaret 2001. Collective Preferences, Obligations and Rational Choice, *Economics and Philosophy* 17/1: 109–119.
- Gilbert, Margaret 2006. Who's to Blame? Collective Moral Responsibility and Its Implications for Group Members, *Midwest Studies in Philosophy* 30/1: 94–114.
- Goldman, Holly S. 1978. Doing the Best One Can, in *Values and Morals*, ed. Alvin Goldman and Jaegwon Kim, Boston: Reidel.
- Graham, Peter A. 2011. Ought and Ability, *The Philosophical Review* 120/3: 337–382.
- Held, Virginia 1970. Can a Random Collective Be Morally Responsible? *Journal of Philosophy* 67/14: 471–81.
- Jensen, Mark 2009. The Limits of Practical Possibility, *The Journal of Political Philosophy* 17/2: 168–184.
- Kutz, Christopher 2000. Acting Together. *Philosophy and Phenomenological Research* 61/1: 1–31.
- Lawford-Smith, Holly 2012. The Feasibility of Collectives' Actions, *Australasian Journal of Philosophy* 90/3.
- List, Christian and Philip Pettit 2011. *Group Agency*, Oxford: Oxford University Press.
- Ludwig, Kurt 2007. Collective Intentional Behavior from the Standpoint of Semantics, *Nous* 41/3: 355–93.
- May, Larry 1992. *Sharing Responsibility*, Chicago: University of Chicago Press.
- May, Larry and Robert Strikwerda 1994. Men in Groups: Collective Responsibility for Rape, *Hypatia* 9/2: 134–51.
- Mellema, Gregory 2006. Collective Responsibility and Qualifying Actions, *Midwest Studies in Philosophy* 30/1: 168–175.
- Miller, David 2004. Holding Nations Responsible, *Ethics* 114/2: 240–68.
- Pettit, Philip 2007. Responsibility Incorporated, *Ethics* 117/2: 141–201.
- Pettit, Philip 2010. Groups with Minds of Their Own, in *Social Epistemology: Essential Readings*, ed. Alvin I. Goldman and Dennis Whitcomb, Oxford: Oxford University Press.
- Pettit, Philip and David Schweikard 2006. Joint Actions and Group Agents, *Philosophy of the Social Sciences* 36: 18–39.
- Sinnott-Armstrong, Walter 1984. 'Ought' Conversationally Implies 'Can', *The Philosophical Review* 93/2: 249–61.
- Smith, Holly 2010. Subjective Rightness, *Social Philosophy and Policy* 27/2: 64–10.
- Sudgen, Robert 2000. Team Preferences, *Economics and Philosophy* 16/2: 175–204.

- Tuomela, Raimo 2006. Joint Intention, We-Mode and I-Mode, *Midwest Studies in Philosophy* 30/1: 35–58.
- Vranas, Peter B.M. 2007. I Ought, Therefore I Can, *Philosophical Studies* 136/2: 167–216.
- Wendt, Alexander 2004. The State as Person in International Theory, *Review of International Studies* 30/2: 289–316.