

## ARTICLE

# Group motivation

Jessica Brown

Arché Research Centre, Department of Philosophy, University of St Andrews.

## Correspondence

Jessica Brown, Arché Philosophical Research Centre. Department of Philosophy, University of St Andrews, Edgecliffe, The Scores, St Andrews, Fife, KY16 9AR, Scotland.

Email: [jab30@st-andrews.ac.uk](mailto:jab30@st-andrews.ac.uk)

## Abstract

In this paper I discuss a key issue for group moral responsibility, namely whether we can make sense of a group acting for one reason rather than another. The notion of acting for one reason rather than another is central to standard accounts of individual agency and responsibility; and also determines whether an individual is blameworthy or praiseworthy for an action. Thus if we model group responsibility on individual responsibility, we need to be able to make sense of a group acting for one reason rather than another. In this paper, I raise problems for both summative and inflationary accounts of what it is for a group to act on a reason, before suggesting several potential solutions at the end.

## 1 | INTRODUCTION

We routinely treat groups, including governments and corporations, as agents with beliefs and aims who are morally responsible for their actions. For instance, we might blame an oil company for an oil spill pointing out that they knew the risk of their profits-first policies. Many philosophers take it that such ascriptions are literally true. They argue that not only individuals, but also groups, can have mental states, act, and be morally responsible (e.g. Gilbert 1989, Tollefsen 2003 and 2015, Mathiesen 2006, Isaacs 2011, List & Pettit 2011). In doing so, they go beyond the uncontroversial claims that groups can be causally and legally responsible; or that it's instrumentally useful to hold groups morally responsible. For these phenomena can come apart from moral responsibility. For

This is part of a project on group responsibility funded by the Leverhulme Trust. Various versions of this paper were presented at Washington University St Louis, Zürich, the Roots of Responsibility Project and St Andrews. Thanks to these audiences and colleagues for helpful feedback and comments, especially to Sandy Goldberg, Katherine Hawley, Allan Hazlett, Matthew McGrath, and Justin Snedegar.

-----  
This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. *Noûs* published by Wiley Periodicals LLC

example, a young child might be causally responsible for damage to a neighbour's property but neither morally nor legally responsible.

In this paper I discuss a key issue for group moral responsibility, namely whether we can make sense of a group acting for one reason rather than another. The notion of acting for one reason rather than another is central to standard accounts of individual agency and responsibility; and also determines whether an individual is blameworthy or praiseworthy for an action. Thus if we model group responsibility on individual responsibility, we need to be able to make sense of a group acting for one reason rather than another. Despite the importance of the notion of acting for a reason, most of the literature on group responsibility has focused on other issues. Here, I start to consider how best to understand what it is for a group to act for a reason.

In the next section (2), I briefly review the importance of the notion of acting for a reason for standard accounts of agency and responsibility, both at the individual and group level. In subsequent sections, I turn to consider the main options for providing an account of group action for a reason, namely summative and inflationary accounts. I argue that major objections face each style of approach (3-5). At the end, I discuss a variety of ways of attempting to defend our treatment of groups as morally responsible despite these results.

## 2 | ACTING FOR A REASON AND GROUP MORAL RESPONSIBILITY

In addressing the moral responsibility of groups, I will focus on the kind of group which defenders of group moral responsibility take to be the best candidate for moral responsibility, namely organised groups with well-ordered decision procedures such as companies or governments (as opposed to, say, a group composed of all left-handed Scottish people).<sup>1</sup> I will assume in what follows that such groups have intentional states to focus on whether they are morally responsible. These are separate issues. It's possible to have intentional states and yet not be morally responsible (consider very young children or animals). In addition, even granting that groups have intentional states, a variety of objections have been raised to the claim that they are morally responsible, including 1) that treating them as morally responsible would have problematic consequences for individual responsibility<sup>2</sup>; 2) that groups lack the capacity to have certain kinds of mental states/attitudes required for responsibility, e.g. the capacity to care, have good or ill will, or have emotions (e.g. Haney 2004); and 3) that groups lack certain cognitive capacities required for responsibility, e.g. to understand the meaning that their acts convey (McKenna 2006). Here I focus on a different concern to these established objections to group moral responsibility, that it's hard to give an account of what it is for a group to act for one reason rather than another. As we will see, the notion of acting for a reason is central to standard accounts of agency and responsibility. So it is a problem for group responsibility if we cannot make sense of what it is for a group to act for one reason rather than another.

On the standard view of individual agency, associated with Davidson, the category of action is understood by appeal to intentional action, explained as action for a reason. In addition, whether a subject is morally praiseworthy or blameworthy for an action depends on the basis of that action.<sup>3</sup> In Kant's example, if a grocer gives his customers the correct change merely because of his desire to increase profits and his belief that it will increase his profits, then he is not praiseworthy for giving them the correct change. Further, influential approaches to moral responsibility appeal to the notion of acting for a reason. This is perhaps most obvious in those accounts which take it that there are substantial conditions for moral responsibility. Standardly, these accounts require that the action was under the agent's control (the control condition) and that the agent was aware

of what she was doing (the epistemic condition). The control condition is frequently understood as requiring the ability to perform actions for, or on the basis of, reasons. For example, many hold that the kind of control required involves “reasons-responsiveness” which in turn requires the ability to perform “the relevant action intentionally (i.e., for a reason)” (Fischer & Ravizza 1998: 64; see also Pettit 2007:175). In addition, the epistemic condition is standardly understood as allowing that a subject may be excused for wrongdoing if she’s not only blamelessly ignorant that she’s doing wrong but also acted from that ignorance.<sup>4</sup>

The connection between responsibility and acting for a reason is perhaps less obvious in those accounts, inspired by Strawson, which hold that facts about who is morally responsible are dependent on our practice of holding people responsible (e.g. Wallace 1994). But, as Strawson himself noted, the reactive attitudes, such as resentment, indignation and gratitude, “are essentially natural human reactions to the good or ill will or indifference of others towards us, as displayed in their attitudes and actions” (2003: 80). And whether some action displays good or ill will in part depends on the reason for which it was done (for instance whether your treading on my toes was done from a desire to hurt me, or to alert me to a greater danger, or simply by mistake when dancing).

Given the centrality of the notion of acting for a reason to accounts of individual agency and responsibility, it’s unsurprising that some accounts of group agency and responsibility explicitly appeal to the notion of acting for a reason. French (1995) claims that agency requires the capacity “to act intentionally – to do something on purpose or for a reason” (10). Pettit argues that for a group to be morally responsible it must have “the control necessary for being able to choose between options of the basis of judgements about their value” (2007:175). Tollefsen (2015) says: “to ask if groups can have mental states is to ask if they are agents that act for reasons” (53). Indeed, it seems plausible that whether a group is blameworthy or praiseworthy for some action depends on the reason for which the group acted. For instance, if we think that a company’s action of helping the environment is motivated merely by financial interest, and not concern for the environment, we don’t think they are praiseworthy for helping the environment. Instead, we might accuse them of “greenwashing”.

Despite the importance of the notion of group action for a reason, there has been relatively little discussion of what makes it the case that a group acts for a reason as opposed to other issues, such as what it is for a group to have a mental state or act (though one exception is Pettit’s (2007) appeal to programming discussed in section 5 below). In the rest of the paper, I consider the main options for providing an account of group action for a reason.

### 3 | A SUMMATIVE ACCOUNT

Summative approaches hold that a group has a certain property if and only if some of its members have that same property. For instance, on a summative approach to group belief, a group believes that *p* if and only if certain members of that group believe that *p*. By contrast, on an inflationary approach to group belief, a group can believe that *p* even if none of its members believe that *p*. To investigate the notion of group action for a reason, let’s start with a summative account of what it is for a group to act for a reason.

Suppose that a group performs some action  $\alpha$  where the group’s doing  $\alpha$  is constituted by one or more members of the group doing  $\alpha_{1-n}$ . For instance, perhaps a string quartet gives a charity performance for Cancer UK, where the string quartet’s performance is constituted by the playing of its individual members. On a summative approach, the basis of the string quartet’s performance

is determined by the basis of the actions of its individual members which constitute that performance. Thus, the string quartet's performance is based on the desire to raise money for Cancer UK if and only if the actions of its members which constitute its performance are based on that desire. More formally, these reflections may encourage the following simple summative account of what it is for a group to act from an aim:

SA): A group's acting, which is constituted by some of its members doing  $\alpha_1, \alpha_2, \alpha_3 \dots \alpha_n$ , is based on the aim, A, if and only if each of those members who do  $\alpha_1, \alpha_2, \alpha_3 \dots \alpha_n$  have aim A and perform the relevant actions on the basis of A.

Such a summative account may be inspired by the thought that a group acts only through the action of its members. Thus, to see the group's motivation for its action, we need to look at the motivation of the actions of its members which constitute the group's action.

However, the simple summative account runs into obvious problems when we consider structured groups which are our focus here. In such groups, such as governments, corporations or charities, different members have different roles and responsibilities, and they may play their part in the group's action for very different reasons. Indeed, some members may act from motivations which are rogue with respect to the broader group. For example, suppose that evil individuals set up a charity (let's call it, "Bad"), with the sole end of drawing dispossessed refugees into terrorism. One of the ways in which they do so involves fostering relations with dispossessed refugees by providing them with food before moving on to inculcate them into their violent form of political extremism. While most of Bad's members and all of its steering committee have terrorist aims, it sometimes manages to co-opt as members well-meaning individuals who lack terrorist aims. For instance, suppose that Bad feeds a particular homeless refugee, Reg, where Bad's feeding Reg is constituted by the activity of one of its workers, Innocent. We may suppose that Innocent is entirely ignorant of the terrorist ends of other group members and feeds Reg out of a sincere wish to help the needy. The proposed summative account has the implausible result that Bad feeds Reg out of a sincere desire to help the needy. By the description of the case, Bad's action of feeding Reg is constituted by Innocent's feeding Reg where Innocent is motivated by a sincere wish to help the needy. So by the sufficiency direction of SA), it follows that Bad's action of feeding Reg is motivated by a sincere wish to help the needy. But, that's implausible given the evil intentions of most of the members of the charity including its steering committee.<sup>5</sup>

The proposed simple summative account not only faces problems in the sufficiency direction but also the necessity direction. To see this, consider a different charity, "Good", most of whose members, and all of its steering committee, are motivated by the aim of helping the needy. Despite this, some of its members join the charity from other motivations, whether to impress others, or perhaps even exploit the charity for their own ends. On a particular occasion, Good feeds a homeless woman Heather where the charity's feeding Heather is constituted by one of its workers, Proudie, feeding Heather. However, Proudie's membership of the charity, and his feeding of Heather on this occasion, is not motivated by a sincere desire to help the needy, but instead motivated by self-interest. As a result, in this example, Good fails to meet the proposed necessary condition for feeding Heather on the basis of a sincere desire to help the needy. For, Proudie, whose action constitutes Good's feeding Heather is not motivated by a sincere desire to help the needy but instead self-interest. However, it seems plausible that the action of Good in feeding Heather is indeed motivated by the sincere desire to help the needy. For its part of the case that most of the members of this charity, and all of its steering committee, are motivated by the aim of helping the

needy. Indeed, we may suppose that many members make personal sacrifices to participate in the charity's work.

We can reinforce the criticism of SA) by considering its consequences for whether the two charities count as praiseworthy or blameworthy for their actions. According to a right reasons approach to moral worth, whether an agent is blameworthy or praiseworthy for an action depends on their motivation. In particular, an agent is praiseworthy for doing the right thing if and only if they do it for the reasons which make it right. However, perversely, SA) holds that Bad, but not Good, does the right thing for the right reason. In particular, SA) holds that Bad feeds Reg out of a sincere wish to help the needy. As a result, on a right reasons approach to moral worth, Bad would be praiseworthy for feeding Reg. By contrast, SA) has the result that Good does not feed Heather from a sincere wish to help the needy. Thus it isn't praiseworthy for feeding Heather. But surely these are just the wrong results about this pair of charities.

So far we've seen that the simple summative account faces problems for both its necessity and sufficiency direction even for the simplest kind of case in which a group's doing  $\alpha$  is constituted by the action of a sole member of the group. The example of Bad illustrates the way in which the basis of a group's actions are intuitively not determined by the motivations of the footsoldiers of the group whose actions constitute the group's action. I now consider three replies to this objection to summative accounts. The first attempts to deal with the objection by reference to the aims of the relevant groups; the second attempts to do so by focusing on the motivations of key members rather than footsoldiers; and the third attempts to do so by suggesting that the set of actions which constitutes the group's action is broader than has so far been suggested. I will discuss each in order.

First, one might wonder if one could get around some of the problems for the summative account by appeal to the aims of the relevant charities. After all, it may be said that the problems arise because the summative account doesn't pay attention to the aims of the relevant charities, as opposed to those of its individual members. For instance, it may be suggested that we would get the right result on the following revised summative account which adds in appeal to the aim of the group:

SA\*): A group's  $\alpha$ ing, which is constituted by some of its members doing  $\alpha_1, \alpha_2, \alpha_3 \dots \alpha_n$ , is based on the aim, A, if and only if 1) the group has the aim A and 2) each of those members who do  $\alpha_1, \alpha_2, \alpha_3 \dots \alpha_n$  have aim A and perform the relevant actions on the basis of A.<sup>6</sup>

Of course, SA\*) won't help with the worries about the necessity direction of SA) since they arose from the intuition that a group could perform an action on the basis of some aim even though that aim is not shared by the member whose action constitutes the group's action.<sup>7</sup> But it may be argued that the revised account helps with at least the sufficiency direction of SA). In our original counterexample, we saw that it's not sufficient for the evil charity, Bad, to feed the refugee Reg on the basis of a sincere desire to help the needy that the group's action of feeding Reg is constituted by the action of Innocent who is indeed motivated by the sincere desire to help the needy. It might be thought that this counterexample doesn't work against SA\*) which adds the condition that the relevant aim be the aim of the group. For, plausibly the evil charity Bad simply doesn't have the aim of helping the needy but only the aim of promoting its evil terrorist ends.

However, I will argue that even if the original Bad example doesn't provide a counterexample to the sufficiency direction of SA\*), a variant example will work instead. To set up the objection, notice that an individual can have some motivation and yet not act on it on a particular occasion. For example, we can imagine that an individual both desires to help the needy and also

desires her own financial self-interest. Nonetheless, even if on some occasion she does help a rich elderly neighbour, it's an open question whether her motivation was her desire to help the needy or instead her own financial self-interest (if helping is likely to increase the likelihood of a substantial legacy). Similar points apply to the case of groups and will undermine the sufficiency direction of SA\*). For example, suppose that a church organisation both desires to help needy people and also desires its own financial self-interest. Suppose that on some occasion the church does help a needy person where the group's action is constituted by some member of the group helping a needy person, Caroline. Just as in the case of individuals, it's an open question whether the group acts on its desire to help needy people or instead on its desire for its own financial self-interest. Furthermore, this isn't settled just by looking at the motivations of the particular member of the group, say Aidan, whose actions constitute the group's actions of helping Caroline. Even if Aidan is motivated only by a desire to help needy people, it's still an open question whether the group's action is so motivated. Depending on how the case is fleshed out, it may be more or less plausible that the group's action is motivated by the desire to help needy people. For instance, perhaps this is a case in which the group would financially benefit by helping Caroline (e.g. by an expected legacy). Furthermore, we may suppose that whenever its own financial interests conflict with the aim of helping needy people, the church always looks after its own financial interests. In these circumstances, it seems intuitive that the group does not help Caroline on the basis of its desire to help needy people but instead on the basis of its desire for its own financial self-interest. Crucially, whether the group acts from its financial motivation or its desire to help needy people is simply not determined by the motivations of Aidan whose action constitutes the group's action.

Having seen that the first defence of the summative account fails, let's consider a second which appeals to a distinction familiar from the literature on groups between "operative members" and other members. In the kind of organised groups on which we are focusing, different members have different roles and responsibilities. Some – the operative members – have much more responsibility and power than others (compare the executive committee of the charity and its outreach workers). The problems for summative accounts seem to arise from the fact that a group's actions may be constituted by the actions of "footsoldiers" in the organisation where it's implausible that the motivations of the footsoldiers determine the basis of the group's action. For the footsoldiers may have motivations which are "rogue" with respect to the group's motivations. As a result, it may be tempting to think that a more plausible summative account could be provided by focusing on the basis of the actions of the operative members of the group. Indeed, those who argue that groups have mental states and can be morally responsible often take the group's mental states and responsibility to hinge on the operative members (e.g. Tuomela, Gilbert). However, this tempting suggestion faces a serious difficulty: the actions of a group may be constituted by the actions of members of the group other than the operative members. For instance, the operative members of some charity might never actually step onto the streets to feed the homeless even though the charity feeds the homeless. Thus, appeal to the motivations of operative members seems unlikely to provide a general account of group reasons for action. In addition, to the extent that there is more than one operative member with different motivations, the suggested solution faces a problem discussed later on that it may make responsibility arbitrary. For instance, if there are 10 board members all with different motivations, how many must be motivated by some reason for the group to be so motivated (6, 7, 8, 9?). To the extent that any answer seems arbitrary, the suggested solution will make responsibility arbitrary.

Unlike the second response which restricts the relevant actions to those of the operative members, the third response suggests that the set of member actions which constitute a group's action is broader than has been suggested so far. In the key counterexample, we suppose that the group's

action, say feeding a particular homeless person, is constituted by the action of just one member of the group, namely the person on the street handing out the food. This generated the problem given that the footsoldier handing out the food on the street might have motivations which are rogue with respect to the group in general. However, it may be suggested that when properly understood, the set of actions which constitute the group's action of feeding a homeless person on the street is much broader than merely the action of the person handing the food out and includes a large number of activities by other members of the group including fundraising, food ordering, food preparation, and food transport et cetera. It may then be suggested that focusing on this larger group of member actions will avoid the counterexamples. For instance, given that most of the members of the evil charity Bad are motivated by terrorist ends, it may be suggested that focusing on this larger group of actions gets the result that Bad's action of feeding Reg on some particular occasion is motivated by terrorism and not compassion for the needy.

However, this appeal to a broader set of actions faces problems. First, the proposal is unlikely to deal with all cases of group actions. For, in some cases, it's plausible that a group's action is constituted by the action of one member, for a group could endow a member with the power to act on its behalf (e.g. consider the spokesperson or chief legal officer of the company). Second, it's not entirely clear how to distinguish which actions are constitutive elements of the group's action and which are merely background conditions. Third, even setting that aside, a deeper problem is that if one focuses on a larger set of actions, the actions of different members may have many different and varied motivations. This raises the difficult question of what proportion of the set of member actions must be motivated in some way for the group's action to be so motivated. To see the problem, let's consider an example.

Suppose that a university awards an undergraduate student, X, a classification of 2.2 for their philosophy degree. There are large number of member actions which make it the case that the university awards a 2.2 degree to a particular student. For instance, various academic staff mark the student's work over the course of their degree. The marks are then considered by department-level examination committees which make a recommendation about what degree classification to award the student in the light of department and university level policies. These department and university level policies are themselves the result of actions of different committees across the university. The department's decision is then forwarded to the relevant university unit which formalises the department's decision. There are thus a very large number of actions by different members of the university which combine to make it the case that the student is awarded a 2.2. They include actions by the academics who marked the work, actions of members of the department-level examination board, actions of administrators in registry, actions of members of the various high-level committees of the university who determine the marking and classification policies to which the department works.

Plausibly, the relevant actions of the different members may be motivated in a wide variety of ways. The relevant actions of some members may be motivated by financial or professional self-interest; others by a sincere desire to award grades fairly; others by less attractive motivations such as building up the department or securing the financial stability of the university (by ensuring students get high marks). Once we realise the diverse and varied motivations members may have, the account would have to specify what proportion of members need to be motivated in their relevant actions by some aim, for the group to be so motivated. We surely wouldn't want to require that for a group to be motivated in doing  $\alpha$  by  $r$  every relevant action of every member must be motivated by  $r$ . For example, that would have the result that if just one member out of the entire university is motivated in their relevant actions by, say, self-interest then the university's award of the 2.2 cannot be motivated by a concern with fairness even if every other member does their relevant

action out of fairness. On the other hand it is surely too weak to say that a group is motivated in doing  $\alpha$  by  $r$  if at least one relevant action of some member is motivated by  $r$ . For example, that would have the result that if just one low-level member of the university does their relevant action from a sense of fairness, then the university's award of the degree is motivated by fairness even if every other member of the university is motivated instead by, say, financial self-interest. But once we have ruled out these two extreme positions, it's far from clear what is a plausible threshold for the proportion of relevant actions by members which must be motivated by  $r$  for the group's doing  $\alpha$  to be motivated by  $r$ . Indeed, picking a particular threshold, says 75% or 76% seems arbitrary. But surely we don't want it to be arbitrary whether a group is praiseworthy or blameworthy for some action.

We've now seen that serious problems face summative accounts of what it is for a group to act from a reason. On a constitutive approach, the basis of a group's action is determined by the basis of the actions of its individual members which constitute the group action. But this approach runs into difficulties where the group's action is constituted by the actions of members with motivations which are rogue with respect to the group. Appealing to a broader set of actions such as the set of member actions which make it the case that the group acts faces the problem that members may have diverse and varied motivations. In the light of this, it's hard for the summative approach to avoid arbitrary results about group motivation and responsibility. So in the next section I examine inflationary accounts which hold that a group may act on the basis of a reason  $r$  even if no member has  $r$  or acts on  $r$ .

## 4 | INFLATIONARY ACCOUNTS

Inflationary accounts of group phenomena allow that a group can have a property which no member of the group has. For example, inflationary accounts of group belief allow that a group can believe that  $p$  even though no member of the group believes that  $p$ . Similarly, an inflationary account of group action for a reason would allow that a group can act for the reason  $r$  although no member acts from  $r$ , whether that's because no member has  $r$  or, if they do, no member acts from  $r$ . There are two main ways of developing inflationary accounts, joint acceptance approaches inspired by Gilbert and approaches which focus on solving the so-called "discursive dilemma" such as List and Pettit's approach. We'll start with the joint acceptance approach before moving on to List and Pettit's approach.

To illustrate the key features of joint acceptance accounts, consider group belief. On Gilbert's joint acceptance account, "a population,  $P$ , believes that  $p$  if and only if the members of  $P$  are jointly committed to believe as a body that  $p$ " (Gilbert 2004). Gilbert explains that "a joint commitment is created only when each of the parties has, in effect, openly expressed his or her personal readiness to be party to it" (100). Although Gilbert's account speaks of "members", she and other defenders of inflationary accounts allow that a group can believe that  $p$  even if not all members of the group are jointly committed to  $p$ , but rather the operative members are so committed (e.g. Tuomela 2004, Gilbert 2006). A group member could express their readiness to believe that  $p$  as a body in formal ways, such as voting, or by informal ways such as a group discussion which arrives at an agreed position. Importantly, members of a group could jointly commit to believe as a body that  $p$  even though none of them believes that  $p$ . For instance, a jury could agree that its verdict is that the defendant is innocent since the legally admissible evidence doesn't support a guilty verdict, even though each juror individually believes the defendant is guilty taking into account the inadmissible evidence. In the case of a commitment to jointly believe that  $p$  as a body, the parties commit

to behaviour expressive of the belief that *p*, such as asserting that *p* with appropriate confidence in relevant contexts (101-2).

Joint acceptance accounts face many existing objections<sup>8</sup>, but here I offer a new objection: that it has difficulties in providing a plausible account of when groups act on the basis of a certain reason. To illustrate, consider the following example, Bribery, in which a company is considering whether to buy certain safety equipment. Their decision depends on the following conjunctive proposition: the relevant risk is serious and the equipment is affordable and the equipment is effective (or that the equipment is cost-effective for short). Let's suppose that the three-person executive board of the company constitutes the operative members of the company and they need to decide whether to buy the safety equipment. Suppose that none of the three board members believes that the equipment is cost-effective. Nonetheless, each has been separately and privately bribed by the manufacturer of the equipment to argue for its purchase. Accordingly, when the three board members discuss the issue together, each argues in favour of the equipment in such a way that they agree that their joint position is that the equipment is cost-effective, and then implement this decision by purchasing the equipment. (No one mentions the bribe and none of them know that each have been bribed just as they have.) Thus we have a case in which, according to an inflationary account, the group believes that the equipment is cost-effective even though no members of the group believe this. Furthermore, since the group's belief that the equipment is cost-effective leads to the purchase of the equipment, this raises the worry that the inflationary account is committed to holding that the group purchases the equipment for the reason that the equipment is cost-effective. However, this seems to be intuitively the wrong result about the case. Given that each of the three board members is motivated by the bribes they have accepted, it's hard to see this as a case in which the group purchases the equipment on the grounds that it's cost-effective.

Before discussing how this objection appears to affect another key inflationary account of group belief, it's useful to note that the bribery example is effective regardless of one's take on the controversial issue of whether reasons are factive. It's perhaps natural to interpret the example on the assumption that the equipment is not cost-effective so that the group's belief that the equipment is cost-effective is false. But, on a factive view of reasons, false propositions cannot be reasons and, thus, cannot be an agent's reason for acting. However, the bribery example still poses a challenge even on a factive view of reasons. For we can simply reframe the relevant question not as whether the group acted for the reason that the equipment is cost-effective but, instead, whether the group acted for some related true proposition, e.g. that they believe the equipment to be cost-effective. But, given the role of bribery in the decision, it is counterintuitive to suppose that the group acted for the reason that they believed the equipment to be cost-effective. So, from now on, I will set aside the issue of factivity, to focus on how the same objection applies to List and Pettit's inflationary account.

List and Pettit's (2011) core argument for their inflationary view arises from so-called discursive dilemmas in which ascribing mental states to a group on a simple majority basis would have the result that the group has an inconsistent set of beliefs. They christen such cases "dilemmas" because they think they force us to choose between a number of attractive desiderata on any account of group belief. While they canvas a number of solutions to discursive dilemmas, one of the potential solutions they suggest is premise-aggregation according to which if a group believes each of a set of premises which entails the conclusion, the group believes the conclusion. Here I focus on the premise-aggregation approach.

To do so, let us reimagine our bribery example so that according to the group's constitution, the group's position on the premises is determined by majority voting but its position on the

conclusion is determined by premise aggregation (if the group believes each of the propositions that the risk is serious (s), the equipment is effective (e), and the equipment is affordable (a), then it believes their conjunction (s and e and a)). As before, no member of the executive board individually believes that the equipment is affordable and effective against a serious risk. However, the bribe leads them to vote as follows:

	s	e	a	(s and e and a)
M1	Yes	Yes	No	No
M2	Yes	No	Yes	No
M3	No	Yes	Yes	No

The way in which they vote differs from their own individual honest opinions on the matter. For example, M1's honest view is that while the equipment is effective, it is not affordable and nor is the risk serious. However, given the bribe they've taken, they vote that the risk is serious, that the equipment is effective but not affordable. Similarly, while M2's honest view is that it is a serious risk, but the equipment is neither effective nor affordable, given their bribe they vote in favour of the risk being serious, the equipment being affordable but not effective. Similarly, M3 votes in a way which doesn't reflect their honest view. Given the way in which each member votes, there is a majority of votes for each of the premises so that the group believes each of the premises. Further by the premise aggregation method, the group believes the conjunction of the premises even though none of the members of the group individually believes the conjunction. (Notice that List and Pettit allow that the group's beliefs can be determined by voting, e.g. 2011:44.)

We see then that just like the joint acceptance account, the premise aggregation view has the result that the group believes that the equipment is cost-effective even though no members believe that the equipment is cost-effective. Furthermore, we may suppose that the group has set up a process whereby group belief arrived at via premise aggregation automatically feeds into company actions. For instance in this case, the group's belief that the equipment is affordable and effective against serious risk automatically generates a purchasing instruction for the purchasing department. This raises the concern that, like the joint acceptance account, the premise-aggregation view is committed to the claim that the group buys the equipment on the basis of its belief that it's affordable and effective against serious risk even though, intuitively, that's the wrong result in this case. In the next section, I consider how defenders of the inflationary view may attempt to deal with the objection.

## 5 | PROGRAMMING, CAUSATION AND COUNTERFACTUALS

Defenders of inflationary views may attempt to resist the objection in a number of ways. Of particular relevance might seem to be List and Pettit's discussion of how a group has the control to choose between options. So it's with their discussion that I start, before moving on to consider appeal to causal and counterfactual conditions.

As part of their defence of the idea that groups can be responsible for their actions, List and Pettit argue that groups have the control required for choosing between options (2011: 155). In an earlier article, Pettit (2007) gives a version of this control condition as involving the idea that groups have the control required for choosing between options on the basis of their normative judgements. In

explaining how groups meet the control condition, List and Pettit discuss the worry that whatever a group agent does is done by an individual agent and that excludes the group's having control over which options it performs. List and Pettit's solution appeals to the notion that a higher-level event can "programme for" some effect even though the programme is implemented by lower-level events. For example, we might truly say that the water's boiling causes the flask which contains it to collapse even though the activities of various water molecules are clearly also causally relevant to the collapse. On their account, the water's boiling is causally relevant to the event of the collapse in virtue of the fact that given that the water is boiling "it is more or less inevitable that there will be some constituent molecule, maybe this maybe that, that has a position and momentum sufficient to induce a crack in the surface of the flask" (Pettit 2007: 191; List & Pettit 2011:162-163). Similarly they suggest that a group can have control over its actions by programming for certain actions to occur:

The group may control in a reason-sensitive way for the performance of a certain action by some members, maybe these maybe those. It will do this, by maintaining a constitution for the formation and enactment of its attitudes, arranging things so that some individual or individuals are identified as the agents to perform a required task, and other individuals are identified as agents to ensure that should the performers fail, there will be others to take their places as backups (Pettit 2007: 192; see also List & Pettit 2011: 163).<sup>9</sup>

However, even if we accept List and Pettit's account of how groups may control their actions, their programming explanation fails to show that in our problematic Bribery case, the group does not purchase the safety equipment on the basis of the consideration that the risk is serious and the equipment is affordable and efficacious. For in Bribery, the company does have a constitution for the formation and enactment of its attitudes, in particular its constitution is that the executive board votes on relevant premises which determine its overall view via premise-aggregation. By this process, the group counts as believing that the equipment is affordable and efficacious against a serious risk, and this belief is then inputted into the group's purchasing activities in such a way that some individual is identified as the individual to purchase the equipment. So by the programming account, it seems to count as purchasing the equipment on the basis of its belief that it is affordable and efficacious against a serious risk. But, this seems the wrong result given that the group's view arises from the votes of individual members none of whom believe that the equipment is cost-effective and all of whom vote as they do on the basis of the bribe.

Indeed, it seems that the programming approach would treat equally Bribery and another company, Honest, in which everything is the same except that each board member votes honestly in line with their own opinions. In both Bribery and Honest, the group may have set up its constitution so that the group's view on the premises is determined by majority voting by individual board members and the group's view on the conclusion is determined by premise aggregation. Furthermore, we may suppose that the group's view so determined is automatically inputted into the group's procedures for implementing decisions. Thus, the programming solution would judge that in both Bribery and Honest, the group purchases the equipment on the basis of its belief that it is cost-effective. But this is a problematic result given that, in Bribery but not Honest, the executive members of the board individually act on the basis of financial self-interest and none of them believe that the equipment is cost-effective.

Of course, there is a difference between Bribery and Honest in terms of the reasons why individual board members vote as they do: in Bribery, they vote on the basis of financial self-interest;

in Honest, we may suppose that they vote on the basis of their own assessment of the evidence. However, these individual differences in motivation are no part of the group-level constitution for the formation and enactment of attitudes which are at the heart of List and Pettit's programming notion. The group-level constitution for the formation of its attitudes determines the group's view on whether the equipment is cost-effective as a function of member votes on the relevant premises, but allows no role for the reasons why individual members vote as they do. Furthermore, it would seem problematic for List and Pettit to appeal to these differences at the individual level, for appeal to programming is supposed to be an account of how we can see the group, rather than its individual members, as controlling its actions. While individual member's reasons explain why they vote as they do which in turn determines the group's view, the programming approach is not concerned with the causal precursors of the group's view but with whether the group's view itself causally explains the group's action. Compare List and Pettit's account of high-level causation on which group control is supposed to be modelled. If we are attempting to explain how the water's boiling causes the flask to collapse it seems irrelevant to appeal to the cause of the water's boiling, say the fact that the gas was turned on underneath the flask. Turning on the gas may explain why the water is boiling, but that doesn't help show that the water's boiling itself has further causal effects.

It seems, then, that appeal to programming fails to provide an adequate account of group action for a reason. However, despite this failure, it may be argued that inflationary accounts have other ways of showing that, in Bribery, the group does not act for the reason that the equipment is affordable and efficacious against a serious risk. In particular, a defender of an inflationary view might hope that they can secure this result by appeal to some of the standard tests for whether an individual acts on a certain consideration, e.g. causal and counterfactual tests. For example, it may be argued that Kant's grocer gives correct change out of self-interest rather than because it's the right thing to do on the grounds that 1) the cause of his giving correct change is his belief that it's in his financial self-interest; or 2) the following counterfactual is true: if it hadn't been in his financial self-interest to give correct change, he wouldn't have done so. Of course, there are well-known problems with causal and counterfactual tests for acting on a reason. However, even setting those aside, we can see that appeal to these tests doesn't help show that, in Bribery, the company does not purchase the equipment on the basis of its belief that the equipment is affordable and effective against serious risk.

To apply the causal test to a group, we need an account of when a group's action is caused by one reason rather than another. There is little discussion of this in the literature except List and Pettit's account which Pettit hoped would show that a group can have reason-sensitive control of its action. But we have already seen that it fails to help. For the company's belief that the equipment is affordable and effective against serious risk programs for the purchase of the equipment, where programming is List and Pettit's way of showing how the group's belief is causally efficacious with respect to the action.

Nor does appeal to counterfactuals help. On the proposed counterfactual test, a subject acts from a reason only if the following counterfactual is true: if the subject hadn't had that reason, then she wouldn't have acted. Applied to the bribery example, the company acts on the basis of their belief that the equipment is cost-effective only if the following counterfactual is true: if the company hadn't believed that the equipment is cost-effective, then it wouldn't have purchased it. But, it seems that even in the bribery example, the relevant counterfactual may be true. The company may well be so set up that the group's beliefs as constituted by the outcome of votes amongst the executive board are translated into action so that it's true that if the group did not believe that the equipment is affordable and effective against serious risk, then it wouldn't have

purchased the equipment. For instance, if the board members had voted honestly in such a way that the outcome of premise aggregation was that the group did not believe the equipment to be affordable and effective against serious risk, then the group would not have purchased the equipment.

The defender of an inflationary view may instead appeal to the notion of the ultimate or fundamental reason for the group's adoption of the belief that the equipment is affordable and effective against serious risk. To the extent that the process leading to the adoption of this belief is the result of financial self-interest, it may be suggested that the group's purchasing of the equipment is ultimately due to self-interest. In defending this suggestion, the defender of an inflationary view may compare individuals and groups. Even if some reason is the proximate cause of an individual's action, plausibly the ultimate cause may be something else entirely. For instance, even if the proximate cause of an industrialist's decision to make his own son head of the business is his sincere belief that his son is the best candidate, the ultimate reason for this belief and the consequent action may be his interest in keeping family control of the business. Likewise, the defender of an inflationary view might hope to argue that they can get the intuitively correct result that the group's purchase of the equipment is ultimately the result of the group's financial interest by arguing that it is the group's financial interest which is the ultimate cause of the group's belief that the equipment is affordable and effective against a serious risk.

However, there are problems with this appeal to the ultimate group motivation of the group's action of purchasing the equipment. On the premise aggregation view, the group's reason for holding the belief that the equipment is cost-effective is not determined as a function of the individual member's reasons for voting as they did. Instead, the group's view on the conclusion that it is cost-effective is determined by the group's view on the premises. Thus, the proximate cause of the group's belief that the equipment is cost-effective is plausibly the group's belief in the relevant premises: that the equipment is affordable; that the equipment is effective; and that the risk is serious. Of course, we can push back the question further and ask what is the group's reason for adopting a belief in each of the premises. So far, our story doesn't say anything about that. To see that it's important to carefully distinguish between the group and individual level. Of course, it is part of the story that the financial self-interest of each individual board member leads them to endorse a premise they don't sincerely believe as an individual. But, on an inflationary account, that tells us nothing about the group's reason for adopting each premise. For on an inflationary account, it is neither necessary nor sufficient for a group to believe some proposition that its members believe that proposition. Furthermore, it's quite possible that the financial self-interest of individual board members comes apart from the group's financial self-interest. It may be that by purchasing the equipment they make themselves individually millionaires but ruin the company long-term. Thus, there is no inference from the fact that individual board members vote from individual self-interest to the claim that it's in the group's financial self-interest to buy the equipment, or that the group buys the equipment because that's what ultimately serves its self-interest.

In conclusion, it seems that inflationary accounts have difficulty getting the intuitively correct result about a group's motivation for its actions. In our bribery example, they allow that a group can be motivated by some reason, such as the belief that the equipment is cost-effective, even though no member of the group believes that and the relevant actions of operative members are motivated by individual self-interest. Furthermore, we saw that we couldn't overcome this problem either by appeal to standard causal and counterfactual tests, or by appeal to the notion of the ultimate reason for the group's belief.

## 6 | GROUP AGENCY AND RESPONSIBILITY RECONSIDERED

So far, we've seen that there are problems facing the main options for providing an account of group action for a reason. On the summative approach, a group's reason for acting is revealed in the reasons why members undertake relevant actions. For example, it may be suggested that the basis of a group's action is determined by the basis of the actions of its individual members which constitute the group action. But, as we saw, this approach runs into difficulties when members have motivations which are rogue with respect to the group. Appealing to a broader set of actions such as the set of member actions which make it the case that the group acts faces the problem that members may have diverse and varied motivations. In the light of this, we saw that it's hard for the summative approach to avoid arbitrary results about group motivation and responsibility. By contrast with summative approaches, inflationary accounts allow that a group can act for a reason even if none of its members act for that reason. But by allowing the group's motivation to float free of the motivation of its members, inflationary accounts give counterintuitive results about key cases. For instance, it allows that in our bribery example, the group acted for the reason that the equipment was affordable and effective against serious risk even though the group's action results from its operative members taking bribes.

The problems we've seen with developing an account of what it is for a group to act for a reason raise concerns about the idea that groups are responsible agents given that the very notions of agency and responsibility are often explained by appeal to the notion of acting for a reason. On what we can call the "reasons view", intentional action is explained as action for a reason; responsibility requires the ability to act for reasons; and whether one is praiseworthy or blameworthy for an action depends on the reason for which it was done. If we cannot find a satisfactory approach to determining group motivation then, on the reasons view, this undermines our practice of blaming and praising groups. For to the extent that we cannot determine group motivation, we cannot determine whether a group is praiseworthy or blameworthy. Thus, our judgements that groups are praiseworthy or blameworthy are unjustified. If we want to defend our practice of holding groups morally responsible, we have two broad options: to provide a better account of group motivation or to argue that groups are morally responsible even if they are not agents who act for reasons. I sketch out each of these options below.

First, we might try to provide a better account of group motivation. In the case of organised groups, such as companies and governments, we might wonder whether we could do so by looking at their institutional structure: the way in which their systems and processes are set up. If, say, a widget manufacturer is to avoid polluting the local river system this will involve designing and implementing the right systems and processes, such as designing the manufacturing system to avoid producing pollutants and/or to appropriately deal with any pollutants produced. Further, since systems – including human operatives – don't always function as intended, the manufacturer will need a variety of procedures to try and increase the likelihood of the system as a whole functioning properly (e.g. appropriate policies to train and incentivise employees, and maintain plant) together with processes to detect and mitigate failure (spot checks, warning systems et cetera). We might wonder if we could read a group's motives in its institutional structure. To the extent that our widget manufacturer has implemented extensive and costly procedures for attempting to avoid environmental pollution, it may seem tempting to think that this shows that it's motivated by the aim of avoiding environmental pollution. Indeed, we might think that institutional structure of this kind can play one of the key roles which was supposed to be played by the notion of acting for a reason in the moral worth literature. Part of the reason for insisting that an

action is praiseworthy only if done from the right motive is to deal with the problem of accidental right-doing, such as Kant's grocer who gives his customers the correct change only out of self-interest. To the extent that our widget manufacturer avoids polluting the environment because of an extensive and expensive set of systems and procedures, then it seems no accident that the widget manufacturer avoids polluting the environment.

While examining institutional structures may look promising, it will have the effect of divorcing group motivation completely from the motivation of individual members of the group. For we can imagine two widget companies who have precisely the same institutional structures for avoiding environmental pollution where all the operative members of one company are motivated by genuine concern for the environment whereas all the operative members of the other company are only driven by desire to avoid regulatory sanctions. To the extent that we would be inclined to treat the two companies differently, then this might seem to undermine the idea that we can identify a group's motivation merely by looking at institutional structures and without looking at the motivation of individual members. Thus, finding a satisfactory account of group motivation may not be an easy task.

If we cannot provide a satisfactory account of group motivation we might instead attempt to argue that groups can be properly held morally responsible even if they are not agents who act for reasons. In particular, a number of those working in the group literature have suggested that loose collectives of individuals can be held morally responsible even though they lack the kind of organisation that makes it plausible that they together constitute an agent which can act for reasons. Some do so by arguing that agency is not required for moral responsibility (e.g. Wringer). Others accept that agency is required for moral responsibility but attempt to show how loose collectives which do not themselves constitute agents can nonetheless have moral obligations (e.g. Björnsson 2014, Pinkert 2014). For instance, consider the collection of all humans currently existing. Plausibly, this collection of agents is not organised in such a way that it is an agent with mental states. Instead, the only agents are the individual human beings who are members of this loose collective. Nonetheless, that's compatible with the suggestion that all humans currently existing have joint or shared obligations, say to protect the earth from climate change, and that they may fail to do so in a blameworthy way. We might hope to apply these accounts of how many individuals can together have joint obligations to show that the members of organised groups can have joint obligations even if organised groups fail to be agents who can act for reasons. To the extent that we can do that, then we might be able to defend a practice of praising or blaming members of organised groups even despite the problems in providing an account of group motivation.

It seems, then, that despite the results of the paper, there may well be ways to defend our practice of treating organised groups as morally responsible, whether by providing a better account of group motivation or by arguing that organised groups are morally responsible even if they don't constitute agents who can act for reasons. Which of these options for defending group moral responsibility is more plausible will have to wait for another occasion.

## 7 | CONCLUSION

In this paper, I've examined a key issue for group agency and responsibility, namely what it is for a group to act for one reason rather than another. The notion of acting for one reason rather than another is central to standard accounts of individual agency and responsibility, and also determines whether an individual is blameworthy or praiseworthy for an action. In addition, leading proponents of the idea that groups are morally responsible explicitly appeal to the idea that groups

can act for reasons. Despite this, there has been little explicit discussion of what it is for a group to act for one reason rather than another. In this paper, I've examined the prospects for the two main potential accounts of what it is for a group to act for one reason rather than another – summative and inflationary accounts. The problems facing these accounts raise a serious challenge to our practices of holding organised groups to be morally responsible agents, who can be praised or blamed for their actions. At the end, I outlined several ways of responding to this problem including appealing to group organisational structure, revising our account of agency, or appealing to accounts of the shared obligations of disorganised collections of individuals.

## NOTES

- <sup>1</sup> E.g. Pettit & Schweikard 2006:33; Lawford-Smith 2015:232; Collins 2019:949.
- <sup>2</sup> In response, defenders of group responsibility point out that they can also acknowledge individual responsibility and allow that these can come apart so that, e.g., a group's being responsible is compatible with different members being individually responsible to different degrees. See inter alia Gilbert 2006, Isaacs 2011, Hess 2018.
- <sup>3</sup> For example, it has been suggested that one is praiseworthy for an action only if it is done from the motive of duty; or for the reason that makes it right (Arpaly, Markovitz) or from concern for the right together with knowledge of the right (Sliwa).
- <sup>4</sup> E.g. Rosen 2003, Nettleman 2007, Peels 2014, Alvarez & Littlejohn 2016, Weatherson 2019. The distinction between acting in ignorance and from ignorance harks back to Aristotle. For example, suppose that I feed my arch-enemy arsenic-laced coffee. Even if I'm ignorant that it contains arsenic, that doesn't excuse me if I didn't act from my ignorance (e.g. if knowing that it contains arsenic would have made no difference to my action).
- <sup>5</sup> A further worry about the sufficiency direction of SA is that it may not even make sense to suppose that the group acts for the reason which motivates the relevant individual. For instance, the relevant member of a charity might feed a refugee from the desire to impress his girlfriend. But groups don't have girlfriends. Thanks to Glock for raising this kind of worry.
- <sup>6</sup> Mathiesen suggests that various actions by members of a group constitute the group's acting on the basis of its aims and beliefs to the extent that those members of the group take on, and act from, the group's goals and beliefs (249-550).
- <sup>7</sup> For this to be a summative account, condition 1 must be given a summative reading. If not, then the account would be hybrid, involving a summative and a non-summative condition. But it would still face the problem mentioned in the main text.
- <sup>8</sup> Some argue that the conditions in the joint acceptance account are not sufficient for genuine belief, but only mere acceptance (e.g. Wray 2001). Lackey (2016) argues that the joint acceptance account cannot provide an adequate account of group lies or justified belief.
- <sup>9</sup> In their joint monograph, List and Pettit say that "as a group agent, we are only committed to the weaker view, that exercising control involves exercising causal control" but not the additional requirement endorsed earlier by Pettit that "the agent's normative judgements play an appropriate causal role in generating or controlling action" (2011: 225, n.110)

## REFERENCES

- Alvarez, M. & Littlejohn, C (2016). When ignorance is no excuse. In P. Robichaud and JW. Wieland (Eds.), *Responsibility – The Epistemic Condition*. Oxford: Oxford University Press.
- Arpaly, N. (2003). *Unprincipled Virtue*. Oxford: Oxford University Press.
- Björnsson, G. (2014). Essentially shared obligations. *Midwest Studies in Philosophy*, 38: 103–120. <https://doi.org/10.1111/misp.12019>
- Collins, S. (2019). Collective responsibility gaps. *Journal of Business Ethics*, 154, 943-954. <https://doi.org/10.1007/s10551-018-3890-6>
- Fischer, JM. & Ravizza, M (1998). *Responsibility and Control*. Cambridge: Cambridge University Press.
- French, P. (1995). *Corporate Ethics*. Orlando, Florida: Harcourt Brace.
- Gilbert, M. (1989). *On Social Facts*. Princeton, NJ: Princeton University Press.
- Gilbert, M. (2004). Collective epistemology. *Episteme*, 1, 95-107. <https://doi.org/10.3366/epi.2004.1.2.95>

- Gilbert, M. (2006). Who's to blame? *Midwest Studies in Philosophy*, 30: 94–114. <https://doi.org/10.1111/j.1475-4975.2006.00130.x>
- Haney, MR. (2004). Corporate loss of innocence for the sake of accountability. *Journal of Social Philosophy*, 35(3), 391-412. <https://doi.org/10.1111/j.1467-9833.2004.00240.x>
- Hess, K. (2018). Who is responsible? *Midwest Studies in Philosophy*, 1–23. <https://doi.org/10.1111/misp.12087>.
- Isaacs, T. (2011). *Moral Responsibility in Collective Contexts*. Oxford: Oxford University Press.
- Lackey, J. (2016). What is justified group belief? *Philosophical Review*, 125(3), 341-396. doi: [jphil198683183](https://doi.org/10.1111/j.1467-9833.2016.00130.x)
- Lackey, J. (2020). *Epistemology of Groups*. Oxford: Oxford University Press.
- Lawford-Smith, H. (2015). What “we”? *Journal of Social Ontology*, 1(2), 225-249. <https://doi.org/10.1515/jso-2015-0008>
- List, D. (2005). Group knowledge and group rationality. *Episteme*, 2, 25-38. <https://doi.org/10.3366/epi.2005.2.1.25>
- List, D. and Pettit, P. (2011). *Group Agency*. Oxford: Oxford University Press.
- McKenna, M. (2006). Collective responsibility and an agent meaning theory. *Midwest Studies in Philosophy*, 30(1), 16-34. <https://doi.org/10.1111/j.1475-4975.2006.00126.x>
- Mathiesen, K. (2006). We are all in this together. *Midwest Studies in Philosophy* 30: 240–55. <https://doi.org/10.1111/j.1475-4975.2006.00137.x>
- Markovits, J. (2010). Acting for the right reasons. *Philosophical Review*, 119, 201-242. <https://doi.org/10.1215/00318108-2009-037>
- Nettleman, N. (2007). *Blameworthy Belief*. Dordrecht: Springer.
- Peels, R. (2014). What kind of ignorance excuses? *Philosophical Quarterly*, 64: 478–496. <https://doi.org/10.1093/pq/pqu013>
- Pettit, P. (2007). Responsibility incorporated. *Ethics*, 117, 171-201. <https://doi.org/10.1086/510695>
- Pettit, P and Schweikard, D. (2016). Joint actions and group agents. *Philosophy of the Social Sciences*, 36, 18-39. <https://doi.org/10.1177/0048393105284169>
- Pinkert, F. (2014). What we together can do. *Midwest Studies in Philosophy*, 38(1), 187-202. <https://doi.org/10.1111/misp.12023>
- Rosen, M. (2003). Culpability and ignorance. *Proceedings of the Aristotelian Society*, 103, 61-84.
- Strawson, P. (2003). Freedom and resentment. In G. Watson (Ed.) *Free Will*, 2nd edition. Oxford: Oxford University Press: 72–93.
- Tollefsen, D. (2003). Participant reactive attitudes and collective responsibility. *Philosophical Explorations*, 6(3), 218-234.
- Tollefsen, D. (2015). *Groups As Agents*. Polity Press.
- Tuomela, R. (2004). Group knowledge analysed. *Episteme*, 1: 109–27. <https://doi.org/10.3366/epi2004.1.2.109>.
- Wallace, RJ. (1994). *Responsibility and the moral sentiments*. Massachusetts: Harvard University Press.
- Weatherston, B. (2019). *Normative Externalism*. Oxford: Oxford University Press.
- Wray, B. (2001). Collective belief and acceptance. *Synthese*, 129, 319-333. <https://doi.org/10.1023/A:1013148515033>.
- Wringe, B. (2020). Global obligations, collective capacities, and “ought implies can”. *Philosophical Studies*, 177(6), 1523-1538.

**How to cite this article:** Brown J. Group motivation. *Noûs*. 2021;1–17.  
<https://doi.org/10.1111/nous.12366>