Journal of Ethics and Social Philosophy
VOL. 25, NO. 3 · SEPTEMBER 2023
https://doi.org/10.26556/jesp.v25i3.2590
© 2023 Author

ATTRIBUTIONIST GROUP AGENT RESPONSIBILITY

Adam Piovarchy

MUCH WORK has been carried out showing that group agents exist as distinct entities not merely constituted by aggregating the set of individuals who make them up. Many philosophers believe we can also talk meaningfully about their possessing duties and carrying out actions. Paradigmatic examples of genuinely group agents include nation-states and corporations. But once we have established that group agents exist, that they can have duties, and that they can perform *bona fide* actions, an important question arises: Can group agents be *morally responsible* for violating said duties through their actions (or omissions)? Here, I am concerned with moral responsibility in a backward-looking sense, where an agent is responsible in virtue of what they have done independent of forward-looking considerations such as whether holding them responsible will produce good effects. To say that an agent is morally responsible for a wrong (right) action or omission is to say that they are an appropriate target of blame (praise) for that action or omission. As is standard, I will focus on blame for wrongdoing, given that the risks of incorrectly blaming are typically much higher than the risks of incorrectly praising. For now, as a first pass to help home in on our target phenomenon and to be ecumenical with respect to existing theories of blame, I will understand blame as a negative reactive attitude, which is generally unpleasant to be targeted with, and which communicates disapproval of the agent's conduct.

Intuitively, it seems like group agents can be blameworthy. We blame Volkswagen for its widespread intentional violation of emissions laws. We blame governments for failing to pass laws that reduce the damage caused by climate change. We call on such group agents to exhibit certain kinds of responses, such as apologizing and compensating victims, and we blame them even further if they do not. But demonstrating that group agents can be morally responsible requires that we spell out what features group agents must possess in order to be blameworthy or praiseworthy. This is standardly provided by taking group agents to be constituted by—and responsible in virtue of—certain well-ordered decision-making structures that are responsive to reasons. Since group

agents are the kinds of things that can appreciate moral reasons, and since they can control themselves in response to those reasons, they exhibit the kind of control that is emblematic of moral agency and which therefore makes them an appropriate target of blame and praise.

Though this standard line of argument gets many things right, parties to these debates may not be locating the blameworthy-making features of group agents in the right place. This becomes particularly salient when we notice that some group agents seem to lack the capacity to respond to certain kinds of considerations—and so cannot act on those considerations—and, rather than being excused, seem to be blameworthy precisely in virtue of this fact. The existence of such agents calls for a revised understanding of what it is that makes group agents responsible.

This paper proceeds as follows. First, I will outline a standard way of accounting for group agency and responsibility. Though particular accounts of group agency differ on the details, these will not be relevant for my argument. I will then present two objections. One is from Thompson, arguing that group agents cannot be responsible because they cannot take people as the objects of their attitudes.[1] The other is a new objection from myself: it is that group agents who consistently do the wrong thing due to their decision-making structure seem intuitively blameworthy, but current ways of understanding responsibility are committed to excusing such agents. I will then argue that avoiding these objections requires us to adopt an *attributionist* theory of group responsibility. On this account, group agents are responsible when their actions are attributable to them in such a way that reflects their evaluative judgments. Group agents are blameworthy when this evaluative judgment is objectionable, and importantly, evaluative judgments can be objectionable even if the agent lacks the ability to avoid wrongdoing or recognize some moral considerations.

1. GROUP AGENCY AND GROUP AGENT RESPONSIBILITY

The standard story for how there can be group agents takes a functionalist approach. To be an agent, one must be capable of having representational states and motivational states, and be capable of acting on the basis of these states. One must also meet some minimum standards of rationality, such as having a certain degree of consistency among one's beliefs and motivations. Group agents can possess these states and meet these standards by having certain well-ordered procedures or decision mechanisms in place, such as

1 Thompson, "The Moral Agency of Group Agents."

majority-rule voting.[2] The results of these procedures determine what the group agent "believes" or "decides." That there really is an agent existing over and above the decisions of individual group members is strongly evidenced by the fact that group agents can believe and decide things that none of the individual group members believe or decide.

If we grant that group agents exist and can perform actions, the next question regards which kinds of actions they are morally responsible for performing. In the most thorough treatment of this question to date, List and Pettit propose that group agents are morally responsible when the following criteria are met:[3]

> *Normative Significance*: The agent faces a normatively significant choice, involving the possibility of doing something good or bad, right or wrong.
>
> *Judgmental Capacity*: The agent has the understanding and access to evidence required for making normative judgments about the options.
>
> *Relevant Control*: The agent has the control required for choosing between the options.[4]

To give an example of how this works in practice, we can easily see that Volkswagen qualifies as responsible for its widespread violation of emissions laws. The decision to violate emissions laws is normatively significant because increasing emissions imposes nontrivial costs on others, which Volkswagen does not have a *prima facie* right to impose. Volkswagen has an understanding of the costs of increasing emissions and breaking the law, and access to evidence required for making normative judgments about its options. It also possessed control over its actions—it could have freely chosen to comply with the law, and it freely chose to violate the law, without any compulsion or coercion. Since it meets the above criteria, Volkswagen seems blameworthy for violating emissions laws.

2 List and Pettit, *Group Agency*; French, *Collective and Corporate Responsibility*.

3 List and Pettit, *Group Agency,* 155.

4 List and Pettit, *Group Agency,* 155. Similar sentiments are endorsed by Gilbert ("Collective Guilt"). Ways to adopt something like the standard story of agency with a different conception of responsibility can be found in Baddorf, "Phenomenal Conciousness" and Tanguay-Renaud, "To Fill or Not to Fill." Some philosophers widen the scope of group agency to include some kinds of collectives, such as Gilbert and Pilchman, "Belief, Acceptance, and What Happens in Groups"; Tuomela, Hakli, and Mäkelä, *Social Ontology in the Making*; and Tuomela, *Social Ontology*. An anonymous reviewer asks whether this paper's argument is relevant for such accounts. While I do not have the space to canvass the similarities and differences of these approaches to List and Pettit's, so long as these accounts allow that group agents can experience local structural deficits (explained below) in ways that do not undermine their agency altogether, and remain capable of expressing objectionable evaluative judgments, then the argument will apply to these accounts too.

2. PROBLEMS FOR THE STANDARD STORY

There are a number of ways to object to accounts of group agency, and this has generated a number of defenses in turn. For considerations of space, my focus will be limited to a select few; however, these particular objections are important. While other objections to group agency are typically responded to by finessing our account of group agency or identifying properties possessed by both individual agents and group agents, handling the objections raised in this paper instead requires that we reconsider the nature of moral responsibility and what it is that ultimately makes group agents blameworthy for wrongdoing.

2.1. Group Agents and Persons as Intentional Objects

The first objection is that group agents simply do not possess certain kinds of properties putatively necessary for moral responsibility. Thompson argues that group agents are unable to have certain kinds of emotions that Strawson takes to be essential to moral responsibility, namely, reactive attitudes such as guilt and resentment.[5] Thompson takes guilt and resentment to view the same wrong from different perspectives, with resentment being the second-personal perspective of the agent wronged and guilt being the first-person perspective of the wrongdoer. He believes that resentment *qua* blame has the function of bringing a perpetrator's moral understanding of their actions into alignment with the blamer's and the rest of the moral community by generating guilt and remorse.

Group agents are capable of functional equivalents of the epistemic and motivational components of guilt and remorse, in that group agents can have beliefs like "I have culpably violated a norm," and they can engage in apologies.[6] But Thompson argues that this is not enough. When we blame, we do not simply want the functional equivalents of guilt and remorse; it will not be sufficient for our target to go through the motions, acting *as if* they feel guilt and remorse. It is not enough that perpetrators simply believe they are blameworthy and desire to make amends. Rather, we want them to *care*, and this is something which group agents are unable to do.

Thompson's key argument is that moral emotions require certain intentional objects, and the objects of some reactive attitudes are *people*. Guilt that is not directed at one's self is not truly guilt, for instance.[7] The problem here is not

5 Thompson, "Moral Agency of Group Agents"; Strawson, "Freedom and Resentment."

6 Bjornsson and Hess, "Corporate Crocodile Tears?"

7 For alternative accounts of collective guilt, see Gilbert, *Joint Commitment*; cf. Ziv, "Collective Guilt Feeling Revisited." Hindriks also develops a noteworthy account of the moral emotions of group agents ("Collective Agency"), though for objections, see de Haan, "Collective Moral Agency and Self-Induced Moral Incapacity."

simply that group agents are not moral agents because moral agency requires phenomenal experiences, and group agents are incapable of these;[8] such an argument would beg the question against functionalist approaches to agency. Rather, the problem lies in group agents being able to only take propositions as their intentional objects. If group agents cannot take people as their intentional objects, they cannot care about people, and so they are not genuinely capable of experiencing guilt. Because they cannot genuinely care, they are psychologically abnormal, outside the bounds of our moral community, and therefore we can only respond with Strawson's "objective attitude" toward them.[9] They are a thing to be managed, rather than participants in our moral practices. Group agents seem more analogous to psychopaths, whom various philosophers take to be excused.[10]

2.2. Group Agents and Local Deficits

A second challenge to the responsibility of group agents concerns the existence of group agents who are intuitively blameworthy but whom the standard story will excuse. Though a lot of attention has been given to whether group agents possess the general capacity to deliberate, consider reasons, and act on the basis of those reasons, a problem which has not yet been considered concerns the possibility of what I shall call *local structural deficits* in decision-making capacity. Even if we grant that group agents can be morally responsible in general, there may be group agents who, due to the structure of their decision-making process, lack the ability to make certain kinds of decisions or act on certain kinds of reasons, and so cannot be responsible for failing to make certain kinds of decisions.[11]

Suppose, for instance, that when forming a company, group members design a decision-making structure that specifically precludes the group agent

8 Hindriks, "How Autonomous Are Collective Agents?"; Tollefsen, "Participant Reactive Attitudes."

9 Strawson, "Freedom and Resentment." Two anonymous reviewers helpfully suggest that we might avoid this objection by formulating caring in another way, such as valuing something and acting appropriately in light of that valuing. For instance, a corporation can value its employees and buy safety equipment for them to keep them safe, even if this reduces the company's long-term profits. Alternatively, we might think of caring as patterns of attention that cluster around certain kinds of issues in their deliberation.

10 Nelkin, "Psychopaths, Incorrigible Racists, and the Faces of Responsibility"; Watson, "The Trouble with Psychopaths." This summary is quite quick, and some aspects, as written, call for more clarification. For reasons that will become clear, I return to this argument and elaborate on the details below.

11 List and Pettit hint at such a possibility but think it is unlikely to happen in practice (*Group Agency*, 159).

from considering reasons that favor the environment or performing actions that would sacrifice profits for the environment. If we like, perhaps the group members' constitution has a clause stipulating that such reasons are simply inadmissible in the decision-making procedure, or that the group agent will self-destruct if such an action is performed, or that the group must now fulfill some kind of chain of steps that continues in an infinite recursion until the group members remove that reason from consideration. While we can certainly argue that the group members are blameworthy for this in virtue of their control over their actions and their awareness of those actions' consequences in creating the group agent, it is still tempting to think that the resulting group agent is itself also blameworthy.

Such structures would prevent group agents from meeting List and Pettit's criteria for many actions they perform. This could occur in two ways. The first is that group agents might, in virtue of certain local structural deficits, lack control over some kinds of actions and so cannot make certain choices. This would limit what options the group agent has available, and thus, according to List and Pettit's third criterion, the agent cannot be responsible for failing to avoid wrongdoing as no alternative was available. Alternatively, group agents with local structural deficits could fall afoul of List and Pettit's second criterion, in that its restricted options prevent the agent from forming certain kinds of normative judgments about other, unavailable options.[12] Even if one does not endorse List and Pettit's particular account, group agents would also fail to be responsible according to *control* accounts of moral responsibility, which hold that agents cannot be blameworthy for wrongdoing when they lack a certain kind of control over their actions, typically the capacity to avoid wrongdoing.[13]

The problem is that there appear to be group agents with local structural deficits that prevent them from making certain kinds of judgments or performing certain kinds of actions, and who yet, to many people, still seem blameworthy. For example, when a villager in the Amazon rainforest sees ACME Co. destroying the trees around them, polluting groundwater, and bribing officials to get away with this, it seems very appropriate for said villagers (and us) to blame ACME Co. even once they understand how this group agent truly lacks the ability to do otherwise. For many of us, finding out that the group agent has the kind of decision-making procedure that prevents it from reducing profits to

12 In a similar theme, Albertzart argues that group agents cannot be responsible because they do not truly have autonomy ("Monsters and Their Makers"). Though they can act freely, they cannot deliberate about which ends to adopt in the same way that human persons can.

13 See, e.g., Brink and Nelkin, "Fairness and the Architecture of Responsibility"; Fischer and Ravizza, *Responsibility and Control*.

save the environment does not quite seem like the kind of fact that now makes our blame inappropriate or unfair.[14] Such agents do not seem analogous to a company that, say, cannot avoid wrongdoing because the law prohibits it from doing so or because it lacks necessary resources. And yet, if we think that group agents need to possess a certain kind of control over their actions in order to be responsible, namely the ability to do the right thing for the right reasons, then we will have to accept that such agents are, in fact, excused whenever their wrong stems from such local structural deficits (or exempted from our responsibility practices altogether if those deficits are large enough).[15]

One might be tempted to try to explain away our intuition at this point. Perhaps our blame is misfiring and ought to be reserved for the group members for intentionally creating such an insensitive agent. Perhaps we are responding to the wrongness or badness of our imagined group agent's actions, which are not in question and which have been shown to affect people's judgments of culpability.[16] However, such options might have less pull on us when we are aware of an alternative account with which we can offer a principled justification for holding such group agents responsible while also preserving the intuition that agents who lack control over their actions are typically excused.

14 I appreciate that not all readers will share this intuition, particularly consistent control theorists. But it is a hard sell to argue that our blaming responses toward ACME Co. (and psychopaths) ought to be the same as they are toward Wonka Co. (and young children) who are more clearly excused of wrongdoing. These kinds of results at least motivate reconsidering the ultimate bases of blameworthiness. A common reply in cases like these is to argue that the group agent *could* have avoided their lack of capacity at some prior point, which we can "trace" their blameworthiness back to. De Haan, for instance, uses a tracing approach to argue that group agents can be blameworthy for "self-induced moral incapacity," such as gradually sliding into incapacity as a result of group members' corporate greed ("Collective Moral Agency and Self-Induced Moral Incapacity"). There are significant objections to such moves; see Shabo, "More Trouble with Tracing"; Smith, "Attitudes, Tracing, and Control"; Agule, "Resisting Tracing's Siren Song"; and Vargas, "The Trouble with Tracing." For one, they seem to get the phenomenology of our blaming wrong. The target of our blame simply is not failing to ensure that at some point in the future they will have the capacities required to avoid cutting down the rainforest (though this might be an additional source of blameworthiness). Relatedly, for many kinds of actions we want to trace culpability back to, it simply is not the case that that specific wrongdoing was reasonably foreseeable. And even if it was foreseeable, tracing explanations misrepresent the degree of blame we experience. Tracing seems to entail, e.g., that someone who takes heroin one time, knowing it has a risk of addiction, is thereby *fully* blameworthy for *all* wrongdoing that occurs as a result of their addiction since there was one decision said wrongdoings can be traced back to.

15 Wallace, *Responsibility and the Moral Sentiments*.

16 Knobe, "Intentional Action in Folk Psychology."

3. AN ATTRIBUTIONIST ACCOUNT OF GROUP RESPONSIBILITY

Let us reconsider for a moment the group agent that has local structural deficits. Perhaps the group agent is locally blind—it cannot even consider certain kinds of reasons—or perhaps it is locally constrained—it can recognize some considerations in some cases but cannot perform certain kinds of actions. The puzzle is that typically, lack of control or capacity undermines culpability: if I cannot swim, or I am too far away to help, or I am ignorant that anyone needs my help, it seems to be these facts that explain why I am not culpable for failing to save a drowning child. Likewise, if a company fails to give aid to others because it is prohibited by law or because it lacks the means to help them, it will be excused. But when ACME Co. poisons the water supply because its decision-making structure requires it to maximize profits, a plea to be considered excused seems much less convincing than that which would be offered if ACME Co. had no idea a chemical was poisonous or was forced to dump poison because the government mandated it. "Sorry I poisoned your water supply, but I was unable to consider the harm it would cause you as a reason to abstain from harming you" seems to have much less effect tempering our outrage. Indeed, it is tempting to think this blindness is precisely what we want to blame this agent for.

Thompson is not the only person to think group agents are analogous to psychopaths; this similarity has also been noted by Hindriks and Bakan.[17] Though it is true that some philosophers take psychopaths to be excused precisely because they lack the capacity to understand moral reasons, another line of argument is that psychopaths differ from other agents we take to be excused due to a lack of capacity to avoid wrongdoing. Although they cannot understand moral reasons, psychopaths are still capable of making assessments of what kinds of things are or are not reason giving. They are still agents capable of guiding themselves by what reasons they take to be present. And importantly, they are capable of understanding the effects that their actions have on others, e.g., that stabbing someone will cause a lot of pain, frustrate the victim's desire to go on living, and result in death. These features suggest a difference between psychopaths and other agents who lack the capacity to do the right thing. Though they cannot understand the concept of moral status, and so cannot form the attitude "your moral status is a reason to not harm you," they can form the attitude "the fact that this would cause you pain is *not* a reason to abstain from harming you."

According to attributionist theories of moral responsibility, while agents are *typically* excused for wrongdoing when they lack a certain kind of control

17 Hindriks, "How Autonomous Are Collective Agencies?"; Bakan, *The Corporation*.

over their actions or certain kinds of capacities, they are not excused *in virtue of* this lack of control or capacity.[18] Rather, what makes them excused is that many kinds of lack of control prevent an action from being attributable to the agent in the right kind of way. In particular, lack of control often prevents an action from expressing the agent's evaluative judgment about other agents.[19] When someone fails to give aid because they are tied up, cannot swim, or are unaware that someone needs aid, the failure to give aid does not reveal anything objectionable about the agent's take on the person in need.

Importantly, not all capacity deficits are like this. Talbert argues that when we reflect on the agency of psychopaths, it is evident they are different from other agents we typically take to be excused.[20] Given that the psychopath can understand all of the nonmoral facts and express judgments about those facts, this is sufficient to make blaming them appropriate. It is their *denial* of our importance that we find objectionable and which we want to blame them for. And this reasoning is just as appropriate in the case of group agents. When ACME Co. poisons my water supply to increase profits, it is not that they show me ill will in particular. I barely even enter into their deliberation. But it is precisely this *lack* of concern that I and others care about.

This makes them very different from Wonka Co., which poisons the river because it does not understand what effects its effluent has (and has no reason to think it ought to check). The latter's actions do not express any evaluative take on the merits of doing things that poison me, and this is why they are not blameworthy. But ACME Co. is aware that the poison will harm us, and their knowledge of this fact, combined with the decision to dump poison anyway, means that they do have the attitude that the harm to us is *not* a reason to avoid dumping poison. Such attitudes are precisely what make ACME Co. blameworthy.

Whereas most philosophers who think group agents can be blameworthy locate this blameworthiness in the group agent being able to perform actions, the attributionist approach locates it in the group agent being able to have evaluative judgments. This seems notable because, as highlighted earlier, a key

18 This is not simply referring to the attributionist "face" of responsibility, as an earlier reviewer thought, which is often adopted by control theorists and developed in more detail by Shoemaker (*Responsibility from the Margins*). Though Smith, "Responsibility as Answerability," emphasizes responsibility as answerability, she and Talbert, "Blame and Responsiveness," take an agent's evaluative judgments to ground blameworthiness.

19 Hieronymi, "The Force and Fairness of Blame"; Smith, "Moral Blame and Moral Protest." See also Scanlon, *What We Owe to Each Other*. Note, though, that his views have since changed (see Scanlon, *Moral Dimensions*).

20 Talbert, "Blame and Responsiveness to Moral Reasons."

argument in favor of thinking that group agents are genuine agents and not reducible to the aggregate of its group members is that group agents can have attitudes none of its group members themselves endorse. This is a strength of the account: even if Volkswagen's group members individually believe that violating emissions laws is immoral, and strongly preferred that this not take place, the group agent Volkswagen remains blameworthy, and this seems to be because of how that group agent evaluated the merits of its options. In particular, the contribution to climate change that its actions would produce (and which it knew, or should have known, they would produce) is the kind of thing that reveals an objectionable attitude toward those people who will be negatively affected by climate change.

This is not to say that the actions of group members are irrelevant for our assessments, however; an attributionist model of moral responsibility generates some new insights for thinking about the ways that group members' actions influence group agent blameworthiness. Consider the familiar idea that blameworthiness comes in degrees. One factor relevant here is the degree of wrongdoing, which control-based accounts of responsibility can accommodate. But another relevant dimension concerns accounting for the intuition that actions can be more or less blameworthy in virtue of how strongly they are endorsed or how attributable they are to the agent's evaluative orientation. We typically think that someone who experiences significant internal conflict and then commits wrongdoing is less blameworthy than someone who knowingly commits wrongdoing with enthusiasm (though conflict alone surely does not get one off the hook). Control accounts might try to explain this by invoking difficulty as a factor that is relevant to blameworthiness, and which is often present when agents' experience does not fully endorse their actions.[21] But an attributionist approach seems to do a better job of directly accommodating degree of endorsement by taking group agents to be more or less blameworthy in virtue of the extent to which the action was endorsed by its members. For example, it seems that in many cases, all else being equal, if a government's immoral decision is the result of 100 percent of voters voting in favor, this is more blameworthy than an otherwise equivalent decision produced by only 51 percent of voters. The latter action is less attributable to the group agent, even if it remains sufficiently attributable to make the agent blameworthy.

However, the relationship between degree of support among group members and degree of blameworthiness is not always simple. Readers can no doubt recall various instances in which governments failed to act in ways that voters supported, but which seemed to make said governments *more* blameworthy,

21 Nelkin, "Difficulty and Degrees of Moral Praiseworthiness and Blameworthiness."

rather than less. This initially seems to be a problem for my above observation. But group agent theorists emphasize that the attitudes and decisions of group agents are determined not only by group members' votes, but also by how the votes are aggregated.[22] Group agents, which can easily act in ways that most group members strongly reject, are agents in which members have less control over the relevant attitudes and thus are more poorly designed. Poor design does not exculpate group agents, but it does allow us to see how acting in ways with lower support among group members can reduce blameworthiness in some circumstances and increase it in others. In cases where members' votes seem to have insufficient influence on group agent attitudes and actions, owing to a suboptimal aggregation or decision-making procedure, the failure to influence attitudes suggests the group agent has, in their agency, something like a *bias* toward forming certain judgments, and this is what explains our blaming.[23] That the group agent regularly forms certain attitudes or commits wrongdoing *despite* the group members' votes can show that the decision-making procedure, which is a stable part of the group agent's makeup, is having an outsized criticizable effect, and this is what explains our tendency to increase blame.

There is one last feature which may be affecting our intuitions that is worth identifying, and this concerns the level of stability in decision-making that the decision-making procedure allows. Group members might continually fail to have sufficient influence on the group agent's attitudes and decisions, but those attitudes and decisions might not manifest something like a bias because the resulting attitudes and decisions are too inconsistent, unstable, or haphazard. In short, we might have discordance that suggests the agent is less responsive to reasons altogether.[24] While it is common to talk of "being an agent" as if it were a threshold notion, agency, in fact, comes in degrees, evidenced by there being no clear point between birth and adulthood in which one becomes a

22 First-past-the-post voting, for instance, tends to produce different results than those of mixed-member proportional representation, even if each group member's vote remains the same under both systems.

23 Here I do not mean something analogous to implicit bias, which has received a lot of attention and is commonly taken to be defined by the way that it *conflicts* with the agent's explicit attitudes and has little impact on agential decision-making. I mean bias in the traditional sense in which someone, e.g., continually favors their own group despite there being no adequate justification for this, or selectively interprets evidence and misrepresents challenges to their view due to motivated reasoning.

24 A subtle point: there is a sense in which the incorrigible racist is not responsive to reasons in that he will not change when we argue with him. But what attributionists take to matter is that the agent's attitudes are responsive to what the agent takes to be the case, rather than whether their attitudes accurately reflect the reasons they, in fact, have (Smith, "Attitudes, Tracing, and Control," 125–26).

morally responsible agent. Likewise, a group agent whose decision-making procedure leads to decisions that are too haphazard seems to be less of an agent and could be more analogous to a young child or someone with certain mental disabilities. To be sure, agents need not be perfectly consistent; Smith emphasizes that someone can fear a spider while also sincerely claiming to believe the spider is perfectly safe, and both of these attitudes (along with the charge of irrationally holding inconsistent attitudes) will be attributable to the agent without impugning her status as an agent.[25] But too much inconsistency can put one out of the agency and responsibility game altogether.

4. GROUP AGENTS AND JUSTIFYING BLAME

Now that we have considered how an attributionist account of responsibility avoids the local structural deficits objection, we can also see how it avoids the inability to care objection. The fact that group agents cannot form attitudes that have people as their object is no barrier to blaming them. What concerns us is their take on what things are or are not reason giving, and this is something they are able to do.[26]

Admittedly, this result relies on adopting a different account of blame and blameworthiness. This may be unsatisfying for some readers who interpret Thompson's argument as a conditional (*if* Strawsonian accounts of moral responsibility are correct, group agents cannot be blameworthy) and take attributionists to be simply rejecting the premise. We would have more reason to support the attributionist story if there were independent objections to Thompson's argument, or reasons to not grant the premise. These can be provided by examining some background considerations regarding Strawsonian accounts of moral responsibility that Thompson is somewhat unclear on, which any future treatments on the responsibility of group agents should be sensitive to.

On Strawsonian accounts, our reactive attitudes are responses to others' quality of will or level of regard. But most philosophers take the reactive attitudes to respond to, and thus occur *downstream* of, the blameworthy-making features of agents (though the reactive attitudes can be good evidence of blameworthiness). On this *response-independent* interpretation of blameworthiness, showing that group agents cannot experience resentment, guilt, or indignation does not yet show that blame is inapt, as these are things that occur after someone has displayed poor quality of will. That we cannot create guilt in them now

25 Smith, "Attributability, Answerability, and Accountability," 579.

26 Note this account also avoids worries that without consciousness, group agents are analogous to zombies (O'Madagain, "Group Agents").

(meaning blame might be ineffective or pointless) will not affect whether they are blame*worthy*. Instead, Thompson would need to show that moral emotions are essential to displaying blameworthy quality of will (or essential to some other factor that grounds blameworthiness) and that group agents cannot experience these emotions.

Taking reactive attitudes to be downstream of blameworthy-making features of agents is distinct from the position most commonly attributed to Strawson, known as a *response-dependent* conception of responsibility.[27] On this account, the reactive attitudes are crucial to understanding responsibility because there is no external justification for holding people responsible beyond the fact that ill will or lack of regard is what properly trained human reactive attitudes respond to.[28] This seems to be closer to what Thompson has in mind. On this account, the reactive attitudes are not just an inseparable part of holding agents responsible; they are also constitutive of those agents *being* responsible. It is our proneness to experiencing the reactive attitudes to an agent's quality of will that makes the agent responsible.

Holding responsible is typically thought to be linked to being responsible via the demands our blame expresses.[29] Thompson, however, links them via blame's purported function. He thinks that blame's function is to produce moral alignment of the wrongdoer and victim's understanding of the wrong. In particular, it aims to produce guilt, which is linked to resentment (blame) because guilt and resentment view the same wrong from different perspectives. Importantly, he takes guilt to involve caring, and caring to require affective attitudes, evidenced by the fact that our blame does not cease if wrongdoers behave as if they feel guilt.

With this background clarified, we are now in a position to note some costs to the overall argument. One is that a response-dependent account leaves Thompson in something of a minority position; many philosophers and folk alike take our blame to be responding to facts about wrongdoers that are independent of our actual (properly trained) blaming practices which ground blameworthiness. For example, when asking what *makes* him blameworthy, *why* that feature makes him blameworthy, and whether he is *really* blameworthy,

27 For a review of the trickiness of articulating his exact position, see Todd, "Strawson, Moral Responsibility, and the 'Order of Explanation.'"

28 Shoemaker, "Response-Dependent Responsibility."

29 As Strawson puts it: "the making of the demand *is* the proneness to such attitudes" ("Responsibility as Answerability," 207). For more on this point, see also Darwall, *The Second-Person Standpoint*. Cf. McKenna, who argues that ability to make demands is intertwined with the ability to respect demands because our blaming interactions are similar to conversations that require understanding meaning (*Conversation and Responsibility*).

answering with an appeal to the mere fact that we have tendencies to blame agents like him seems unsatisfying.[30]

More substantive costs concern the proposal that blame's function is to produce moral alignment by making our target care, which is identified by reflecting on our practices of blaming and what its point is.[31] As Thompson points out, "We do not merely accept a form of behaviorism. We do not demand that others merely act *as if* they experience guilt for their actions."[32] But there are three problems here. The first is that if we are appealing to our actual practices to determine the proper objects of our blame, it seems that many people *do* blame group agents. Some readers who are skeptics about the responsibility of group agents may not personally have the intuition that such agents are responsible, but it is undeniable that many members of our moral community routinely blame group agents in a manner identical to the way they blame individual group agents. Claims about what *our* moral practices are require some explanation for why, though blame toward individual agents has the aim of making them care, seemingly identical blame toward group agents either does not have this function or can only be interpreted as a mistake.

Relatedly, the observation that we do not accept mere behaviorism seems difficult to square with the observation that many people's blame toward group agents does subside in response to how the group agent responds to its earlier wrongdoing. Group agents are capable of offering apologies, attempting to repair relationships, and signaling that their failures will not happen again. And, indeed, failures to do these things typically generate even more blame from us.

Finally, observing we want something for our blame to cease does not show that thing is what makes our blame appropriately begin. Such reasoning seems

30 That this is not how most philosophers writing on the responsibility of group agents would think about moral responsibility is evidenced by how the literature as a whole has progressed. As noted earlier, philosophers have been concerned with investigating whether group agents possess the same kind of relevant features that justify considering them *bona fide* agents, who perform actions, and whom we can *justify* blaming and holding responsible, or whether there are relevant differences between group agents and ordinary agents (e.g., ability to care). But for the response-dependent approach, this framing may be wrongheaded as there is no independent justification for our reactive attitudes. If we have granted that group agents are agents, the facts about which agents are responsible is determined by whatever (properly developed and informed) human sensibilities deem to be fitting targets of blame. And this, it seems, is to be determined by looking inward to our actual moral practices and the fittingness conditions of our emotions, rather than considering principled arguments regarding functionalism, agency and intentional objects.

31 It should also be noted that appeals to the benefits of aligning understanding (Thompson, "The Moral Agency of Group Agents," 524) risk appealing to an external source of justification.

32 Thompson, "The Moral Agency of Group Agents," 526.

to imply that if someone committed wrongdoing and felt guilt immediately after doing it (perhaps they even felt guilty in the lead-up to the wrong, realizing and caring about how the wrong would affect us), then our blame would be pointless. But if our blame is pointless, then it seems like our blame ought not even begin, meaning the agent *would not be blameworthy*. But that is clearly not the case; agents who commit wrong while experiencing considerable guilt, and a proper understanding of their actions, can nevertheless be appropriate targets of blame.

5. MORAL ALIGNMENT AND PROTEST

The previous section's objections concerned the way in which we use our moral practices to determine who counts as blameworthy. But even if we set aside our concerns with justification and blameworthiness, taking blame's function to be moral alignment seems to also generate a few notable discrepancies with our blaming practices. In particular, we do not think our blame is inapt or unjustified even if we know it will be ineffective, and there are a variety of ways that blame aimed at producing alignment can be rendered ineffective.[33] Perhaps our target is simply incorrigible. It is also difficult to see how any alignment could be achieved when we blame historical figures or the dead, or when no one is around to see our blame. Another difficulty is that blame seems to achieve little when our target already feels guilt and so seems to have already acquired the same understanding as us. And sometimes we blame expressly without the aim of producing any alignment. We often feel outrage and are not interested in what our target has to say in response or how they think about what they did. Sometimes it is apt to storm off and not talk to the wrongdoer. Sometimes it is apt to continue blaming after they feel guilty. And even if group agents did somehow achieve consciousness and could thereby take persons as intentional objects or feel guilt, it does not seem like the character of our blame would dramatically change; alignment still would not be the goal. Just like when we blame politicians on the television, most of the time we do not expect our blame toward corporations or governments to have much effect on the target. We do not expect any response from them to us *qua* individuals at all, and we do not always make an effort to make sure said agents notice our blame. If moral alignment is the goal, blame is often an inefficient way to do it.

One can try to account for these discrepancies by arguing that we are taking a paradigm-based approach to our theorizing, or argue that blame is a speech act and speech acts can be unexpressed, or try to locate the alignment in

33 Tierney, "Guilty Confessions."

someone other than the target of our blame, as Fricker does.[34] But at a certain point, such discrepancies look more like counterexamples and seem against the spirit of the Strawsonian emphasis on our actual moral practices. These observations at least motivate looking for an alternative.

An attributionist account of blame is much better placed to make sense of these aspects of our blaming tendencies. Rather than taking blame's function to be producing moral alignment between wrongdoer and blamer, attributionists instead take blame to be a form of *protest*.[35] Our blame's focus is on the wrongdoing and the threat that the attitude expressed in the wrongdoing poses to our moral standing. In blaming, we are standing up for ourselves or another victim. Although protest is communicative, it does not only communicate to the wrongdoer. It also expresses outrage to other members of the moral community (which may include group members of the group agent) who have a role in maintaining moral standards. Unlike the moral alignment account, which seems to make blame pointless once alignment is achieved, communication more clearly does not merely function to transmit information. Many forms of communication have an expressive point, such as telling a spouse you love them even when this was never in doubt. Protest also makes better sense of our goal in blaming. The moral alignment account had trouble accounting for the fact that many instances of blame seem to not explicitly aim at alignment, such as storming off, ceasing interaction altogether, or blaming when alignment has been achieved. But the protest account does much better because protest denounces or repudiates the attitude that was expressed in the group agent's wrong. When I storm off in response to your wrongdoing, I am protesting against the attitude that was expressed, whether or not you, in particular, get the message and whether or not you already perfectly understand how the wrong affected me. This, in turn, allows us to make sense of how blame that is directed at a group agent but which is not performed in a way likely to evoke a response (e.g., because they are a foreign government, and you are just in a university classroom) has not thereby misfired. Protest against something that is not the case, however, does seem pointless, and this matches the thought that blaming someone who has not culpably done anything wrong is inappropriate.

Even if one accepts that attributionist accounts of moral responsibility and blame can overcome the objections raised against the possibility of group agent moral responsibility, one might object to attributionist accounts on independent grounds. One might argue that attributionist accounts are implausible,

34 Fricker, "What's the Point of Blame?"

35 Hieronymi, "Articulating an Uncompromising Forgiveness"; Smith, "Moral Blame and Moral Protest"; Talbert, "Moral Competence."

citing various factors often used to support control-based accounts of moral responsibility.

While I cannot resolve the debate between control theorists and attributionists here, it is worth noting that traditional objections to attributionist accounts of moral responsibility are much less persuasive when it comes to considering the responsibility of group agents. One of the main fault lines in the debate over moral responsibility concerns what ultimately justifies blaming culpable wrongdoers.[36] Control theorists typically argue that the ability to avoid wrongdoing is needed because blame is a negative treatment or sanction which is against one's interests. For blame to be apt, it must be *deserved*, and people do not deserve things which they could not avoid.[37] This is partly what grounds the thought that psychopaths are not blameworthy: since they lacked the capacity to understand moral reasons, they could not choose to act on those reasons and so cannot be blameworthy for failing to do so.

However, a concern for desert is much less pressing when considering the responsibility of group agents. There is strong support for the idea that group agents do not merit as much consideration as ordinary agents, and List and Pettit believe we should not extend the same rights to group agents that we give to individuals, such as the right to vote.[38] Additionally, some control theorists argue that part of what makes blame unpleasant and deserved is that it induces guilt, understood here as a *pained* recognition of what wrongs one has done.[39] Such arguments also cannot be used to justify blaming group agents because their lack of phenomenal consciousness means they are unable to experience pain. And given that pain seems to be bad in virtue of its phenomenal quality, a functional analogue of pain is unlikely to be a sufficient alternative for our theory.

If we are already much less concerned with the interests of group agents, and our ordinary reason to be careful with our blame—that it induces pained recognition of wrongdoing—simply does not apply, then arguments that we should favor a control requirement on blameworthiness for group agents over a protest-based account of blame (and an attributionist account of blameworthiness)

36 Rudy-Hiller, "It's (Almost) All about Desert."

37 Levy, *Hard Luck*; Nelkin, "Desert, Fairness, and Resentment." Group agents can "deserve" things like prizes in virtue of meeting the conditions stipulated in contests, but this is not the same as *basic* desert, which is what most control theorists take to be at issue (Pereboom, *Free Will, Agency, and Meaning in Life*). For some reflections on how attributionists justify blame, see Hieronymi, "I'll Bet You Think This Blame Is about You."

38 Hindriks, "How Autonomous Are Collective Agents?"; List and Pettit, *Group Agency*, 180–82.

39 Carlsson, "Blameworthiness as Deserved Guilt."

are much harder to get off the ground. It is perhaps even possible that one could be a control theorist about individual moral responsibility while being an attributionist about group agent moral responsibility, but I will set aside the possibility of defending such a position.[40]

6. CONCLUSION

This paper has argued that an attributionist account of moral responsibility is well-suited to make sense of our practices of blaming group agents and holding them morally responsible. Even though group agents cannot experience guilt, cannot feel pain, and can sometimes lack the ability to avoid wrongdoing, these factors are not barriers to them being blameworthy. This is because group agents are the kinds of things which, in virtue of their reasons-responsive decision-making structure, are able to make assessments about what kinds of actions are worthwhile and, importantly, what kinds of considerations are *not* reason giving. When their actions reflect attitudes or assessments that are objectionable, group agents are blameworthy, and our blame toward them is warranted.

University of Notre Dame, Australia
adam.piovarchy@nd.edu.au

REFERENCES

Agule, Craig. "Resisting Tracing's Siren Song." *Journal of Ethics and Social Philosophy* 10, no. 1 (January 2016): 1–24.

Albertzart, Maike. "Monsters and Their Makers: Group Agency without Moral Agency." In *Reflections on Ethics and Responsibility: Essays in Honor of Peter A. French,* edited by Zachary J. Goldberg, 21–35. Cham, Switzerland: Springer, 2017.

Baddorf, Matthew. "Phenomenal Consciousness, Collective Mentality, and Collective Moral Responsibility." *Philosophical Studies* 174, no. 11 (November 2017): 2769–86.

Bakan, Joel. *The Corporation: The Pathological Pursuit of Profit and Power*. New York: Free Press, 2004.

Björnsson, Gunnar, and Kendy Hess. "Corporate Crocodile Tears? On the

40 Elsewhere I have endorsed control accounts of responsibility; these arguments should be considered independent.

Reactive Attitudes of Corporate Agents." *Philosophy and Phenomenological Research* 94, no. 2 (March 2017): 273–98.

Brink, David O., and Dana K. Nelkin. "Fairness and the Architecture of Responsibility." In *Oxford Studies in Agency and Responsibility*, vol. 1, edited by David Shoemaker, 284–314. Oxford: Oxford University Press, 2013.

Carlsson, Andreas Brekke. "Blameworthiness as Deserved Guilt." *Journal of Ethics* 21, no. 1 (March 2017): 89–115.

Darwall, Stephen L. *The Second-Person Standpoint: Morality, Respect, and Accountability*. Cambridge, MA: Harvard University Press, 2006.

de Haan, Niels. "Collective Moral Agency and Self-Induced Moral Incapacity." *Philosophical Explorations* 26, no. 1 (2023): 1–22.

Fischer, John Martin, and Mark Ravizza. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press, 1998.

French, Peter A. *Collective and Corporate Responsibility*. New York: Columbia University Press, 1984.

Fricker, Miranda. "What's the Point of Blame? A Paradigm Based Explanation." *Noûs* 50 no. 1 (March 2016): 165–83.

Gilbert, Margaret. "Collective Guilt and Collective Guilt Feelings." *Journal of Ethics* 6, no. 2 (June 2002): 115–43.

———. *Joint Commitment: How We Make the Social World*. New York: Oxford University Press, 2013.

Gilbert, Margaret and Daniel Pilchman. "Belief, Acceptance, and What Happens in Groups: Some Methodological Considerations." In *Essays in Collective Epistemology*, edited by Jennifer Lackey, 189–212. Oxford: Oxford University Press, 2014.

Hieronymi, Pamela. "Articulating an Uncompromising Forgiveness." *Philosophy and Phenomenological Research* 62, no. 3 (May 2001): 529–55.

———. "The Force and Fairness of Blame." *Philosophical Perspectives* 18, no. 1 (December 2004): 115–48.

———. "I'll Bet You Think This Blame Is about You." In *Oxford Studies in Agency and Responsibility*, vol. 5, edited by D. Justin Coates and Neal A. Tognazzini, 60–87. Oxford: Oxford University Press, 2019.

Hindriks, Frank. "Collective Agency: Moral and Amoral." *Dialectica* 72, no. 1 (March 2018): 3–23.

———. "How Autonomous Are Collective Agents? Corporate Rights and Normative Individualism." *Erkenntnis* 79, no. S9 (October 2014): 1565–85.

Knobe, Joshua. "Intentional Action in Folk Psychology: An Experimental Investigation." *Philosophical Psychology* 16, no. 2 (2003): 309–24.

Levy, Neil. *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*. New York: Oxford University Press, 2011.

List, Christian, and Philip Pettit. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press, 2011.

McKenna, Michael. *Conversation and Responsibility*. New York: Oxford University Press, 2012.

Nelkin, Dana K. "Desert, Fairness, and Resentment." *Philosophical Explorations* 16, no. 2 (2013): 117–32.

———. "Difficulty and Degrees of Moral Praiseworthiness and Blameworthiness." *Nous* 50, no. 2 (June 2016): 356–78.

———. "Psychopaths, Incorrigible Racists, and the Faces of Responsibility." *Ethics* 125, no. 2 (January 2015): 357–90.

O'Madagain, Cathal. "Group Agents: Persons, Mobs, or Zombies?" *International Journal of Philosophical Studies* 20, no. 2 (2012): 271–87.

Pereboom, Derek. *Free Will, Agency, and Meaning in Life*. New York: Oxford University Press, 2014.

Rudy-Hiller, Fernando. "It's (Almost) All about Desert: On the Source of Disagreements in Responsibility Studies." *Southern Journal of Philosophy* 59, no. 3 (September 2021): 386–404.

Scanlon, T. M. *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, MA: Belknap Press of Harvard University Press, 2008.

———. *What We Owe to Each Other*. Cambridge, MA: Belknap Press of Harvard University Press, 2000.

Shabo, Seth. "More Trouble with Tracing." *Erkenntnis* 80, no. 5 (October 2015): 987–1011.

Shoemaker, David. "Response-Dependent Responsibility; or, A Funny Thing Happened on the Way to Blame." *Philosophical Review* 126, no. 4 (October 2017): 481–527.

———. *Responsibility from the Margins*. Oxford: Oxford University Press, 2015.

Smith, Angela M. "Attitudes, Tracing, and Control." *Journal of Applied Philosophy* 32, no. 2 (May 2015): 115–32.

———. "Attributability, Answerability, and Accountability: In Defense of a Unified Account." *Ethics* 122, no. 3 (April 2012): 575–89.

———. "Moral Blame and Moral Protest." In *Blame: Its Nature and Norm*, edited by D. Justin Coates and Neal A. Tognazzini, 27–48. New York: Oxford University Press, 2013.

———. "Responsibility as Answerability." *Inquiry* 58, no. 2 (2015): 99–126.

Strawson, Peter. "Freedom and Resentment." *Proceedings of the British Academy* 48 (1962): 187–211.

Talbert, Matthew. "Blame and Responsiveness to Moral Reasons: Are Psychopaths Blameworthy?" *Pacific Philosophical Quarterly* 89, no. 4 (December 2008): 516–35.

———. "Moral Competence, Moral Blame, and Protest." *Journal of Ethics* 16, no. 1 (March 2012): 89–109.

Tanguay-Renaud, Francois. "To Fill or Not to Fill Individual Responsibility Gaps?" In *The Legacy of Ronald Dworkin*, edited by Wil Waluchow and Stefan Sciaraffa, 71–98. Oxford: Oxford University Press, 2016.

Thompson, Christopher. "The Moral Agency of Group Agents." *Erkenntnis* 83, no. 3 (June 2018): 517–38.

Tierney, Hannah. "Guilty Confessions." In *Oxford Studies in Agency and Responsibility*, vol. 7, edited by David Shoemaker, 182–204. Oxford: Oxford University Press, 2021.

Todd, Patrick. "Strawson, Moral Responsibility, and the 'Order of Explanation': An Intervention." *Ethics* 127, no. 1 (October 2016): 208–40.

Tollefsen, Deborah Perron. "Participant Reactive Attitudes and Collective Responsibility." *Philosophical Explorations* 6, no. 3 (2003): 218–34.

Tuomela, Raimo. *Social Ontology: Collective Intentionality and Group Agents*. New York: Oxford University Press, 2013.

Tuomela, Raimo, Raul Hakli, and Pekka Mäkelä, eds. *Social Ontology in the Making*. Boston: De Gruyter, 2020.

Vargas, Manuel "The Trouble with Tracing." *Midwest Studies in Philosophy* 29, no. 1 (September 2005): 269–91.

Wallace, R. Jay. *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press, 1994.

Watson, Gary. "The Trouble with Psychopaths." In *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon,* edited by R. Jay Wallace, Rahul Kumar, and Samuel Freeman, 307–31. New York: Oxford University Press.

Ziv, Anita Konzelmann. "Collective Guilt Feeling Revisited." *Dialectica* 61, no. 3 (2007): 467–93.