

## Group Responsibility

Christian List\*

February 2019, with minor edits in October 2021

**Abstract:** Are groups ever capable of bearing responsibility, over and above their individual members? This chapter discusses and defends the view that certain organized collectives – namely, those that qualify as group moral agents – can be held responsible for their actions, and that group responsibility is not reducible to individual responsibility. The view has important implications. It supports the recognition of corporate civil and even criminal liability in our legal systems, and it suggests that, by recognizing group agents as *loci* of responsibility, we may be able to avoid “responsibility gaps” in some cases of collectively caused harms for which there is a shortfall of individual responsibility. The chapter further asks whether the view that certain groups are responsible agents commits us to the view that those groups should also be given rights of their own and gives a qualified negative answer.

A central feature of our social world is our practice of assigning responsibility to one another for our actions. We treat one another as responsible agents, and on this basis, we praise one another for some actions and blame one another for others. The practice of assigning responsibility underlies both civil and criminal law. This raises the question of what the *loci* of responsibility are: who or what are the entities capable of bearing responsibility? Are individual human beings the only such entities, or are groups, especially organized collectives, appropriate targets for the assignment of responsibility as well? Given the influential roles that many collectives play in the social world, from states and corporations to courts and committees, this question is of great importance for morality and politics as well as for the law.

The aim of this chapter is to describe and defend the view that certain organized collectives – namely, group agents – can be held responsible for their actions, over and above their individual members, and that group responsibility is not reducible to individual responsibility.<sup>1</sup> This view has some important implications. There are some things that happen in the social world, including some harms, that cannot easily be attributed to any single individual or even any number of individuals. If individual human beings were the only entities capable of bearing responsibility, then we might not be able to identify anyone who bears sufficient responsibility

---

\* I am very grateful to Hein Duijf for helpful comments on this chapter.

<sup>1</sup> The literature on the responsibility of collectives is large, and I cannot provide a comprehensive review here. See, among others, Feinberg (1968), Held (1970), French (1984), May (1987), Erskine (2001), Copp (2006), Pettit (2007), List and Pettit (2011, ch. 7), Pasternak (2013), and Tollefsen (2015). For an overview, see Smiley (2017).

in such cases. We might be faced with “responsibility gaps” or “responsibility voids”: situations where some significant harm has occurred, but there is a shortfall of attributable responsibility.<sup>2</sup> If, on the other hand, there are responsible group agents to whom the actions in question can be attributed, we may be able to avoid such responsibility gaps. And indeed, many jurisdictions recognize some forms of corporate civil and even criminal liability, holding certain collective entities responsible over and above their individual members. If the notion of group responsibility can be defended, then this legal practice stands on a firm footing.

The chapter is structured as follows. In Section 1, I will introduce the problem of group responsibility in more detail and put this chapter’s question into sharper focus. In Section 2, I will say more about the conditions an entity must meet in order to count as “fit to be held responsible”. While adult human beings are normally assumed to meet those conditions, it is much less clear whether collectives could do so too. In Sections 3 and 4, I will introduce the idea of group agency and explore whether group agents can be fit to be held responsible. In Section 5, I will look at the relationship between individual and group responsibility and outline the case for the irreducibility of group responsibility. In Section 6, finally, I will ask whether the view that certain groups are responsible agents also commits us to the view that those groups should be given rights of their own, and I will give a qualified negative answer.

An important disclaimer is due: my focus in this chapter is on the question of whether certain groups can be *loci* of responsible agency in their own right; it is not on what responsibilities individuals have in the context of their participation in collective actions or their membership in certain groups. I will touch on the latter question only in passing. In taking the present approach to the subject of group responsibility, I am building on my earlier work on group agency, especially my joint work with Philip Pettit.<sup>3</sup> Others would undoubtedly approach the subject from different angles, but given space constraints, I cannot cover the subject in its entirety here.

## **1. The problem of group responsibility**

Our lives are profoundly affected by the activities of many collectives. These include not just firms and commercial corporations, but also states, other governmental and non-governmental organizations, courts, health-care providers, universities, expert panels, clubs, and many

---

<sup>2</sup> See, e.g., Braham and van Hees (2011, 2018), Collins (2017, 2018), Duijf (2018a), and Buell (2018).

<sup>3</sup> See List and Pettit (2011) and List (2018a,b). See further Pettit (2007), from whom I am borrowing the language of “fitness to be held responsible”. The present approach has many precursors in the literature; recall footnote 1. The chapter draws significantly on List (2018b), which, for the academic record, should be regarded as prior.

informal groups. All of these make a difference to people's lives through their decisions and actions, and sometimes cause harms, either by accident or as more or less predictable externalities or side-effects of what they do. Examples readily come to mind: from the Deepwater Horizon oil spill in the Gulf of Mexico, caused by BP's drilling operations, through the environmental degradation caused by big industry, to the recent Windrush scandal in the UK, where the rights of many legal residents were infringed due to a change in the country's immigration policies. As is widely acknowledged, when harms result from the actions of collective entities, it is not always easy to identify one or several individuals to whom all of the relevant responsibility can be attributed. It is possible that most, or even all, of the individual members of the collectives in question have acted in ways that were reasonable, given their circumstances, and are therefore individually blameless or only blameworthy to a small extent, and yet their individual contributions have jointly produced a very harmful outcome. In this way, the sum-total of individual responsibility may fall short of the full amount of responsibility that we might intuitively find appropriate for the harm caused. There seems to be a shortfall of attributable responsibility: a partial or even complete "responsibility gap" or "responsibility void".<sup>4</sup>

An illuminating example, given by Philip Pettit, is that of the Herald of Free Enterprise, a car ferry in the English Channel which sank on its way from Zeebrugge to Dover in 1987.<sup>5</sup> Almost two hundred people died. An inquiry found that the operating company, Townsend Thoresen, was extremely sloppy and had very bad safety standards. As a commentator put it, "[f]rom top to bottom the body corporate was infected with the disease of sloppiness".<sup>6</sup> Nonetheless, strikingly, nobody was found legally responsible to an extent that seemed commensurate with the harm that had occurred. The same commentator noted: "the primary requirement of finding an individual who was liable ... stood in the way of attaching any significance to the organizational sloppiness that had been found by the official enquiry".<sup>7</sup> Because the ferry company was not considered a *locus* of responsibility, there was no agent to

---

<sup>4</sup> Recall the references in footnote 2. Interestingly, responsibility gaps might be viewed as involving a violation of what is sometimes called "moral harmony". In the words of Portmore (2018, p. 318), an account of morality is "harmonious" if "the agents who satisfy [it], whoever and however numerous they may be, are guaranteed to produce the morally best world that they have the option of producing." Perhaps, then, responsibility gaps could occur only against the background of an insufficiently harmonious criterion of individually reasonable action. That said, the envisaged scenarios seem familiar enough from the perspective of commonsense morality to be worth investigating. I am grateful to Hein Duijf for pressing me to address this point.

<sup>5</sup> See Pettit (2007) and List and Pettit (2011, ch. 7).

<sup>6</sup> See Colvin (1995), quoted in Pettit (2007).

<sup>7</sup> Ibid.

whom the relevant responsibility could be assigned. The harm was treated, in effect, as if it had been the result of a natural disaster or bad brute luck, not the result of anyone's wrong-doing.

One view about responsibility, which we may call "individualism", would vindicate this outcome. According to an individualist view about responsibility, only human beings, especially adult human beings, are appropriate targets for the assignment of responsibility; collectives are not. On such a view, it is appropriate to ask how various individuals have contributed to a harmful outcome and, if negligence, recklessness, or wrong-doing can be established, to hold them individually accountable. But, on the individualist view, the collective itself cannot be held responsible over and above its members. If, even after the most careful accounting, the sum-total of individual responsibility falls short of the amount of responsibility that would seem commensurate with the harm done, then we would have to bite the bullet and accept that no further responsibility can be assigned.

A competing view is "collectivism". According to it, collectives may sometimes be held responsible in their own right, independently of the individuals involved. The key idea is that there can in principle be collective responsibility without any corresponding individual responsibility. In limiting cases, a collective entity could be held responsible for some outcome without any of its members bearing any individual responsibility. Of course, the details of when this is so, and what exactly it means, need to be spelt out further, and I will turn to them later.

The collectivist view is sometimes motivated by the observation that unless we are prepared to assign responsibility to certain collectives in their own right, the occurrence of some undesirable responsibility gaps is inevitable in practice. However, from the premises that

- (1) on the individualist view, there may sometimes be responsibility gaps, and that
- (2) responsibility gaps are undesirable,

we cannot derive the conclusion that

- (3) the collectivist view is correct.

The argument from (1) and (2) to (3) would be a non-sequitur. Premises (1) and (2) are insufficient to establish collectivism about responsibility. Perhaps it is simply an unfortunate fact about morality that there is sometimes a shortfall of individually attributable responsibility, and it is wishful thinking to expect that in every such case one could close the gap by identifying a responsible collective.

Indeed, we can easily imagine cases of collectively caused harms for which it is implausible to hold anyone – whether individuals or the collective – responsible. Imagine, for instance, a freak accident in a town centre, where, due to an unusual and unforeseeable confluence of pedestrian traffic, a stampede occurs and people get injured, but neither the pedestrians

involved, nor the police, nor the local government did anything wrong. If this truly is a freak accident, then there is no responsibility to be assigned, given that there was nothing wrong with anyone's conduct: no negligence or reckless, let alone intentional wrong-doing. Furthermore, in this case, the relevant collective – the various pedestrians – seems far too amorphous to count as a plausible *locus* of responsibility.<sup>8</sup>

In fact, one line of reasoning would suggest that it is actually a misnomer to call this a responsibility *gap*. There simply is no responsibility here. It is no different from a natural disaster, as far as responsibility is concerned. We would not normally look for a responsible agent in the case of a genuine natural disaster.<sup>9</sup> Now, the proponents of an individualist view about responsibility might insist that the case of the Herald of Free Enterprise is similar. If, as we may stipulate for the sake of argument, there was no wrong-doing on the part of any individuals, then the ferry disaster might be like the stampede or a natural disaster.

If we still want to argue that the ferry company should be held responsible as a collective entity, independently of the extent of any wrong-doing by relevant individuals, we need a more careful account of the conditions that an entity must meet in order to count as “fit to be held responsible”. Once we have specified those conditions, we are in a better position to adjudicate whether collectives could be bearers of responsibility in their own right.<sup>10</sup>

## 2. The conditions for responsibility

Adult human beings are paradigmatic examples of entities that we consider capable of bearing responsibility for their actions: they qualify as “fit to be held responsible”, as Philip Pettit puts it.<sup>11</sup> Are adult human beings the only such entities, or could other entities also qualify?

---

<sup>8</sup> On the distinction between organized and random collections, and on whether the latter can bear responsibility, see also Held (1970).

<sup>9</sup> I speak of “genuine” natural disasters here in order to set aside those disasters that have a plausible anthropogenic source, such as anthropogenic-climate-change-induced disasters. The latter arguably do raise questions of responsibility.

<sup>10</sup> Hein Duijf (in correspondence) has suggested the following more nuanced analysis of the (formally invalid) argument from (1) and (2) to (3). First, note that the claim that there may sometimes be responsibility gaps, as asserted by premise (1), implicitly relies on some background theory of responsibility. After all, it is claimed that individuals are *not* responsible, and yet that something has happened that would somehow warrant the attribution of responsibility. Now, if we further accept the claim that such responsibility gaps are undesirable, as asserted by premise (2), then we may well conclude that our background theory of responsibility is in some ways defective, and that we should revise it. At this point, several avenues are open to us. One would be to deny that there are any responsibility gaps. A second would be to revise the criteria for the attribution of individual responsibility. And a third would be to extend the *loci* of responsibility, by recognizing collectives as responsible agents. Finally, one might combine two or more of these avenues. Against the background of this analysis, the argument from (1) and (2) to (3) is still invalid, but premises (1) and (2) lend some support to the conclusion that *something* needs to be revised, even though accepting (3) is not the only possible way this revision could go. See further Duijf (2018b).

<sup>11</sup> See Pettit (2007).

Obviously, we don't want to answer this question just by stipulation. Rather, we would like to appeal to some general conditions for fitness to be held responsible. These will hopefully vindicate our assumptions in the paradigmatic human case, while helping us to clarify what is going on in other cases. For the purposes of this article, I will treat the following conditions as necessary and sufficient for fitness to be held responsible.<sup>12</sup>

**Moral agency:** The entity is an intentional agent and has the capacity to make normative judgments about its choices and to be guided by those judgments.

**Relevant information:** The entity has the information needed for assessing its choices normatively, or it has at least reasonable access to such information.

**Free choice:** The entity has the freedom and control required for making its choices.

To see that these conditions are indeed required for fitness to be held responsible, let us briefly go through them. The first condition – moral agency – is evidently necessary. Entities that do not qualify as intentional agents, such as rocks, arm chairs, or hammers, do not even begin to qualify as candidates for bearing responsibility. But intentional agency alone is not enough. There are many entities, most notably some non-human animals, that plausibly qualify as *intentional* agents without being *moral* agents. They exhibit systematic, goal-directed behaviour, but they lack the capacity to make judgments about what is right and wrong, permissible and impermissible, and to act on the basis of those judgments. And for this reason, we do not hold them responsible for anything. The second condition – relevant information – is equally necessary. If someone lacks the information – or even reasonable access to the information – required for a normative assessment of his or her choices, perhaps due to the constraints of his or her environment, then it would not be justified to hold him or her responsible (and certainly not fully responsible) for choices made in those circumstances. Finally, the third condition – free choice – captures the widely accepted idea that freedom and control, of some sort, are necessary for responsibility. If someone has no freedom or control in relation to his or her actions – perhaps he or she lacks alternative possibilities or has no causal control over what he or she does – then it would seem problematic to hold him or her fully responsible for those actions.<sup>13</sup> Consistently with this, freedom-or-control-undermining

---

<sup>12</sup> The conditions build on the ones in List and Pettit (2011, ch. 7), which, in turn, are variants of conditions in Pettit (2007).

<sup>13</sup> There are some complications here that I cannot go into, given space constraints. For instance, Frankfurt (1969) has argued that alternative possibilities are not needed for responsibility. I accept that one may debate the precise nature of the “free choice” requirement for responsibility, but I assume that at least some suitably qualified such requirement is needed. For further discussion, see also Alvarez (2009) and List (2019, pp. 168–169).

constraints, whether external or psychological, are often thought to be responsibility-undermining as well. While these considerations suggest that the three conditions are each necessary for fitness to be held responsible, I will also assume that they are jointly sufficient.

It should be evident that it is very much a contingent question which entities meet these conditions. While adult human beings paradigmatically satisfy them – or so we tend to believe – we must not assume from the outset that some non-human entities could not do so as well. There is no conceptual reason, for instance, why future sophisticated robots could not also satisfy them; nor should we rule out the possibility that suitably organized collectives could do so. Of course, much depends on the details of specific cases. It would be implausible to suggest that the collection of pedestrians whose uncoordinated movements accidentally caused a stampede constitutes a responsible agent. Random collections are likely to fall short on each of the three conditions. But arguably, the satisfaction of the conditions is less implausible in the case of a suitably organized collective, such as a commercial corporation or a state. To see whether certain collectives can indeed meet them, we must turn to the idea of group agency.

### 3. Group agency

A group agent, in brief, is an organized collective that constitutes a goal-directed, intentional agent in its own right, over and above its individual members.<sup>14</sup> Examples of group agents are firms and corporations, courts, churches, universities, governments, non-governmental organizations, and even states. We ascribe beliefs and desires, and other intentional attitudes, to those entities and commonly treat them as agents. In fact, the law treats some of those entities as artificial persons with certain obligations and a legal status, distinct from the obligations and the legal status of their members.

Now some people might object that the ascription of agency and attitudes to collectives is just metaphorical. In a much-quoted passage, Anthony Quinton, for instance, writes:

“Groups are said to have beliefs, emotions and attitudes ... But these ways of speaking are plainly metaphorical. To ascribe mental predicates to a group is always an indirect way of ascribing such predicates to its members ... To say that the industrial working class is determined to resist anti-trade union laws is to say that all or most industrial workers are so minded.”<sup>15</sup>

---

<sup>14</sup> On group agency, see, e.g., French (1984), Rovane (1997), Pettit (2001, ch. 5, 2003), Tollefsen (2002, 2015), List and Pettit (2006, 2011), and Tuomela (2013).

<sup>15</sup> See Quinton (1975, p. 17).

On Quinton's view, any talk of group agency or group attitudes is just a shorthand for something that could be equally expressed by reference to the group's members. However, as I will now explain, there are reasons for taking a realist view about group agency, according to which at least some groups can truly qualify as intentional agents. Perhaps Quinton is right that the industrial working class, like the group of pedestrians in my earlier example, is too amorphous to qualify as a genuine agent. But the situation is different in the case of suitably organized collectives, such as firms and corporations as well as the other examples mentioned at the beginning of this section. Such collectives can be shown to meet the defining conditions for intentional agency, and they are in fact routinely treated as agents in the social sciences.

Let me begin with the defining conditions for agency. One might be tempted to define group agency in some kind of *sui generis* way – say, by describing how a collection of individuals must act together to qualify as a group agent – but this approach would be rather *ad hoc*. It would leave open what it is that human beings, group agents, and perhaps other artificial agents (such as robots and AI systems) have in common. A more systematic approach is to introduce general conditions for intentional agency, not conditions tailor-made for the corporate case, and then to ask whether certain groups could meet them. For present purposes, I will adopt a relatively simple definition. I will define an *intentional agent* as an entity, in some environment, that meets three conditions:<sup>16</sup>

**Representation:** The entity has representational states, which encode its “beliefs” about how things are.

**Motivation:** The entity has motivational states, which encode its “goals” or “desires” as to how it would like things to be.

**Action:** The entity has the capacity to interact with its environment, in a way that systematically involves those states: the entity is responsive to information from its environment and “acts” in pursuit of its goals or desires in line with its beliefs.

Clearly, a great variety of entities can meet these conditions, from human beings and non-human animals to certain kinds of robots and AI systems. Each of these can be readily subsumed under the belief-desire-action structure described. Even a system as simple as a thermostat can be interpreted as meeting the conditions: it has representational states, representing the actual temperature of its environment, motivational states, representing a

---

<sup>16</sup> This definition, which expresses a version of the Humean belief-desire conception of agency, draws on List and Pettit (2011, ch. 1) and List (2018a,b). On “belief-desire-intention agency”, see Bratman (1987).



target temperature, and a capacity to interact with its environment, by measuring the temperature and regulating the heating so as to bring the actual temperature in line with the target. It is important to recognize that intentional agents lie on a spectrum, and that their capacities can range from very basic to extremely complex. Human beings, for instance, have a great variety of internal psychological states that go beyond beliefs and desires, and their capacities give them *moral* agency, over and above mere intentional agency; more on this later.

Given all this, the claim that certain organized collectives can meet the defining conditions for intentional agency should be fairly uncontroversial. Indeed, the theory of the firm in economics depicts firms and commercial corporations as belief-desire agents, which rationally maximize expected profits in line with certain beliefs about their commercial environment.<sup>17</sup> Likewise, the so-called “realist” theory of international relations depicts states as belief-desire agents, which rationally pursue their goals and objectives in line with their beliefs about the international world.<sup>18</sup> In game-theoretic models of the Cold War, for instance, the United States and the Soviet Union often featured as “rational players”, with certain strategic objectives and certain beliefs about how their opponent will behave. Their behaviour was then explained in much the same way in which we would ordinarily explain the behaviour of a rational individual, such as a consumer in the marketplace. Of course, those explanations are somewhat simplistic, but they illustrate that taking an “intentional stance” towards certain collectives – explaining their behaviour by viewing them as intentional agents – is a reasonable explanatory strategy.<sup>19</sup>

A critic might say that all of this is consistent with interpreting the ascription of agency to the collectives in question as nothing more than an instrumentally useful move. Treating groups *as if* they were agents might be explanatorily useful, but that does not imply that group agency is a real phenomenon. There is, however, a fairly straightforward philosophy-of-science argument for realism about group agency:<sup>20</sup>

**Premise 1:** The ascription of intentional agency to certain collective entities is indispensable in (social-)scientific explanations of their behaviour.

**Premise 2:** If the ascription of some property to an entity is indispensable in scientific explanations of that entity’s behaviour, then we are – at least *prima facie* – warranted in

---

<sup>17</sup> See, e.g., Mas-Colell, Whinston, and Green (1995).

<sup>18</sup> See Waltz (1979) and Snidal (1985).

<sup>19</sup> On the “intentional stance”, see Dennett (1987). On taking that stance towards collectives, see also Tollefsen (2015).

<sup>20</sup> For previous versions of this indispensability argument, see List (2018a,b). A similar argument is implicit in other works on group agency, including Tollefsen (2002) and List and Pettit (2011).

assuming that the entity really has that property. (The warrant is only *prima facie* because it could still ultimately be defeated by other considerations.)

**Conclusion:** We are – at least *prima facie* – warranted in assuming that the relevant collective entities really are intentional agents.

The argument is certainly valid (the conclusion follows logically from the premises), and I will here assume that a case can be made for each premise. The argument is analogous to arguments for realism about other properties in the sciences. For instance, from the fact that the ascription of certain physical forces to certain entities is indispensable in physical explanations, we often derive a *prima facie* warrant for realism about those physical forces.<sup>21</sup> The upshot is that, insofar as some of our best theories in the social sciences ascribe intentional agency to certain collectives, it is acceptable to treat group agency as a real phenomenon, at least *prima facie*.

Before turning to the question of whether certain groups can also qualify as *moral* agents, I should distinguish group agency from two related phenomena: joint action and collective action. *Joint action* is the coordinated action of several individuals under a structure of shared or common intentions, for instance when several individuals jointly implement a shared plan.<sup>22</sup> Standard examples are carrying a piece of furniture downstairs together or going for a walk together. *Collective action* is the cooperative action of several individuals more generally, whether or not this is based on any shared or common intentions.<sup>23</sup> Successfully cooperating in the provision of a public good, despite incentives to free-ride, would be an example. Neither of these phenomena involves a group agent in its own right. To make sense of joint or collective action, we need not view the group itself as an intentional agent. The *loci* of agency are still the individuals.<sup>24</sup>

#### 4. Can groups be fit to be held responsible?

We have seen that there is a case to be made for the claim that certain organized collectives qualify as intentional agents in their own right. This is so at least for those collectives for which the ascription of agency is explanatorily indispensable, such as firms, corporations, and states, among others. By contrast, unorganized collectives such as the group of pedestrians in the

---

<sup>21</sup> This derivation implicitly relies on what is sometimes called “the naturalistic ontological attitude”, the view that science is our best guide to ontological questions – questions about which entities and properties are real. See, e.g., Quine (1977) and Fine (1984).

<sup>22</sup> On this phenomenon, see, e.g., Bratman (1999, 2014), Gilbert (1989), and Tuomela (2007).

<sup>23</sup> On the problem of collective action, see Olson (1965) and Ostrom (1990).

<sup>24</sup> Gilbert (1989) introduces the idea of a plural subject in her account of joint action. My view is that there are many cases of joint action where we do not need to postulate a group agent.

example of the stampede would not warrant such an ascription. Now, if we restrict our attention to those groups that constitute intentional agents, the possibility that they might also qualify as moral agents becomes a live option. Although some of the most familiar cases of non-human agency, such as animal agency, fall short of *moral* agency, I have already pointed out that there is no conceptual reason why some non-human entities could not also qualify as moral agents. What distinguishes a moral agent from a mere intentional one is the capacity to make normative judgments about its choices – judgments about what is right and wrong, permissible and impermissible – and to be guided by those judgments. There is no reason why corporations and other organized collectives could not have institutional procedures and mechanisms that enable them to make such judgments and to let these judgments guide their actions. Whether an organized collective is like this depends entirely on how it is designed and structured.<sup>25</sup>

What about the other conditions for fitness to be held responsible? Recall that moral agency alone is not enough. The entity in question must also have access to the information needed for a normative assessment of its choices, and it must have the freedom and control required for making those choices. The access-to-information condition is relatively unproblematic. Just as individuals with the right cognitive capacities are capable of accessing and processing relevant information about their choices, provided their informational environment is not too constrained, so are suitably organized group agents. For instance, a commercial corporation is no less well placed to access and process information about the likely consequences and negative externalities of its actions than any individual human being would be.

The condition of free choice is trickier. Can group agents genuinely have freedom and control in making their choices? I cannot do justice to this question here, but I want to point out that, at least given the possibility of intentional group agency, a positive answer is not too far-fetched. I take freedom and control to require, at a minimum, two things. First, the relevant agent must be able to choose between alternative possibilities. And second, when the agent makes a choice, the agent's intentional state must be the difference-maker of that choice: whatever the agent does must not be some sub-agential fluke (such as an uncontrollable reflex or spasm). Can group agents meet these requirements?<sup>26</sup>

As far as the requirement of alternative possibilities is concerned, I see no reason why they could not. Group agents can certainly be faced with choices between alternative courses of action. Their choices need not be determined in a way that would have made it impossible for

---

<sup>25</sup> For recent works on corporate moral agency, see, e.g., Björnsson and Hess (2017) and Pasternak (2017).

<sup>26</sup> I have previously discussed these questions in List and Pettit (2011, ch. 7) and in List (2019).

them to act any differently from how they did. When a corporation makes a risky decision that leads to some disaster, it may be perfectly reasonable to say that this corporation could have acted otherwise: the alternative, less risky course of action may have been entirely open to the corporation; perhaps its executive board even considered it, but dismissed it.

Second, group agents can also be said to have control over what they do. Just as – in the case of individual agency – the difference-making cause of an individual’s action is arguably that individual’s intentional state, so the difference-making cause of a corporate action can plausibly be the group agent’s intentional state, for instance as generated through its corporate decision-making structure. Although the claim that group agents exercise causal control over their actions is vulnerable to the same kinds of philosophical challenges that arise for mental causation in the case of individuals, the available philosophical responses to those challenges in the individual case arguably carry over to the corporate one.<sup>27</sup> I cannot do more than flag these issues here. At any rate, I want to note that there is, by now, a large philosophical literature that assumes that these obstacles can be overcome and that recognizes the possibility of corporate moral responsibility. Peter French’s work on corporate responsibility is a *locus classicus*.<sup>28</sup>

In sum, there are theoretical resources available for arguing that some group agents are fit to be held responsible. The recognition of corporate responsibility, in turn, may be combined with an appropriate penal code at the corporate level. The punishment for corporate wrongdoing may include, for instance, fines and other sanctions targeting the assets of the relevant corporate entities, as well as restrictions on their operating permissions. In the same way in which the punishment for an individual’s criminal wrong-doing may be the (perhaps temporary) removal of certain liberties, so the punishment for a corporate crime may be a certain kind of operating restriction or ban for the group agent.

The present considerations suggest that society might protect itself against the risk of corporate responsibility gaps by legally requiring that group agents operating in high-stakes settings should be organized so as to meet the conditions for responsible moral agency. According to this proposal, for a powerful group agent to receive its corporate charter or operating licence, it would have to demonstrate that it has procedures in place to ensure its

---

<sup>27</sup> My reasoning relies on the account of mental causation in List and Menzies (2009). The solution to the corporate problem of mental causation suggested in List and Pettit (2011, ch. 7) is subtly different. On agential control as difference-making, see also Himmelreich (2015).

<sup>28</sup> See French (1984) and the other references in footnotes 1 and 2.

fitness to be held responsible for its actions.<sup>29</sup> In the same way in which universities and other organizations receive their accreditation only if they comply with certain legal requirements, so powerful group agents in general would be allowed to operate only if they are compliant with the conditions for fitness to be held responsible. The details of how this could work are complicated, but the present proposal would offer one way to avoid the kind of situation that arose in the example of the ferry disaster. Indeed, we can now see that *if* the ferry company truly fell short of the conditions for fitness to be held responsible, then – according to the proposal just sketched – the company should never have been permitted to run a ferry service.

### **5. How does group responsibility relate to individual responsibility?**

The claim that there can be group responsibility over and above individual responsibility – and sometimes even in the absence of individual responsibility – invites a natural objection. Group agents are collections of individuals, and they cannot act without their members' contribution. Anything a group agent does depends on what its members do, under the group's organizational structure. Insofar as those members are themselves fit to be held responsible, shouldn't group responsibility somehow translate into member responsibility? In slightly more metaphysical terms, since all group-level facts (facts about what a group agent does) presumably "supervene" on individual-level facts (facts about what the members do), so the objection goes, it must be possible to "reduce" group-level responsibility to individual-level responsibility.<sup>30</sup>

The response is that this objection rests on a misconception. It overlooks the fact that, in real cases of group agency, unlike in cases of mere joint or collective action, the group itself is a *locus* of agency over and above its members. This means that the agent of the group's actions is the group itself, understood in "corporate" terms, not its individual members, understood in "distributive" terms. While the group's actions are *causally* and even *ontologically* dependent on the members' contributions, it does not follow that those actions can also be *agentially* attributed to them. In a case of genuine group agency, the actions of the group agent are not correctly classified as actions of its members.

For a structural analogy, consider the relationship between the actions of an individual person and the biological processes within the person's organism. Even if we accept the standard view that the person's actions are causally and ontologically dependent on the

---

<sup>29</sup> For a related proposal, involving a "developmental rationale" for holding group agents responsible, see List and Pettit (2011, ch. 7).

<sup>30</sup> This objection is obviously motivated by a broader background assumption of methodological individualism. For a discussion of methodological individualism, see, e.g., Heath (2015).

underlying bodily processes, it would be a mistake to treat any of the constituent body parts – say, the visual cortex, the cerebellum, or the muscles – as “agentially responsible” for those actions. The visual cortex, the cerebellum, and the muscles all contribute to the person’s performance of his or her actions, but these body parts do not qualify as *agents* of those actions. The *locus* of agency is the person as a whole, not any components of the underlying organism. Of course, in this example, the relevant body parts do not qualify as intentional agents themselves, and so it would never occur to us to ascribe any agential responsibility to them. But even if the function of one of these body parts – say that of the cerebellum – were performed by a hypothetical system that did qualify as an agent, it would still not follow that the *locus* of agency for the person’s overall actions was anything other than the person as a whole. Similarly, if we take the notion of group agency seriously, we must distinguish between the actions of a group agent and the actions of its members, and we must not confuse the former with the latter. Just as an individual person’s actions are “located” at the systemic level of the person as a whole, not the level of any body parts, so the actions of a group agent are “located” at the systemic level of the group as a whole, not the level of its individual members.

A version of this point was already recognized by Virginia Held in a 1970 article, in which she discussed the crucial difference between organized groups and mere random collections:

“[I]f organized group *G* is morally responsible for the failure to do *A* it does not follow that member *M* of *G* is morally responsible for the failure to do *A*. If a random collection *R* can be represented as a set equivalent, say, to *M* & *N* & *Q*, then, if *R* is morally responsible, we would seem to be able to conclude that *M* is morally responsible & *N* is morally responsible & *Q* is morally responsible. On the other hand, if these same members formed an organized group, the group could not be adequately represented as equivalent simply to *M* & *N* & *Q*, because its depiction would have to include the decision method by which the members act as a group. And the distribution of moral responsibility over such a combination would not seem plausible.”<sup>31</sup>

That said, depending on the details of the case, some individuals may bear individual responsibility for the roles they have played in relation to the group agent’s conduct. They may do so to the extent that they have knowingly and voluntarily played some of the following roles:

- *enacting roles*, for example as managers, officials, or representatives of the group agent,

---

<sup>31</sup> See Held (1970, p. 480).

- *authorizing roles*, for example as directors, board members, owners, share-holders, or even regulators, or
- *organizational-design roles*, for example as founders, policy makers, or institutional designers.<sup>32</sup>

How exactly the occupants of these roles are responsible for what the group agent does is a contingent matter that may differ from case to case. But as I have already emphasized, there is no reason to assume that the sum-total of their individual responsibility will always match the responsibility of the group agent as a whole. In short, recognizing individual responsibility in the context of a group agent is no substitute for recognizing group responsibility.

It is worth acknowledging two further complications that arise when we try to “distribute” group responsibility to the underlying individuals. The first is this. Suppose we took the view that individuals are responsible for the difference they have knowingly and willingly made to the group’s actions: if they have made a big difference, they would bear significant responsibility, whereas if they have made only a small difference or no difference, they would bear very little responsibility or none. The problem with this view is that, even if individuals have knowingly and willingly contributed to the group’s actions and have wholeheartedly endorsed them, they might still come out as bearing no responsibility at all. We can easily imagine cases in which the group collectively takes a momentous decision, and yet no-one is individually pivotal. The easiest illustration is given by a democratically organized group that makes its decisions by majority rule. If a decision is reached by a margin of victory of more than one or two votes, then no individual voter is pivotal: no-one can say that if he or she had voted differently, this would have changed the outcome (by flipping the majority). In such a case, no individual qualifies as a difference-maker of the collective outcome. The lesson is that, even in cases in which it seems appropriate to assign responsibility to the group members on the basis of their knowing and willing support for the group’s decision, a simple difference-making criterion would not vindicate this assignment of responsibility. Rather, a more subtle criterion would be needed. One example would be the so-called “INUS” or “NESS” criterion. According to it, a minimal condition for an individual’s action to count as a significant contributor to a collective outcome would be that this individual’s contribution (which we here assume is made voluntarily and knowingly) is “an insufficient but necessary part of a condition

---

<sup>32</sup> Cf. List and Pettit (2011, ch. 7).

which is itself unnecessary but sufficient for the result” or, more generally, “a necessary element of a [causally] sufficient set”.<sup>33</sup> Of course, more could be said about such a criterion.<sup>34</sup>

The second complication is the following. Suppose we thought that at least in democratically organized groups, under background conditions of freedom and voluntariness, it should not be difficult to “distribute” group responsibility to the individual members. The reason is that, for any democratic group decision, we would expect that some individuals (normally a majority) will have voted in support of that decision. Those individuals could then qualify as bearing some responsibility, regardless of whether they were individually pivotal. In particular, we might take the view that a reasonable condition for bearing some individual responsibility for a collective decision is having voluntarily and knowingly voted in support of it. The problem with this view, however, is that it can still lead to responsibility gaps. A much-cited example is that of a collegial court making a decision in a breach-of-contract case.<sup>35</sup> Consider the following scenario. A three-member court has to decide whether a defendant is liable for breach of contract. The relevant legal doctrine requires the court to make this decision on the basis of three background questions: first, did the defendant do the action in question; second, did the action go against the relevant contract; and third, was the contract valid in the first place? The defendant is liable *if and only if* the answer to all three questions is “yes”. If, on the other hand, the defendant did not do the action, or the action was permitted by the contract, or the contract was invalid, then there is no liability. Now we can easily imagine a situation in which each member of the court gives a positive answer to two of the three questions, but not to the third, with different judges disagreeing about which question is to be answered in the negative. The pattern of individual judgments is shown in Table 1.

	Action done?	Prohibited by the contract?	Contract valid?	Liable?
Judge 1	Yes	Yes	No	No
Judge 2	Yes	No	Yes	No
Judge 3	No	Yes	Yes	No
Majority	Yes	Yes	Yes	No

**Table 1: A three-member court**

<sup>33</sup> See Mackie (1965) and Wright (1985), respectively.

<sup>34</sup> See, e.g., Braham and van Hees (2012).

<sup>35</sup> The example goes back to Kornhauser and Sager (1986). For more recent discussions, see List and Pettit (2011).



In this case, the majority answer to all three questions is “yes”, and yet no individual judge thinks that the defendant is liable. If – hypothetically – the court were to use the so-called “premise-based procedure”, it would reach its overall verdict as follows. It would take majority votes on the three premises for liability (was the action done, was it against the contract, was the contract valid), and it would then derive its overall verdict on the case from its collective judgments on those premises. Under this arrangement, there would be no need to take a vote on the issue of liability itself; the verdict would be derived via a form of “collectivized reasoning”, from democratically agreed premises. Given the individual judgments shown in Table 1, the premise-based procedure would support a “liable” verdict, despite the unanimous individual view that the defendant is not liable. This example illustrates that a collective decision may have been reached through a perfectly reasonable democratic procedure, such as the premise-based procedure, and yet no individual may have personally supported the ultimate outcome – here a “liable” verdict. If having voted for the outcome were a necessary condition for bearing some responsibility for it, then none of the individual judges would bear any responsibility in the present case.<sup>36</sup>

Of course, one might object that the premise-based procedure is not the best decision procedure here, but the example generalizes. It can be shown that any democratically organized group that seeks to arrive at consistent collective judgments on several logically related issues cannot generally follow the majority judgment on every issue.<sup>37</sup> Issue-by-issue majority voting would run the risk of producing an inconsistent pattern of judgments, as illustrated by Table 1. If the group is to function as an instrumentally rational group agent, with consistent collective views, then it must sometimes (in some logically possible situations) deviate from the views of a majority of its members. And so, we cannot presume that every group decision will be supported by a majority of group members, to whom the relevant responsibility can then be assigned. Although the details are complicated, I hope the basic point is clear enough. For a further analysis of the present problem, I refer the reader to the literature on judgment aggregation.<sup>38</sup> The bottom line is that there is no general scheme available by which group responsibility can always be distributed to the group’s members: group responsibility is not generally reducible to individual responsibility.

---

<sup>36</sup> The analysis of this case would be more complicated if one took the view that, in cases such as that of the court, voting for an outcome is not a necessary condition for bearing some responsibility for it. One might argue, for instance, that strategically abstaining from voting on a controversial matter does not absolve one of responsibility.

<sup>37</sup> For a review of relevant results and further references, see List and Pettit (2011, ch. 2).

<sup>38</sup> For a review and further references, see again List and Pettit (2011).

## 6. Can groups also have rights of their own?

I have described and defended the view that certain collective entities can bear responsibility in their own right, over and above their individual members. In closing, I would like to ask whether accepting this view commits us to accepting that those groups should be given rights as well. If groups can be duty-bearers, can they also be right-holders?

I will argue that, in one sense, the answer is quite unproblematically positive, but in another, it is clearly negative. To distinguish those senses, I must begin with two clarifications about the notion of a “right”. First of all, the term “right” can refer either to a legal right or to a moral one. Legal rights are those rights that agents have under the legal systems they are subject to. For instance, I have a legal right to vote in my country of citizenship, but I don’t have a legal right to vote in the country I visit as a tourist. Moral rights, by contrast, are rights that agents have from the perspective of morality. We arguably all have the moral right to certain basic liberties and respectful treatment, whether or not we are lucky enough to live in a society in which those rights are legally entrenched. The second clarification is this. When we say that an agent has a particular right, where this entails certain powers or permissions and perhaps duties on the part of others, then the agent might have that right either non-derivatively or derivatively. A non-derivative right is one the agent has by virtue of some underlying moral status. The right might be grounded, for instance, in the fact that the agent matters intrinsically: he or she is an object of moral concern. The natural rights that human beings are thought to have by virtue of their humanity, or by virtue of their status as moral persons, would be examples of non-derivative rights. A derivative right, by contrast, is one the agent doesn’t have by virtue of any underlying moral status; rather, the right derives from something else. A derivative right might be grounded in its instrumental usefulness for promoting something else, which matters intrinsically. For example, a medical doctor might have certain derivative rights in his or her professional role, where those rights are justified by virtue of the fact that they are part of a normative system that serves the interests of society as a whole.

Given these clarifications, I can now explain in what sense group agents can have rights, and in what sense they cannot. My claim is that organized collectives can have derivative rights (in principle, moral and legal), but they cannot have any non-derivative ones. Let me begin with the positive part of this claim. Consider the functions that collective entities perform in the social world. The state provides all sorts of services; universities educate students and contribute to our knowledge of the world; medical institutions look after people’s health; and firms and financial organizations – notwithstanding their downsides – are central to the

economy. It is hard to deny that, given the way society is organized, those entities could not perform those functions unless they had certain rights, such as the right to own or rent property, the right to employ staff, the right to enter contracts and to engage in litigation when necessary, and so on. In recognition of these rights, the law treats the relevant collectives as legal persons and thereby considers them fit to be right-holders and duty-bearers. Of course, the rights in question are all derivative. They are not grounded in any special moral status of those collectives, but they are justified simply by virtue of the fact that, by and large, the normative system of which they are a part serves society's interests. To be sure, there may be plenty of scope for improving that system. The philosophical point is just that, to the extent that giving certain rights to some collectives promotes the interests of the human members of society, group agents can unproblematically have those rights. That is the sense in which organized collectives can have derivative rights.

Let me now turn to the negative part of my claim. I want to suggest that, unlike individual human beings, organized collectives cannot have any non-derivative rights. The reason is that collective entities lack the relevant moral status: they do not matter intrinsically. Their significance is only instrumental. In this respect, group agents differ fundamentally from individual human beings, and arguably also from some non-human animals such as the great apes and other mammals, all of whom – we may reasonably assume – matter intrinsically. But while the differential categorization of group agents, as compared to individual humans and even animals, may seem intuitive, we must show that it is philosophically defensible.<sup>39</sup> At first, one might think that if organized collectives not only meet the conditions for intentional agency, but also qualify as full-blown moral agents, then it is hard to deny them the kinds of rights that individual moral agents have. If we accept collectivism with regard to responsibility, how can we defend individualism when it comes to rights? Wouldn't a robust defence of group agency vindicate, for instance, the US Supreme Court's decision in *Citizens United v. Federal Election Commission* to extend certain free-speech rights to corporations and other organizations, where those rights can then be exercised via corporate activities and monetary spending?

To show that only individual human beings, and perhaps certain non-human animals, but not group agents can have non-derivative rights, we must show that the latter lack the kind of moral status that would ground such rights. To do so, we must appeal to an appropriate background criterion as to what the basis of an entity's moral status is and show that group

---

<sup>39</sup> On the debate about corporate rights and the moral status of collectives, see also Pasternak (2017) and Silver (2018).

agents fail to meet that criterion. There are a number of possible criteria on offer, which approach the present issue in different ways.<sup>40</sup>

One criterion for being an object of moral concern might be membership of the human species. This criterion is clearly satisfied by individual humans and violated by groups. But the criterion is somewhat *ad hoc* and has the unfortunate consequence of “throw[ing] animals under the bus”, as Will Kymlicka puts it.<sup>41</sup> A less restrictive criterion for being an object of moral concern would be the possession of intentional agency. This criterion is satisfied not only by human beings, but also by many non-human animals as well as by group agents. Even robots and AI systems could easily satisfy it. The worry here is that the criterion is too permissive. On this criterion, we would have to extend our moral concern to entities ranging from British Petroleum to self-driving cars. A related but more restrictive criterion would be the possession of moral, not just intentional, agency. But, as we have already seen, group agents can satisfy that criterion, whereas, for instance, the great apes, who arguably matter intrinsically, would lose out.

Can we come up with a better criterion? I think a plausible necessary (though perhaps not sufficient) condition for being an object of moral concern is the possession of phenomenal consciousness or, more broadly, the potential for phenomenal consciousness. Having phenomenal consciousness means being an entity that is a *locus* of subjective experiences: there is something it is like to be that entity, in Thomas Nagel’s words.<sup>42</sup> This criterion would support our assumption that human beings are objects of moral concern and can therefore have non-derivative rights,<sup>43</sup> and it would also allow us to make a similar case for animals. The claim that some non-human animals have subjective experiences, including experiences of pleasure and pain, is fairly uncontroversial. So, what about group agents?

There are strong reasons to believe that group agents do not have anything like phenomenal consciousness or the potential for it. There is nothing it is like to be a group agent.<sup>44</sup> At most, there is something it is like to be a member of a group agent; but that is because the members are *loci* of subjective experiences as individuals; the group as a whole isn’t. In particular, group agents lack the kinds of properties which, according to contemporary neuroscience, are

---

<sup>40</sup> On the grounds of moral status, see Jaworska and Tannenbaum (2018).

<sup>41</sup> See Kymlicka (2017).

<sup>42</sup> See Nagel (1974). I have previously discussed this criterion for having non-derivative rights in List (2018a,b). On phenomenal consciousness, see also Chalmers (1996).

<sup>43</sup> There are some complications here, which I cannot go into. One might ask whether babies or comatose patients would satisfy the present criterion. I think the answer is positive, provided we adopt a sufficiently broad interpretation of “potential” for phenomenal consciousness.

<sup>44</sup> See List (2018); cf. Schwitzgebel (2015).

associated with phenomenal consciousness.<sup>45</sup> Although I cannot go into the details here, phenomenal consciousness appears to occur only in information-processing systems in which there are massive internal high-bandwidth feedback mechanisms, as exemplified by the human cortex. Arguably, this sort of cortex-like informational integration is absent from organized collectives, even if we think of them as socially implemented intelligent systems.<sup>46</sup>

For there to be a phenomenally conscious group agent, we would have to imagine the kind of scenario proposed by Ned Block in his “China brain” thought experiment.<sup>47</sup> In this scenario, each member of a very sizeable population (say, that of China) is tasked with playing the role of one neuron in a biological brain, and the overall neural network is then established via internet connections. Conceivably, this kind of setup might replicate cortex-like informational integration in a social structure, so that the resulting group agent might be a candidate for phenomenal consciousness. But this is evidently a science-fiction scenario, and existing group agents, or plausible future ones, are not like this at all. In short, if the present empirical claims about the physical correlates of phenomenal consciousness are correct, and phenomenal consciousness (or the potential for it) is indeed a necessary condition for being an object of moral concern, then clearly group agents do not meet that condition, and consequently they cannot have any non-derivative rights.

The bottom line, then, is that collective entities may unproblematically have certain *derivative* rights as well as legal personhood, but no *non-derivative* rights or a full-blown moral status. Furthermore, the criterion for adjudicating whether to give any derivative rights to collective entities is whether this serves the interests of those who do qualify as objects of moral concern, such as human beings and perhaps non-human animals. In the case of *Citizens United*, the only way in which extending free-speech rights to corporations could be justified would be by showing that this does indeed serve the interests of human society – a claim that is at least contentious. A mere appeal to the legal personhood of corporations would be insufficient as a basis for extending non-derivative rights to groups.

In conclusion, there is no contradiction in recognizing certain collective entities as moral agents who are fit to be held responsible while at the same time denying that they have the kinds of rights that individuals have. Collectivism about responsibility is perfectly compatible with individualism about rights.

---

<sup>45</sup> This claim is supported, for instance, by integrated information theory (Tononi and Koch 2015), but other theories of the physical correlates of consciousness (such as the neural synchronization theory of Crick and Koch 1990) would support it too.

<sup>46</sup> See List (2018), but see Schwitzgebel (2015).

<sup>47</sup> See Block (1980).

## References

- Alvarez, M. 2009. Actions, Thought-Experiments and the “Principle of Alternate Possibilities”. *Australasian Journal of Philosophy* 87(1): 61–81.
- Björnsson, G., and K. Hess. 2017. Corporate Crocodile Tears? On the Reactive Attitudes of Corporate Agents. *Philosophy and Phenomenological Research* 94(2): 273–298.
- Block, N. 1980. Troubles with Functionalism? In N. Block, ed., *Readings in Philosophy of Psychology, Vol. 1*: 268–306. London: Methuen.
- Braham, M., and M. van Hees. 2011. Responsibility voids. *The Philosophical Quarterly* 61(242): 6–15.
- Braham, M., and M. van Hees. 2012. An Anatomy of Moral Responsibility. *Mind* 121(483): 601–634.
- Braham, M., and M. van Hees. 2018. Voids or Fragmentation: Moral Responsibility For Collective Outcomes. *The Economic Journal* 128(612): F95–F113.
- Bratman, M. E. 1987. *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Bratman, M. E. 1999. *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge: Cambridge University Press.
- Bratman, M. E. 2014. *Shared Agency: A Planning Theory of Acting Together*. Oxford: Oxford University Press.
- Buell, S. W. 2018. The Responsibility Gap in Corporate Crime. *Criminal Law and Philosophy* 12(3): 471–491.
- Chalmers, D. J. 1996. *The Conscious Mind*. New York: Oxford University Press.
- Collins, S. 2017. Filling Collective Duty Gaps. *Journal of Philosophy* 114(11): 573–591.
- Collins, S. 2018. Collective Responsibility Gaps. *Journal of Business Ethics*. Online.
- Colvin, E. 1995. Corporate Personality and Criminal Liability. *Criminal Law Forum* 6: 3–44.
- Copp, D. 2006. On the Agency of Certain Collective Entities: An Argument from “Normative Autonomy”. *Midwest Studies in Philosophy* 30(1): 194–221.
- Crick, F., and C. Koch. 1990. Towards a Neurobiological Theory of Consciousness. *Seminars in the Neurosciences* 2: 263–275.
- Dennett, D. 1987. *The Intentional Stance*. Cambridge, MA: MIT Press.
- Duijf, H. 2018a. Responsibility Voids and Cooperation. *Philosophy of the Social Sciences* 48(4): 434–460.

- Duijf, H. 2018b. *Let's do it! Collective responsibility, joint action, and participation*. PhD thesis, University of Utrecht <<https://dspace.library.uu.nl/handle/1874/364155/>>.
- Erskine, T. 2001. Assigning Responsibilities to Institutional Moral Agents: The Case of States and Quasi-States. *Ethics & International Affairs* 15(2): 67–85.
- Feinberg, J. 1968. Collective Responsibility. *Journal of Philosophy* 65(21): 674–688.
- Fine, A. 1984. The Natural Ontological Attitude. In J. Leplin, ed., *Scientific Realism*, 83–107. Berkeley: University of California Press.
- Frankfurt, H. 1969. Alternate Possibilities and Moral Responsibility. *Journal of Philosophy* 66(23): 829–839.
- French, P. A. 1984. *Collective and Corporate Responsibility*. New York: Columbia University Press.
- Gilbert, M. 1989. *On Social Facts*. New York: Routledge.
- Heath, J. 2015. Methodological Individualism. In E. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Spring 2015 Edition) <<https://plato.stanford.edu/archives/spr2015/entries/methodological-individualism/>>.
- Held, V. 1970. Can a Random Collection of Individuals be Morally Responsible? *Journal of Philosophy* 67(14): 471–481.
- Himmelreich, J. 2015. *Agency as difference-making: Causal Foundations of Moral Responsibility*. PhD thesis, London School of Economics <<http://etheses.lse.ac.uk/3277/>>.
- Jaworska, A., and J. Tannenbaum. 2018. The Grounds of Moral Status. In E. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Spring 2018 Edition) <<https://plato.stanford.edu/archives/spr2018/entries/grounds-moral-status/>>.
- Kornhauser, L. A., and L. G. Sager. 1986. Unpacking the Court. *Yale Law Journal*, 96(1): 82–117.
- Kymlicka, W. 2017. Human rights without human supremacism. *Canadian Journal of Philosophy*. Online.
- List, C. 2018a. What is it like to be a group agent? *Noûs* 52(2): 295–319.
- List, C. 2018b. Group Agency and Artificial Intelligence. Working paper, LSE.
- List, C. 2019. *Why Free Will is Real*. Cambridge, MA: Harvard University Press.
- List, C., and P. Menzies. 2009. Non-Reductive Physicalism and the Limits of the Exclusion Principle. *Journal of Philosophy* 106(9): 475–502.
- List, C., and P. Pettit. 2006. Group Agency and Supervenience. *Southern Journal of Philosophy* 44(S1): 85–105.

- List, C., and P. Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- Mackie, J. L. 1965. Causes and Conditions. *American Philosophical Quarterly* 2(4): 245–264.
- Mas-Colell, A., M. D. Whinston, J. R. Green. 1995. *Microeconomic Theory*. Oxford: Oxford University Press.
- May, L. 1987. *The Morality of Groups: Collective Responsibility, Group-Based Harm, and Corporate Rights*. Notre Dame: University of Notre Dame Press.
- Nagel, T. 1974. What Is It Like to Be a Bat? *Philosophical Review* 83(4): 435–450.
- Olson, M. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, MA: Harvard University Press.
- Ostrom, E. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.
- Pasternak, A. 2013. Limiting States' Corporate Responsibility. *Journal of Political Philosophy* 21(4): 361–381.
- Pasternak, A. 2017. From Corporate Moral Agency to Corporate Moral Rights. *The Law & Ethics of Human Rights* 11(1): 135–159.
- Pettit, P. 2001. *A Theory of Freedom: From the Psychology to the Politics of Agency*. Cambridge and New York: Polity and Oxford University Press.
- Pettit, P. 2003. Groups with Minds of their Own. In F. Schmitt, ed., *Socializing Metaphysics*: 167–193. New York: Rowan and Littlefield.
- Pettit, P. 2007. Responsibility Incorporated. *Ethics* 117(2): 171–201.
- Portmore, D. W. 2018. Maximalism and Moral Harmony. *Philosophy and Phenomenological Research* 96(2): 318–341.
- Quine, W. V. 1977. *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- Quinton, A. 1975. The presidential address: Social objects. *Proceedings of the Aristotelian Society* 76: 1–27+viii.
- Rovane, C. 1997. *The Bounds of Agency: An Essay in Revisionary Metaphysics*. Princeton: Princeton University Press.
- Schwitzgebel, E. 2015. If Materialism is True, the United States is Probably Conscious. *Philosophical Studies* 172(7): 1697–1721.
- Silver, K. 2018. Can a Corporation be Worthy of Moral Consideration? *Journal of Business Ethics*. Online.



- Smiley, M. 2017. Collective Responsibility. In E. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition) <<https://plato.stanford.edu/archives/sum2017/entries/collective-responsibility/>>.
- Snidal, D. 1985. The Game Theory of International Politics. *World Politics* 38(1): 25–57.
- Tollefsen, D. P. 2002. Collective Intentionality and the Social Sciences. *Philosophy of the Social Sciences* 32(1): 25–50.
- Tollefsen, D. P. 2015. *Groups as Agents*. Cambridge: Polity Press.
- Tononi, G., and C. Koch. 2015. Consciousness: Here, There and Everywhere? *Philosophical Transactions of the Royal Society B* 370.
- Tuomela, R. 2007. *The Philosophy of Sociality: The Shared Point of View*. New York: Oxford University Press.
- Tuomela, R. 2013. *Social Ontology: Collective Intentionality and Group Agents*. Oxford: Oxford University Press.
- Waltz, K. 1979. *Theory of International Politics*. Long Grove, IL: Waveland Press.
- Wright, R. W. 1985. Causation in Tort Law. *California Law Review* 73(6): 1735–1825.