

Collective Responsibility Gaps

Stephanie Collins

This is the final draft of a paper whose definitive version will be published in

Journal of Business Ethics.

Abstract

Which kinds of responsibility can we attribute to which kinds of collective, and why? In contrast, which kinds of collective responsibility can we not attribute—which kinds are ‘gappy’? This paper provides a framework for answering these questions. It begins by distinguishing between three kinds of collective (diffuse, teleological, and agential) and three kinds of responsibility (causal, moral, and prospective). It then explains how gaps—i.e., cases where we cannot attribute the responsibility we might want to—appear to arise within each type of collective responsibility. It argues some of these gaps do not exist on closer inspection, at least for some collectives and some of the time.

Introduction

Collective responsibility is full of holes. By this, I do not mean that collective responsibility is incoherent or non-existent. I mean there are many situations in which we are unable to attribute all the responsibility we might pre-reflectively want to attribute to collectives, such as business corporations and states. The diversity of collective responsibility gaps, and the connections and disconnections between them, has not yet been fully appreciated. The aim of

this paper is to examine the breadth of gaps and the challenges they pose for asserting that collectives have responsibility.

A failure to distinguish between types of collective responsibility gaps can lead to confusion. Thus, the main message of this paper is a call for clarity when we discuss collective responsibility gaps. Secondly, I will argue that some types of collective responsibility gap do not exist at all, while others are quite common. The paper also has a third message: that a failure to distinguish different types of collective responsibility gaps might make us more sceptical of collective responsibility than we ought to be. That is, when we notice the existence of one type of collective responsibility gap, we might infer that other types of collective responsibility gap also exist. That inference would be a mistake.

To catalogue collective responsibility gaps, we must first understand the breadth of phenomena captured by the phrase ‘collective responsibility.’ To this end, Section 1 typologises three kinds of collective and three kinds of responsibility. Section 2 motivates the idea each kind of responsibility faces gaps in collective contexts, with reference to the Great Pacific Garbage Patch (GPGP). To remedy this dire prognosis, Section 3 examines the gaps described in Section 2 in detail and argues some can be filled—and some can be filled for some kinds of collective, but not others. The result is that there are many kinds of collective responsibility of gaps, but not as many as meets the eye.

1. Types of Collective, Types of Responsibility

In philosophy, legal theory, and political theory, it has long been acknowledged that ‘responsibility’ is multiply ambiguous (Hart 1968; Feinberg 1970, 130-9; Scanlon 1998, 290-4; Miller 2007, 81-91). More recently in philosophy, it has also been acknowledged that

‘collective’ is multiply ambiguous (Isaacs 2011, 24-27). To illustrate, consider the following three responsibility attributions:

- A. ‘19th century industrialists affected global warming.’
- B. ‘The fossil fuel lobby is to blame for global warming.’
- C. ‘BP has an obligation to clean up its oil spill.’

Each refers to a different collective-type and responsibility-type.

Statement A concerns a *diffuse* collective and *causal* responsibility. Diffuse collectives are groups of agents who are not united either by acting responsively to one another as they work towards a common goal, or by acting under a collective-level decision-making procedure. These include ‘humanity’, ‘affluent consumers’, ‘tech companies’, and ‘developed states.’ The agents in a diffuse collective can be utterly random (for example, ‘me, BP, and Botswana’), though more commonly we’ll be concerned with agents that have a socially significant property in common. In this paper, I am concerned with diffuse collectives whose members include *agential collectives* (defined shortly). Regarding responsibility, Statement A attributes *causal* responsibility—impact or influence without, necessarily, praiseworthiness or blameworthiness. Causal responsibility need not involve actions: a storm can be causally responsible for a fallen power line.

Statement B concerns a *teleological* collective and *moral* responsibility. The collective is teleological in that its members (1) act responsively to one another (insofar as they encounter one another) as they work toward a common goal (oil-friendly public policy), but (2) lack clear procedures for forming decisions, intentions, beliefs, and desires that are attributable to the collective *as such*. When we say ‘the fossil fuel lobby wants or believes X’, we do not mean that this is an official corporate position of some particular entity that can

outlive whoever happens to currently make up the fossil fuel lobby. Instead, we simply mean that all or most fossil fuel lobbyists currently want or believe X.¹

Nonetheless, members of teleological collectives are interlinked: when I say they are disposed to be ‘mutually responsive,’ I mean they are disposed to reinforce, predict, and rely upon each other’s pursuit of the goal that they share. We can analyse this further in different ways. For example, the lobbyists are an example of Michael Bratman’s (1992, 328) “shared cooperative activity,” in which “each participating agent [knows] that the other is attempting to be similarly responsive”; Christopher Kutz’s (2000) acting together through “participatory intentions” to do one’s part in an action or project together with others; and Raimo Tuomela’s (2006, 38) joint intentional action, which “amounts to the group members’ jointly intending X and jointly acting on the basis of their joint intention.” Each of these specifications is more demanding than the minimal notion I will work with (defined by characteristics (1) and (2) above).

With regards to the responsibility evoked in Statement B, the lobbyists’ responsibility concerns something that they were causally responsible for *and* are liable to praise or blame for. This makes the responsibility moral. By ‘is liable to praise or blame,’ I mean ‘deserves or merits praise/blame’, or ‘is a fitting object of praise/blame’. The liability here is backwards-looking and reactive: it responds to what the entity has done in the past, not to what they might do in the future (Strawson 1960).

Finally, Statement C concerns an *agential* collective and *prospective* responsibility. BP has a well-defined collective-level decision-making procedure, the use of which enables us to attribute some decisions (and intentions, beliefs, and desires) to BP as such. An agential collective has a group-level decision-making procedure—a process that takes in reasons and produces aims and instructions—that is operationally distinct from the procedures held

respectively by its members. It is operationally distinct in that its reasons and beliefs (inputs) may not be identical to the reasons and beliefs of any members; its method for processing inputs (e.g., majoritarian, committee-based, or dictatorship rules) is different from the method of any one member when deciding for herself; and the decisions it produces are not the straightforward conjunction of individuals' decisions (List and Pettit 2011). By creating its own bundle of decisions, intentions, beliefs, desires, and so on, the agential collective comes to have what Carol Rovane (1998) calls a "rational point of view." That is: the organisation acts from a particular web of judgments, goals, and intentions, and it updates these in order to maintain both the internal coherence of the web and to achieve the goals via the intentions in accordance with the judgments. This characterisation of agential collectives is neutral between more specific sets of conditions for when a collective has agency (e.g., Gilbert 1989; Pettit and Schweikard 2006, 33; Preda 2012).

That a collective is teleological does not entail that it is agential: a mob is teleological, but not agential. Likewise, a collective can be agential without its members having any goal in common: consider a criminal organisation with a 'cell' structure, which unites members under a procedure while preventing them from having goals in common (List and Pettit 2011, 33). That said, many agential collectives have members that are united around a goal. For example, most members of the agential collective 'BP' act responsively to one another in pursuit of the goal 'making profit.'

Regarding responsibility, Statement C is about what BP has a duty or obligation to do in the future. This makes the responsibility *prospective*. This is the sense of responsibility evoked in debates about corporate social and moral responsibility.² Broadly, we can understand a prospective responsibility as a moral reason that is presumptively decisive in decision-making. That is, if one has a prospective responsibility to do X, then (on pain of blameworthiness) one should presume, in one's practical reasoning, that one ought to do X,

all-moral-things-considered, unless and until one is presented with strong countervailing moral reasons.

The distinction between the three kinds of responsibility is outlined in Figure 1. Moral and causal responsibility attributions are similar in that they are both backward-looking: they both concern what happened in the past. But moral responsibility assertions imply moral judgments about those happenings, while causal responsibility assertions do not. Their common backward-looking nature can be contrasted with prospective responsibility attributions' forward-looking nature. Yet, like moral responsibility, prospective responsibility assertions imply moral judgments. I hope the three example sentences above demonstrate that 'responsibility' can be used in all three senses (see also those cited above on the ambiguity of responsibility). This way of carving up the distinctions brings to light a fourth category: non-moral responsibility attributions that are concerned with future events. Thus we might say: "The upcoming cyclone will be responsible for wild weather." Such predictions are not so much a concern of business ethics, so I leave them aside in this paper.

	Is concerned with past events	Is concerned with future events
Implies moral judgment	Moral responsibility attribution	Prospective responsibility attribution
Implies no moral judgement	Causal responsibility attribution	Predictive causation attribution [not a concern of this paper]

Figure 1. Types of Responsibility

The two tripartite distinctions (between three kinds of collectives and three kinds of responsibility) yield a framework for investigating gaps in collective responsibility. Prima facie, each kind of responsibility is attributable to each kind of collective. This gives rise to nine types of collective responsibility (Figure 2). Statements A, B, and C occupy the top-left, centre, and bottom-right cells. Notice that, in each case, to attribute responsibility to a

collective is to attribute it to the collective as such. That is, when we say that industrialists, the fossil fuels lobby, or BP has *collective* responsibility for some affair, we are not just saying that many or most or all *individual members* have responsibility for it. Instead we are saying the collective *itself* has responsibility. Just what this means—for each collective and each kind of responsibility—will become clearer as the discussion proceeds.

		RESPONSIBILITY		
		Causal Responsibility	Moral Responsibility	Prospective Responsibility
COLLECTIVE	Diffuse collective	E.g., ‘19th century industrialists affected global warming.’	E.g., ‘Present-day motor companies are culpable for causing global warming.’	E.g., ‘Present-day motor companies have a duty to ameliorate global warming.’
	Teleological collective	E.g., ‘The fossil fuels lobby causes lax government policy.’	E.g., ‘The fossil fuels lobby is to blame for global warming.’	E.g., ‘The fossil fuels lobby would do the right thing by ceasing to exist.’
	Agential Collective	E.g., ‘BP had a role in the <i>Deep Horizon</i> oil spill.’	E.g., ‘BP is to be praised for its efforts.’	E.g., ‘BP has an obligation to clean up its oil spill.’

Figure 2: Types of Collective Responsibility

These categories are often obscured by language and reality. As the language in the table demonstrates, different terminology can be used to evoke each of the three kinds of responsibility, which can lead to confusion. Also, real connections between the categories obscure their differences. For example, take moral and prospective responsibility. If one faces moral responsibility, then, we usually assume, one has reneged on a (past) prospective responsibility;³ if one has a prospective responsibility, then one faces the prospect of moral responsibility if one reneges on that prospective responsibility. These connections might

tempt us to equate moral and prospective responsibility. That would be a mistake, since there are also disconnections: one can have a prospective responsibility for repairing a harm while not having moral responsibility for the existence of that harm, and one can have moral responsibility for the existence of the harm while not having prospective responsibility for repairing it (for example, if one is unable to repair it).

2. Responsibility Gaps

2.1 An Example

What does it mean for there to be *gaps* in collective responsibility? In short, it means that we have a pre-theoretic intuition, or (equivalently) a pre-reflective gut reaction, that a particular collective (whether diffuse, teleological, or agential) bears a particular kind of responsibility (causal, moral, or prospective), yet we cannot justify or vindicate that pre-theoretic intuition upon engaging in principled reflection. To be clear, in the cases I have in mind, our intuition or gut reaction is not that “someone or other must be responsible.” If that were the intuition, then we would need to specify exactly *which* collective we are talking about, before we could know that the intuition is not justifiable (this is because the justifiability of the intuition, as we shall see, often hinges on what kind of collective it is). Instead, a collective responsibility gap arises when we have an intuition that *some specific collective* (of one of the three types) is responsible – but it turns out that it cannot be.

To illustrate, consider the Great Pacific Garbage Patch (GPGP). This is a vast area in the North Pacific Ocean—perhaps twice the size of the continental United States (Marks and Howden 2008)—which is, in effect, an “unstrained consommé of small bits of floating plastic” (Kaiser 2010, 150). The bits of plastic that make up the GPGP are in various stages of disintegration. Many are invisible to the human eye, while others are identifiable as bags,

toys, cups, and so on. Most will never fully break down into their molecular parts (NOAA 2017). The precise effects of the GPGP on sea life are uncertain, but various creatures are disposed to ingest the plastic, which in the meantime will have adsorbed toxic chemicals. The plastics and associated chemicals will then enter the food chain (Teuten et al 2009).

The GPGP is an intriguing case for collective responsibility, because it is many miles from any country's coastline and thus no state has jurisdiction over it. Various agential collectives have, however, acknowledged its existence. For example, in 2013, UNESCO and the Italian Ministry of the Environment sponsored a performance art installation aimed at raising awareness of the GPGP (UNESCO 2013). And the Algalita Marine Research Foundation has voluntarily engaged in privately funded research into the GPGP's extent and effects.⁴ But notably lacking is any broader taking of, or attribution of, collective responsibility (whether causal, moral, and prospective)—for example, responsibility placed on states, which regulate plastics production and consumption, or on corporations, which both produce and consume plastic products. The United States' National Oceanic and Atmospheric Administration (NOAA) has described the prospect of cleaning up the debris as a “tremendous challenge” (NOAA 2017). Far from attributing prospective responsibility to collectives such as states or businesses, NOAA's and UNESCO's only suggestions for cleaning up the GPGP target individual action (in short, admonishing individuals to ‘reduce, reuse, recycle’) (NOAA 2017; UNESCO 2013).⁵

One plausible reason for this lack of collective responsibility-attribution is that the GPGP appears to exhibit several types of collective responsibility gap. The rest of this section demonstrates this. A more detailed examination will follow in Section 3.

2.2 Types of Gaps

A. Causal Responsibility Gaps

A causal responsibility gap arises just in case our pre-theoretical intuition is that some particular entity caused some effect, but, upon reflection, it did not. Causal explanations matter for our historical understanding of the world, including how it produces outcomes that harm us. Such explanations are not the primary concern of ethicists. However, victims who are harmed by outcomes are often concerned to know simply *how* something happened, without necessarily knowing *who* is to blame or *who* has a duty to ameliorate the outcome. Thus, causal responsibility does matter ethically, insofar as responding to victim's concerns matters ethically. Let's assume the GPGP is a bad effect. Who—or what—caused it? There are two reasons why *any* intuitive answer to that question might seem indefensible, creating a gap in causal responsibility.

The first is empirical complexity. While some of the ocean's plastic comes from ships, most of it comes from land. Ultimately, it can come from production and consumption that occurs far inland. These causal chains are near-impossible to track. To attribute any part of the GPGP—or, more ambitiously, the whole GPGP—to *any* specific entity is utterly empirically infeasible. Most simply, this is because much of the GPGP is micro-plastics—pieces of plastic between 0.3mm and 5mm in diameter (National Geographic 2017). They often cannot be distinguished one from the other, and so we cannot know whether any particular micro-plastic was eroded off this Coca-Cola bottle rather than that Coca-Cola bottle. And even if a piece of plastic is identifiable as, say, a bottle, it's unimaginable that we would be able to pick it up in the GPGP and trace it to a particular shopper, or even a particular shop or factory.

The second reason for apparent causal responsibility gaps is the conceptual ambiguity of causation. Let's suppose (per impossible) that we know a particular piece of plastic was once part of a particular Coca-Cola bottle, bought by a particular person in a particular shop on a particular day. Who caused the bottle's parts to end up in the GPGP? There are numerous possibilities: Coca-Cola, who designed and marketed the bottle; the consumer, who purchased the bottle; the local government, who didn't install trash cans at the beach where the consumer discarded the bottle; and so on. We might try to say that each of these agents—and more besides—caused it. However, each can argue that, if they hadn't done what they did, the bottle might still have ended up in the GPGP. Coca-Cola can say that someone else would have used bottles from that factory, the local government can say that many people ignore trash cans, and the consumer can point out that plastic in trash cans still makes it to the sea. So in what sense did each of them cause it?

To explain further, suppose you have the pre-theoretical intuition that Coca-Cola caused it. We can problematise this intuition using two popular theories from the metaphysics of causation: the 'probability' theory and the 'process' theory. Under the probability theory, to cause something is to *make that thing more likely* (Schaffer 2016). But had Coca-Cola (say) not developed its products whose bottles ended up in the GPGP, then Pepsi (or some other company) would have had its bottles end up in the bit of the GPGP that Coca-Cola's products in fact ended up in. So Coca-Cola did not raise the probability of (that bit of) the GPGP; if Coca-Cola's bottle hadn't got there, Pepsi's would have. Coca-Cola merely 'pre-empted' Pepsi's causal influence. That is, the GPGP (and even each individual *part* of the GPGP) is over-determined: subtracting the actions of one entity does not reduce the probability of (any part of) the GPGP existing.

Under the 'process' theory, to cause something is to *physically produce that thing*. This theory also seems to give us causal responsibility gaps. Suppose Greenpeace sets out to

clean up the GPGP, but is prevented from doing so because its path is blocked by a fishing trawler. Here, the fishing trawler prevents a would-be-preventer from preventing the GPGP. So, it's natural to have the pre-theoretical intuition that the fishing trawler had some causal role to play in the continued existence of the GPGP. But that role is not one of *physically producing* the GPGP. The process view gives us the result that the trawler is causally disconnected from the GPGP, undermining our intuition. In this case, the process view *under-produces* causation: it doesn't identify all the causation our intuitions tell us are there. Conversely, if we take an extremely wide view of what might count as 'physical production,' then perhaps the trawler does cause the GPGP—but, if we take that wide view, then the process view *over-produces* causation: it will likely tell us that almost every action on Earth causes the GPGP.⁶

That is, under both the 'probability' and the 'process' views, we cannot attribute causal responsibility as often as we like it too. It seems we cannot attribute causal responsibility for the GPGP as we would (pre-theoretically) like, resulting in causal responsibility gaps.

B. Moral Responsibility Gaps

A moral responsibility gap arises when our pre-theoretical intuition is that some particular entity is morally blameworthy or praiseworthy for some effect, but, upon reflection, it is not. Moral responsibility gaps are the most prominent kind of collective responsibility gap in the philosophical literature, being discussed by Jackson (1987), Copp (2007), Pettit (2007), Braham and van Hees (2012), and others. I will focus on blameworthiness (as opposed to praiseworthiness) gaps. These can arise in two ways: first, when *no one* did anything

blameworthy in the lead-up to some harmful outcome; second, when the blameworthy entities' blameworthiness doesn't 'add up' to blame proportionate to the harm.

Young (2011, 126-30) analyses sweatshop labour in these ways. One might think each company that produces or sells cheap clothing is blameworthy for the drudgery of sweatshop work. But Young explains how each company in the supply chain of cheap clothing is operating under pressure and seems not to have any—or at least, not much—blameworthiness for their role. The same might be said about the GPGP. Pre-theoretically, one might think that each corporation that produces and uses plastic is blameworthy for (part of) the GPGP. But each such company is stuck in a competition cycle whereby it must produce plastic trash to survive. It would be too demanding to say they must stop making and selling plastic products. Of course, some companies operate at huge profits—perhaps such companies do have moral responsibility for their contributions (if indeed they do 'contribute'; see the previous section). But others operate within extremely tight margins. So, many companies that produce plastic trash are not morally responsible. Or, at least, they're not morally responsible *enough* that the aggregate of all the companies' moral responsibility is proportionate to the size of the moral outrage the GPGP evokes.

One might wonder: if indeed no corporation is blameworthy, then why do we have the gut intuition that the corporations are blameworthy? I suspect the answer is that humans have a (perhaps unconscious) desire to see *someone* attributed blame for large-scale harmful outcomes. In several psychological experiments, it has been found that people are more likely to believe an outcome was produced intentionally if that outcome was bad, than if that outcome was neutral (Knobe 2003; Nadelhoffer 2004). That is, we have an urge to posit intentions behind harms. It is a short step from positing such intentions, to pinpointing someone's intentions in particular, to blaming that someone. A moral responsibility gap

arises when this ‘someone’ is not actually blameworthy, including because they are not the kind of thing that *can* be blameworthy.

C. Prospective Responsibility Gaps

There are three kinds of prospective responsibility gap: ability, justification, and fulfilment gaps.

An *ability gap* arises when there is some desirable outcome—a good to be produced or harm to be avoided—and the entity (individual or collective) that we (pre-theoretically) believe has an obligation to produce that outcome cannot produce that outcome. Consider the outcome of the GPGP being cleaned up. Organisations like NOAA and Algalita lack the resources needed to produce this. Yet we instinctively want to say they ought to. This is a problem for attributing prospective responsibility to these collectives if we assume that ‘ought’ implies ‘can’: something cannot have a duty to do something it cannot do (Vranas 2007). So, if we face an ability gap, then we cannot attribute collective prospective responsibility as we would like.

Suppose we lack an ability gap. Still, a *justification gap* might arise. This happens if there is an entity that *can* produce the desirable outcome, but there is no compelling justification for why they are *moral required* to produce that outcome. This means we cannot assign prospective responsibility as we might initially want. For example, imagine a private funder—Bill Gates or Warren Buffett—offers sufficient monetary resources to eradicate the GPGP. With the money on offer, there are many organisations (NOAA, Algalita, Greenpeace) able to clean up the GPGP. However, these organisations did not cause the GPGP (whatever that means), nor are they culpable for it, nor do they benefit from its existence. Perhaps they have a duty simply because of their ability (Caney 2010). But mere ability-to-produce-an-outcome can justify prospective responsibilities to produce that

outcome only if the outcome is very morally valuable, relative to other outcomes the agent can produce. So ability will not fill all justification gaps.

Finally, consider *fulfilment gaps* in prospective responsibility. These arise when there are entities who *can* produce the desirable outcome, and there are *justifications* for their duties to produce it. However, they renege. There is a gap between the demands of the prospective responsibility and the reality of the responsibility-bearer's actions. There is a backward-looking component to these gaps: we assess the renegers as morally responsible. However, I categorise these as *prospective* responsibility gaps because the question arises about whether (and if so, on what grounds) other entities have prospective responsibility to 'take up the slack' for the renegers' failures. (Miller 2011, Karnein 2012, Stemplowska 2016).

3. Which Collectives Face Which Responsibility Gaps?

In Section 2, I suggested that each kind of responsibility cannot be attributed as often as we might like. I now want to scrutinise that thought: just how prevalent are these responsibility gaps? This will require considering each collective-type separately, because some gaps are prevalent for some collective-types but not others.

3.1 Diffuse Collectives

Above, I explained that it is highly infeasible to assign causal responsibility for the GPGP to any individual, business, or government agency. This seemed to imply that the GPGP has a causal responsibility gap. But there is another possibility: could a diffuse collective, as a whole, have causal responsibility for the GPGP? Consider the collective composed of all plastics-using agents (individuals and agential collectives). These agents are not united by acting responsively to one another as they work towards a common goal (indeed, their goals

are in opposition, insofar as they are competitors). Nor are they united under a collective-level decision-making procedure. This collective is diffuse. Can we attribute causal responsibility to it? It is tempting to think not. After all, there is no united entity here for responsibility to be attributed *to*.

The metaphysics of causation can help here. I shall use Christian List and Peter Menzie’s difference-making account (a version of the probability account) to show that causal responsibility can be attributed to diffuse collectives, plugging causal responsibility gaps.

Consider the following. What matters for the existence of the GPGP is not that Coca-Cola, Pepsi, Greenpeace, a fishing trawler, a local government, and a consumer—alongside innumerable other agents—each performed certain highly specific actions. As explained in Section 2, the GPGP would still have existed if those actions had been different. Those actions are specific instantiations of a more generally-describable phenomenon: plastics use by plastics-users. If we insist that a causal explanation of the GPGP cannot refer to this diffuse collective, but must instead list all the specific actions of all the specific entities included within that diffuse collective, then we would get a causal explanation that “does not satisfactorily fit the role [of ‘cause’] because it is overly specific and involves extraneous detail.” (List and Menzies 2009, 479-80). To illustrate, List and Menzies give an example from Stephen Yablo:

consider a pigeon that has been trained to peck at all and only red objects. The pigeon is presented with a red target and she pecks at it. As it happens, the target is a specific shade of crimson. What caused the pigeon to peck? Was it the fact that the target was red or the fact that it was crimson? ... The target’s being red is of the right degree of specificity to count as a cause of the pigeon’s action. In contrast, the target’s being

crimson is too specific to count as a cause of the pigeon's pecking: citing it as a cause might give the erroneous impression that the pigeon would not peck at anything non-crimson. (List and Menzies 2009, 480)

The point is this: the target's *being crimson* isn't what makes the pigeon peck. The pigeon would have pecked if the target were maroon, or scarlet, and so on. What makes a difference to the pecking is that the pigeon's target is *some form of red or other*.

Likewise for the GPGP. Coca-Cola's designing bottles isn't what makes a difference to the GPGP existing. Instead, what makes a difference is that some-collection-or-other of entities, taken as a whole, in aggregate, produce large amounts of (what ends up being) plastic trash. The diffuse collective 'plastics-using agents, whoever they happen to be' is the grouping that has the right level of detail to explain how the GPGP occurs. The causal power of this diffuse collective is "realisation insensitive" (List and Menzies 2009, 496; List and Spiekermann 2013; Parfit 1984, 71). That is, regardless of which particular agents happen to be part of this large plastics-using diffuse collective, the GPGP will result; and if the large plastics-using diffuse collective did not exist, then the GPGP will not result. It is this collective that makes the difference.

If we follow List and Menzies in this way, then we can say the diffuse collective 'plastics-users' causes the GPGP, while none of its members does. More generally, we can say that diffuse collectives can have causal responsibility: polluters as a group cause pollution; bargain apparel companies as a group cause the continuation of sweatshop conditions; and so on. Given that diffuse collectives are everywhere and made up of anything (Section 1.2), this result has the power eradicate causal responsibility gaps, assuming we can pinpoint the right diffuse collective in each case. This result also allows us to judge that all statements in the left-hand column of Figure 1 are defensible (at least, pending empirical

demonstration that the collectives named in that column are indeed the collectives with which the relevant effect does occur, and without which it does not occur). By salvaging our ability to provide collective-level historical causal assertions about all three kinds of collective, List and Menzies' account salvages our ability to satisfy victims who are concerned to have historical knowledge of how their situation came about.

So far, so good. But we cannot extend this vindication of diffuse collectives' causal responsibility to their prospective and moral responsibility. The main problem can be explained briefly. It is that diffuse collectives lack the *volitions* necessary for prospective and moral responsibility. To explain, consider that there are numerous outcomes that are *possible*: sweatshop labour could end; global warming could be limited to 2 degrees above 1990 levels; the GPGP could be cleaned up. But for there to be moral or prospective responsibility over those outcomes (or their absence), the outcomes (or absence) must be tied in some way to *volitions* (decisions, intentions, willings, tryings). For there to be moral responsibility for an outcome, there must *have been a past volition that did* (in some way) cause the outcome. For there to be prospective responsibility for an outcome, there must be *the possibility of a future volition, produced via moral reasoning, that could* (in some way) cause the outcome.

Unlike agential collectives, diffuse collectives cannot have their own, group-level volitions, decision, intentions, willings, or tryings (Lawford-Smith 2015, 235)—either in the past or in the future. Of course, we could say that *many* or *most* members of a diffuse collective have individual volitions over the same thing. For example, we could imagine that many or most plastics users acquire an individual volition to eradicate the GPGP. But that would be a statistical fact about the individuals; this does not take the group as a united whole. And if *all* members of the collective were to have the volition, then that would render the collective teleological. (It is innocuous to assume that, when individuals have a goal in common, they will respond to one another, insofar as they encounter one another, in their

individual pursuits of that goal.) With this in mind, it is to teleological collectives that I will now turn.

3.2 Teleological Collectives

Recall that teleological collectives (1) act responsively to one another (insofar as they encounter one another) as they work toward a common goal, but (2) lack clear procedures for making decisions that are attributable to the collective as such. The arguments above about causal responsibility in diffuse groups extend straightforwardly to teleological collectives: if diffuse collectives can be causally responsible, then so can teleological collectives. The unity present in teleological collectives (as compared with diffuse collectives) makes it more likely that, in a teleological collective, the collective level-of-description will be the difference-maker for some effect. This includes such collectives as ‘the plastics lobby,’ ‘democracy-endorsing states,’ and ‘environmental non-governmental organisations.’

Again like diffuse collectives, teleological collectives are poor candidates for prospective responsibilities. While the members of teleological collectives do have volitions with the same content, these groups lack a group-level decision-making procedure that could process obligations and itself produce a volition, via moral reasoning, to perform the action that an obligation requires of the group per se (indeed, the group per se cannot act at all, in virtue of not having such a decision-making procedure). That is, teleological collectives do not have a central core of agency at which it is possible to reason its way to a genuinely group-level volition, or decision, or action. To see this, it is helpful to compare them with agential collectives. An agential collective can decide something (have a volition over it), even if *no member* decides that thing – indeed, even if no member *would* decide that thing. This happens in cases of intra-group compromise. Suppose the Coca-Cola board decides to

say something about the GPGP. Half want to say it is regrettable, but not deplorable; the other half want to say it is deplorable. They all agree ‘awful’ is halfway between ‘regrettable’ and ‘deplorable.’ So, as a compromise, the board decides ‘Coca-Cola will say the GPGP is awful.’ This group-level volition has arisen out of the group’s decision-making procedure: no individual member has decided this. What’s more, if any given individual members *were* to decide for the board, using just her *own* decision-making procedures, then she would have decided differently (choosing either ‘regrettable’ or ‘deplorable,’ not ‘awful’) and would have used a different procedure to decide (using her own deliberative faculties, rather than a compromise procedure). (If she would have taken others’ preferences into account, then this is more like her operationalising the distinctively group-level procedure, than it is like her making her own decision.)

Such uniquely group-level procedures (e.g., compromise) for producing uniquely group-level volitions (e.g. ‘awful’) are possible in agential collectives, but these features cannot arise if members merely act responsively toward a common goal. To illustrate, suppose Algalita and NOAA each adopt the goal of cleaning up the GPGP, and each start to do so responsively (e.g., NOAA starts at the East because Algalita has already started at the West; Algalita then heads South when it notices NOAA heading North). There is not yet any mechanism in place for producing an Algalita-NOAA *distinctively group-level* decision on who cleans which areas. This is because, unlike in an agential collective, any alleged group-level Algalita-NOAA decision: (1) cannot conflict with the individual companies’ private decisions; (2) cannot be other than a stringing together (with ‘and’) of the two companies’ private decisions (‘NOAA will clean the East and North *and* Algalita will clean the West and South’); (3) is compartmentalised: each company has influence only over their own ‘part’ of the decision. Contrast this with our imagined Coca-Cola board decision, where the group

decision (1) conflicted with all members' private decisions, (2) was not made of parts conjoined by 'and,' so (3) did not have parts that were divided amongst members.

One might think decision-making is not the only agential capacity: perhaps teleological collectives have some other agential capacities. But decision-making – and, more specifically, the ability to make decisions via moral reasoning – is necessary to bear prospective responsibility, which is ultimately at issue here. And the above three points apply just as much to beliefs, desires, and intentions as to decisions. Thus, teleological collectives lack many agential capacities that are crucial to prospective responsibility.

Of course, teleological collectives are not entirely like diffuse ones, for at least three reasons. First, many teleological collectives are such that they *would* be transformed into agential collectives *if* all the members aimed at this. (Isaacs (2011) uses 'putative collective' to mark collectives with this potentiality.) Until this has happened, though, there is no distinct source of distinct decisions, at which we can target duty attributions.

Second, when a member of a teleological collective fulfils a prospective responsibility to bring about some good, they will often have to act responsively to other members of the teleological collective, as Algalita and NOAA do in the imagined example. This inter-relatedness of the *content* of members' duties might be thought to suggest that the duties are held by the collective *as such* (Young 2011, 180-1). However, we should not infer this. After all, *all* prospective responsibilities of *all* agents require their bearers to react to other agents. This might be a more pressing consideration within a teleological collective than within a diffuse one, but this is a matter of degree, not of kind. What matters for the existence of a group-level duty is not interrelated *content* of individuals' duties, but the possibility of group-level *reasoning* when deciding whether to produce a volition to act on the duty at all.

Teleological collectives are, then, like diffuse collectives in that they can have causal responsibility but not prospective responsibility.

Third, teleological group members are disposed to rely upon and reinforce one another's endorsement of, and actions towards, the shared goal. It is also rational (at least, more rational than in diffuse collectives) for members of teleological groups to engage in "we-reasoning" around their common goal. When someone we-reasons, he "considers which *combination* of actions by members of the team would best promote the team's objective, and then performs his part of that combination." (Sugden 2003, 167, emphasis original). We can thus view teleological groups thus hold a weak type of shared agency, though this doesn't rise to the capacity for collective-level reasoning. In the imagined example, Algalita's cleaning action encourages NOAA's specific cleaning action, which encourages Algalita's subsequent action, and so on. Each chooses their individual action with the common goal in mind ('cleaning the GPGP') and does their bit of the collective pattern that will produce that result.

In light of this, what should we say about moral responsibility in teleological collectives? Neither option is costless. First, we could decide that the weak type of shared agency is enough for blame. (Using different conceptions of weak shared agency, Sheehy (2007) and Pasternak (2011) each argues along these lines). However, if we say this, then we must take one of two further options: (1) Impute prospective responsibility to teleological collectives. This is something I have argued firmly against (though this is the option Sheehy and Pasternak each endorse). (2) Deny an attractive general claim about the connection between moral and prospective responsibility. This attractive general claim concerns symmetry: "If there is symmetry, then whenever there is [moral] responsibility for an action after the fact, there will also have been an obligation to not have performed that action before the fact." (Lawford-Smith 2015, 241) In other words: to impute blameworthiness to teleological groups, while endorsing my argument that they lack obligations, requires saying

that they can be blameworthy for some outcome without having had an obligation not to produce that outcome.

So much for the first option: imputing blame to teleological collectives. The second option is to point out that the common goals, mutual reliance, mutual reinforcement, and common rational availability of we-reasoning do not make for genuinely group-level volitions. The group cannot have a volition that differs from, and is produced in a way different from, the sum of members' decisions. This is important, because a wrongful volition seems necessary for blame to be justified. What's more, the symmetry claim is highly attractive and (I have argued) teleological collectives lack obligations. Thus, we might deny group-level blame for teleological collectives. This will seem unsatisfactory when some teleological group (such 'the plastics lobby') seems to be particularly connected to a harm. But notice all that we can say in such cases: we can say that each member of the plastics lobby is blameworthy for having the goal of promoting plastics production and consumption, and for encouraging others to have that goal. To explain exactly why this is wrong for each agent, we can appeal to the wrongness of being part of the group that is causally responsible for a bad outcome. This doesn't mean that we must blame each lobbyist for *the entire GPGP* (as Kutz (2000, 122) would be inclined to do). Instead, each agent (whether individual or collective) has performed the wrongful action of making itself part of a harm-producing group.

In my view, denying teleological collectives' blame is the best option. At least, this is true on the assumption that an irreducibly group-level volition is necessary for group-level blame to be justifiable. It will also be the best option if (1) we should retain symmetry between moral and prospective responsibility and (2) I am right that teleological collectives lack obligations.⁷

3.3 Agential Collectives

So far, I have discussed attributing responsibility to diffuse and teleological collectives. These are the trickiest cases for collective responsibility-attributions. When we turn to agential collectives, we can see that they can straightforwardly be attributed all three kinds of responsibility. Indeed, this might seem to follow from arguments for the irreducible moral agency of agential collectives (e.g., given by Copp 2007, Pettit 2007, Rovane 1998, Tollefsen 2015). But it is worth spelling out exactly why. This will allow us to see that there are still some gaps, particularly when it comes to prospective responsibility.

If we apply List and Menzies' argument to diffuse and teleological collectives, then we can also apply it to agential collectives. However, this means that assigning causal responsibility to an agential collective is defensible only when the collective-level phenomenon is what makes a difference to the presence or absence of some effect. For example, imagine that BP's role in the *Deep Horizon* spill was in fact dependent on the actions of just one employee, acting beyond the scope of their role. In that case, the collective-level is not the difference-making level. By contrast, if BP's oversight procedures (for example) more-or-less ensured that some disaster like this would happen at some point, then the collective-level is the relevant causal level. The same is true for diffuse and teleological collectives. But this limitation is particularly worth noting for agential collectives, since we might too readily attribute causal power to the agential collective as such.

For moral responsibility, there are other necessary conditions to add. As alluded to above, moral responsibility requires a certain *quality of will*—a volition, intending, willing, or trying—alongside harm-causing actions. Also alluded to above, agential collectives can have

qualities of will even if some, most, or even all members lack that quality. I mentioned compromises earlier, but discursive dilemmas are the most well-known cases of this kind. Imagine a three-person group that makes decisions via majoritarian voting. It votes on the motions P, Q, and R in turn. Motion P is that humans are not sentient, Q is that humans are nutritious, and R is that humans are a sustainable food source. Assume these are each necessary and together sufficient for it to be permissible to kill and eat humans. A votes for P, for Q, and against R; B votes for P, against Q, and for R; and C votes against P, for Q, and against R. So, by voting on each of P, Q, and R in turn, the group comes to believe in all three motions.⁸ The group thereby comes to believe that it is permissible to kill and eat humans. Being rational, it forms the volitions entailed by this (such as being unfazed if it starts to do so, and intending not to interfere if others do so). In this kind of case, no individual member is morally responsible for the agential collective's coming to this belief—after all, they each individually deny that that all three motions are true—so we must place blameworthiness the group itself (Pettit 2007, 197-8).

However, if we are to thereby blame the collective itself, the collective itself must be a *moral* agent. This requires that the agential collective can respond appropriately to moral reasons. If a group's members are human moral agents, and can recognise moral reasons, then it is natural to assume they can design a group decision-making procedure so that moral reasons are treated as desires, in the sense of having a 'world-to-mind' direction of fit. And it is natural to assume they can design the procedure so that the group, by and large, responds to moral reasons, such that the group takes its *own* measures in response to these reasons (Pettit 2007, 187). Of course, it's unlikely that the group will do this infallibly, but if infallibility were the standard for moral agency, then no human would be a moral agent. So, this is not a conceptual roadblock to agential collectives' moral responsibility. It is, however, a practical roadblock: if some agential collective is constitutionally incapable of recognising moral

reasons (if its procedure is set up so that members cannot present these reasons to the agential collective) then the agential collective cannot be held morally responsible.

If an agential collective is a moral agent—if it can use its decision-making procedure to respond appropriately to moral reasons—then it can have prospective responsibility. This might look like great news for attributing duties regarding such issues as cleaning up the GPGP. However, before we can assign prospective responsibilities to agential collectives, we must ask about *ability gaps* and *justification gaps*. Agential collectives seem powerful when compared with individual agents, so it is tempting to think that they can do anything and everything. Patently this is not the case. Ability gaps abound for agential collectives. In short, this is because agential collectives' actions must be enacted by their individual members. Many agential collectives (e.g., states and business corporations) have a fair amount of control over their members, but this control only goes so far.

Inter-governmental organisations (IGOs) are a clear example here. Very roughly, an IGO is a collective, created by agreement between states, of which the members are states or organs of states. State members are bound by the decision-making procedures of IGOs, so are obliged to follow the role instructions that they are given by the IGO's decision-making procedure (unless the state has entered a "reservation" on that point, upon acceding to the treaty establishing the IGO). This includes abiding by the conventions that the IGO endorses on behalf of its members. However, given that IGOs are run by their members (states), IGOs can only distribute to states those roles that states have given the IGO decision-making procedure the power to distribute—and IGOs can only distribute those roles in the way that states permit through the formulation of the decision-making procedure. Yes, IGOs can and sometimes do provide "warrants," "approval," and "legitimacy" to states' actions (as Harbour (2004, 67) argues in defence of their status as moral agents). But they are not the ones executing the actions and their instructions are not always taken as authoritative. As James

Crawford notes, “[u]nlike states, international organizations *do not possess general competence*: they may only exercise those powers expressly or impliedly bestowed upon them” (2012, 184, emphasis added). IGOs are a particularly clear case. But to some extent or other, this is a problem for all agential collectives. They are at the mercy of their members or some subset of their members. They can be disbanded or have their procedures overridden if enough members (or the right ones) act to make it so. (Of course, members might have obligations to do what they can to ensure their collectives *can*, and *do*, discharge their obligations. I argue for these member obligations in [omitted for review]).

Even if an agential collective is *able* to perform some action, we might face a *justification gap* in giving an account of *why it is obliged* to perform that action. In business ethics, this has long been framed in terms of the debate between shareholder and stakeholder theory. In the terms of this article, that is a debate about when, and why, there are justification gaps in businesses’ responsibilities. According to shareholder theorists, there are justification gaps whenever an action is not in the interests of shareholders (though see Mansell 2013). According to stakeholder theorists, justification gaps are far less widespread, since a prospective responsibility to perform some action can be justified if it is in the interests of a wider set of stakeholders. I do not wish to settle this dispute here. My purpose in mentioning it is to point out that this debate occupies a tiny portion of what we might think about under the rubric of agential collectives’ responsibility, and to show that we can frame this dispute in terms of a dispute about the prevalence of responsibility gaps—more specifically, prospective responsibility gaps, and even more specifically, justification gaps.

Last but not least, consider fulfilment gaps. Almost no matter what one’s view of justification gaps (specifically, whether one is a shareholder or a stakeholder theorist), one will look out into the world and see fulfilment gaps: cases where an agential collective was able to bring about some good, and had a duty to do it, but failed to fulfil that duty. In such

cases, we can place censure upon the agential collective itself. The question will remain of whether other agents (collective or individual) have duties to ‘pick up the slack.’ This is another debate that I cannot settle here, but rather, again, wish to place within my framework: the debate about whether agential collectives have duties to pick up the slack for one another’s failures is a debate in fulfilment gaps, which will, in turn, depend upon one’s position on justification gaps.

The conclusions of Section 3 are summarised in Figure 3.

		RESPONSIBILITY		
		Causal Responsibility	Moral Responsibility	Prospective Responsibility
COLLECTIVE	Diffuse collective	Defensible, if there is collective-level difference-making	Not defensible	Not defensible
	Teleological collective	Defensible, if there is collective-level difference-making	Probably not defensible	Not defensible
	Agential Collective	Defensible, if there is collective-level difference-making	Defensible, if the collective had the relevant quality of will and produced the outcome	Defensible, if the collective has the ability and the duty has justification

Figure 3: Types of Collective Responsibility Gaps

4. Conclusion

What kinds of collective responsibility gaps are there? Many, but not all of them should worry us.

Causal responsibility gaps are the least worrying, because causal responsibility can be attributed to all kinds of collectives. Gaps in causal responsibility arise only if we cannot identify any level of description—individual or collective—that made a difference to the presence of some effect.

Moral responsibility gaps—most pressingly, gaps in blameworthiness—arise when an effect arises from a diffuse collective, such as plastics users, or from a teleological collective, such as the plastics lobby. (Strictly speaking, moral responsibility gaps won't arise for teleological collectives if (1) a rebuttal of moral/prospective symmetry is possible and (2) an irreducibly group-level volition is not necessary for collective blame.) One option for filling these gaps would be to reconceptualise diffuse and teleological collectives as agential collectives, and then attribute moral responsibility to them. However, this runs the risk of mischaracterising diffuse and teleological collectives, and misattributing blame to their members. A better option, it seems, is to reserve our negative reactive attitudes for agential collectives—and for individuals in non-agential collectives.

Likewise, prospective responsibility contains numerous gaps, insofar as duties cannot be attributed to diffuse or teleological collectives, and insofar as agential collectives do not always have the abilities or justifications necessary for duties. I have suggested possible ways of filling these prospective responsibility gaps elsewhere ([omitted for review]). Lastly, for prospective responsibility, there will always be fulfilment gaps: agential collectives will never live up to the moral demands we make of them, even on fairly 'narrow' demands, such as those of shareholder theory.

The diagnosis, then, is mixed. We cannot make as many collective responsibility attributions as we might have hoped, but we can make more than we might have thought.

Compliance with Ethical Standards

Funding: not applicable

Ethical approval: This article does not contain any studies with human participants or animals performed by any of the authors.

References

- Alcañiz Isabella and Timothy Hellwig. 2010. "Who's to Blame? The Distribution of Responsibility in Developing Democracies." *British Journal of Political Science* 41, 389-411.
- Braham, Matthew and Martin van Hees. 2012. "Responsibility Voids." *The Philosophical Quarterly* 61(242), 6-15.
- Bratman, Michael. 1992. "Shared Cooperative Activity." *Philosophical Review* 101(2), 327–341.
- Caney, Simon. 2010. "Climate Change and the Duties of the Advantaged." *Critical Review of International Social and Political Philosophy* 13(1), 203-28.
- Coates, D. Justin and Neal A. Tognazzini (eds). 2013. *Blame: Its Nature and Norms*. Oxford: Oxford University Press.
- Collins, Stephanie. ms. "Obligation and Blameworthiness in Non-agent Groups."
- Copp, David. 2007. "The Collective Moral Autonomy Thesis." *Journal of Social Philosophy* 38(3), 369–388.
- Crawford, James. 2012. *Brownlie's Principles of Public International Law*. Eighth edition. Oxford: Oxford University Press.

- Eriksen M, Lebreton LCM, Carson HS, Thiel M, Moore CJ, Borerro JC, et al. 2014. "Plastic Pollution in the World's Oceans: More than 5 Trillion Plastic Pieces Weighing over 250,000 Tons Afloat at Sea". *PLoS ONE* 9(12), e111913.
- Feinberg, Joel. 1970. *Doing and Deserving: Essays in the Theory of Responsibility*. Princeton: Princeton University Press.
- Harbour, Frances. 2004. "Moral Agency and Moral Responsibility in Humanitarian Intervention." *Global Society* 18(1), 61–75.
- Hart, H. L. A. 1968. *Punishment and Responsibility: Essays in the Philosophy of Law*. Clarendon Press.
- Isaacs, Tracy. 2011. *Moral Responsibility in Collective Contexts*. Oxford: Oxford University Press.
- Jackson, Frank. 1987. "Group Morality." In Philip Pettit, Richard Sylvan & J. Norman (eds.), *Metaphysics and Morality*. Blackwell.
- Kaiser, Jocelyn. 2010. "The Dirt on Ocean Garbage Patches." *Science, New Series*, 328(5985), 1506.
- Karnein, Anja. 2012. "Putting Fairness in its Place: Why there is a Duty to Take Up the Slack", *The Journal of Philosophy* 111, 593-607.
- Knobe, Joshua. 2003. "Intentional Action in Folk Psychology: An Experimental Investigation." *Philosophical Psychology* 16(2), 309-24.
- Kornhauser, L. A. and L. G. Sager. 1993. "The One and the Many: Adjudication in Collegial Courts." *California Law Review* 81, 1-59.
- Kutz, Christopher. 2000. *Complicity*. Cambridge: Cambridge University Press.
- Lawford-Smith, Holly. 2015. "What 'We'?" *Journal of Social Ontology* 1(2), 225-249.
- List, Christian and Peter Menzies. 2009. "Non-Reductive Physicalism and the Limits of the Exclusion Principle." *The Journal of Philosophy* CVI(9), 475-502.

- List, Christian and Philip Pettit. 2011. *Group Agency*. Oxford: Oxford University Press.
- List, Christian and Kai Spiekermann. 2013. "Methodological Individualism and Holism in Political Science: A Reconciliation." *American Political Science Review* 107(4), 629–643.
- List, Christian. 2014. "Three Kinds of Collective Attitudes." *Erkenntnis* 79 (9 Supp), 1601-22.
- Mansell, Samuel F. 2013. "Shareholder Theory and Kant's 'Duty of Beneficence'" *Journal of Business Ethics* 117 (3), 583-99.
- Marks, Kathy and Daniel Howden. 2008. "The World's Rubbish Dump: A Garbage Tip that Stretches from Hawaii to Japan." *The Independent* 5 February 2008.
- Mickler, Steve. 2012. "Blame Liberals: Conservatives Columnists and the GFC in Australia and Canada." *Communication, Politics & Culture* 45(1), 5-23.
- Miller, David. 2007. *National Responsibility and Global Justice*. Oxford: Oxford University Press.
- Miller, David. 2011. "Taking Up the Slack? Responsibility and Justice in Situations of Partial Compliance." In Knight, C. and Stemplowska, Z. (eds), *Responsibility and Distributive Justice*, Oxford: Oxford University Press. Chapter 12.
- Nadelhoffer, T. 2004. "On Praise, Side Effects, and Folk Ascriptions of Intentionality" *Journal of Theoretical and Philosophical Psychology* 24(2), 196-213.
- National Geographic. 2014. "Great Pacific Garbage Patch: Encyclopedia Entry." Updated 19 September 2014. Available at: <http://nationalgeographic.org/encyclopedia/great-pacific-garbage-patch/>
- NOAA (National Oceanic and Atmospheric Administration, United States Department of Commerce). 2017. "What We Know About the 'Garbage Patches.'" Updated 10

- January 2017. Available at: <https://marinedebris.noaa.gov/what-we-know-about-garbage-patches>
- Page, Edward A. and Avia Pasternak. 2014. *Benefiting from Injustice*. Special Issue of *Journal of Applied Philosophy* 31(4), 331-444.
- Pasternak, Avia. 2011. "The Collective Responsibility of Democratic Publics." *Canadian Journal of Philosophy* 41(1), 99-123.
- Pettit, Philip. 2007. "Responsibility Incorporated." *Ethics* 117, 141–201.
- Pettit, Philip and David Schweikard. 2006. "Joint Actions and Group Agents." *Philosophy of the Social Sciences* 36, 18–39.
- Preda, Adina. 2012. "Group Rights and Group Agency." *Journal of Moral Philosophy* 9, 228-54.
- Rovane, Carol. 1998. *The Bounds of Agency: An Essay in Revisionary Metaphysics*. Princeton: Princeton University Press.
- Rudolph, Thomas J. 2003. "Who's Responsible for the Economy? The Formation and Consequences of Responsibility Attributions." *American Journal of Political Science* 47(4), 698-713.
- Scanlon, T.M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Schaffer, Jonathan, "The Metaphysics of Causation", *The Stanford Encyclopedia of Philosophy* (Fall 2016 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2016/entries/causation-metaphysics/>>.
- Sebille, Erik van, Matthew H England, and Gary Froyland. 2012. "Origin, Dynamics, and Evolution of Ocean Garbage Patches from Observed Surface Drifters." *Environmental Research Letters* 7, 1-6.
- Sheehy, Paul. 2007. "Blaming Them." *Journal of Social Philosophy* 38(3), 428-441.

- Stemplowska, Zofia. 2016. "Doing More than One's Fair Share." *Critical Review of International Social and Political Philosophy* 19(4), 591-608.
- Strawson, Peter. 1960 (reprinted 2003). "Freedom and Resentment." Reprinted in Gary Watson (ed.), *Free Will*, 72-93. Oxford University Press.
- Sugden, Robert. 2003. "The Logic of Team Reasoning," *Philosophical Explorations* 6(3), 165–81.
- Teuten E, Saquing J, Knappe D, Barlaz M, Jonsson S, et al. 2009. "Transport and release of chemicals from plastics to the environment and to wildlife." *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 364, 2027–45.
- Tollefsen, Deborah. 2015. *Groups as Agents*. Cambridge: Polity Press.
- Tuomela, Raimo. 2006. "Joint Intention, We-Mode and I-Mode." *Midwest Studies in Philosophy* XXX, 35–58.
- UNESCO (United Nations Educational, Scientific, and Cultural Organisation). 2013. "The Garbage Patch Territory Turns into a New State." Updated 11 April 2013. http://www.unesco.org/new/en/venice/about-this-office/single-view/news/the_garbage_patch_territory_turns_into_a_new_state/#.U71u8fl_u9U
- Waddock, Sandra. 2004. "Parallel Universes: Companies, Academics, and the Progress of Corporate Citizenship." *Business and Society Review* 109, 5–42.
- Werhane, Patricia. 2008. "Corporate Social Responsibility, Corporate Moral Responsibility, and Systems Thinking: Is There a Difference and the Difference it Makes." In G. Flynn (Ed.), *Leadership and Business Ethics (Issues in Business Ethics Volume 25)*, pp. 269–289). Dordrecht: Springer.
- Young, Iris Marion. 2011. *Responsibility for Justice*. Oxford University Press.

¹ In List's (2014) terms, we attribute an 'aggregate attitude' (if the attitude is held by most lobbyists or by the average lobbyist), or perhaps a 'common attitude' (if the attitude is unanimous and is out in the open amongst the lobbying organisations)—but not a 'corporate attitude.'

² The distinction between these is of course disputed. Traditionally, corporate social responsibility is about what a business owes to society at large, while corporate moral responsibility is about what a business owes to a narrower set of entities, such as shareholders and key stakeholders (Werhane 2008, 270-1). Alternatively, sometimes corporate social responsibility is viewed as optional or morally supererogatory, while corporate moral responsibility is mandatory or morally required (Waddock 2004). Yet both concepts concern what the business has *moral reason to do in future*, so they both fall under prospective responsibility.

³ I will question this in Section 3.2.

⁴ <http://www.algalita.org>

⁵ The GPGP is not the only case of this kind. Another glaring example is global warming. Again, it's a harmful outcome, not under any one state's jurisdiction, and for which no collective is being seriously attributed, or is taking, responsibility. I focus on the GPGP case simply because the causal, moral, and prospective processes involved are somewhat simpler. Thus, it is more straightforward for demonstrating the theoretical points. But my argument would, with more complex empirical analysis, almost certainly extend to global warming.

⁶ These problems with the probability and process views are outlined in more detail in Schaffer 2016.

⁷ While I have argued that teleological collectives lack prospective responsibilities, there is perhaps room for debate over whether (1) an irreducibly group-level volition is necessary for

group-level blame to be justifiable and (2) we should retain symmetry between moral and prospective responsibility. Unfortunately, I lack space to pursue these arguments here, which I address in Collins (ms). I thank Avia Pasternak for extensive comments on these issues.

⁸ The discursive dilemma is a generalisation of the doctrinal paradox (Kornhauser and Sager 1993). It is discussed by Pettit and by List in many places, but especially List and Pettit 2011, ch. 2.