

On the Relation Between Collective Responsibility and Collective Duties

NIELS DE HAAN

Abstract

There is good reason to think that moral responsibility as accountability is tied to the violation of moral demands. This lends intuitive support to Type-Symmetry in the collective realm: A type of responsibility entails the violation or unfulfillment of the same type of *all-things-considered* duty. For example, collective responsibility necessarily entails the violation of a collective duty. But Type-Symmetry is false. In this paper I argue that a non-agential group can be collectively responsible without thereby violating a collective duty. To show this I distinguish between four types of responsibility and duty in collective contexts: corporate, distributed, collective, shared. I set out two cases: one involves a non-reductive collective action that constitutes irreducible wrongdoing, the other involves a non-divisible consequence. I show that the violation of individual or shared duties both can lead to irreducible wrongdoing for which only the group is responsible. Finally, I explain why this conclusion does not upset any work on individual responsibility.

1. Introduction

We regularly blame various sorts of groups for their immoral behavior. We blame corporations for causing environmental disasters, riot mobs for destroying the city center, or humanity for exacerbating climate change. In order for blame to be appropriate, the target of blame must be morally blameworthy. We can distinguish between moral responsibility as attributability and moral responsibility as accountability. Moral responsibility as attributability concerns our moral evaluative judgments that are made from an aretaic perspective, meaning we attribute behavior (omissions, attitudes, or character traits) to a person that reflects morally relevant features of that person's self, their excellences and faults, or virtues and vices (Watson, 1996). But there is also the practice of holding agents morally accountable for their behavior, which involves the imposition of demands and a readiness to treat an agent adversely when one fails to comply with these demands. Key to this distinctive stance of holding someone accountable are the negative reactive attitudes resentment, indignation, and guilt (Darwall, 2006; Strawson, 1962; Wallace, 1994). In this paper I only focus on the negative side of

this accountability sense of holding responsible, that is, moral blameworthiness.¹ Is it coherent and appropriate to blame various sorts of groups for their behavior?

A helpful distinction to be made is between scope and degree of responsibility. Scope concerns *what* one is responsible for, whereas degree specifies to what extent or *how much* one is to blame (Zimmerman, 2002). How can we explain locutions such as: ‘The group is blameworthy for ϕ ’? ‘Being blameworthy for ϕ ’ is a predicate of the group, and the scope contains ‘ ϕ ’. The question is whether or not what the group is blameworthy for can always be attributed to its members.

Some groups with sufficient organizational structure, such as states, corporations, and universities, qualify as agents. Group agency can be explicated in terms of a collective decision-making procedure, which allows the group to form representational and goal-seeking states while satisfying desiderata of rationality in a robust manner. The procedure is part of the organizational structure, which further includes rules, policies, and conventions in virtue of which the group coordinates their decision-making and action-taking. In a functional sense akin to individuals, these group agents can understand and process moral reasons and act accordingly (Collins, 2019b; Copp, 2007; Erskine, 2001; French, 1984; Hess, 2006; Hindriks, 2018; Lawford-Smith, 2015; List and Pettit, 2011). Because of this, group agents can have normative competence and be morally responsible. This is called corporate responsibility. Corporate responsibility is an additional non-redundant and non-reductive level of responsibility invoked to avoid ‘deficits in our accountability books’ (Pettit, 2007, p. 194). This does not mean that members are necessarily free from responsibility. Corporate responsibility can be (partially) distributed to members. I call this ‘distributed responsibility’. But because group agents can decide and act in ways irreducible to its members, there are cases where no member is morally accountable for a decision that resulted in wrongdoing (Copp, 2007). Hence, if responsibility can only be ascribed to individuals as enactors or constituters, then there will be ‘gaps in the books’ that we can keep on individuals and there is a shortfall in what we should expect the practice of holding agents morally responsible to deliver (Pettit, 2009, p. 170). Because group agents are legitimate targets for reactive attitudes such as resentment and indignation, we can avoid such responsibility voids.

¹ Henceforth, I use ‘moral responsibility’, ‘moral blameworthiness’, and ‘moral accountability’ interchangeably.

On Collective Responsibility and Collective Duties

But occasionally we think that groups such as a random collection of passengers in a subway, riot mobs, men, the affluent or even humanity are morally responsible as well. These unstructured groups lack the required organizational structure for group agency. Because these groups are not agents, one might think that we must explain such claims necessarily in terms of shared responsibility. The content of the group's responsibility is fully distributed to its members. Shared responsibility is nothing else but aggregated individual responsibility.

In this paper I argue that non-agential groups can be morally responsible (see also Chant, 2015; Feinberg, 1968; Held, 1970). I call this collective responsibility. I will argue that, like corporate responsibility, collective responsibility is needed in order to avoid deficits in the accountability books. This means that what the non-agential group is blameworthy for is not exhausted by the summation of what all members are individually responsible for. The group's blameworthiness for ϕ is non-distributive and purely collective. This does not (necessarily) mean that the members cannot be responsible for their part, but ϕ does not feature in the members' responsibility. Occasionally, we do have negative reactive attitudes towards these social groups. In my view, in some instances it is coherent to blame these unstructured groups and these reactive attitudes can be justified.

This leads us to the following typology²:

	Non-Agential Group	Group Agent
Distributive	Shared Responsibility	Distributed Responsibility
Non-Distributive	Collective Responsibility	Corporate Responsibility

Figure 1. Types of Responsibility in Collective Contexts.

If our practice of holding agents morally accountable is connected to the moral demands under which these agents are placed, then one might expect that we can learn something from such responsibility claims about obligations in the collective realm. In particular, about whether and how unorganized groups can have obligations to

² Stephanie Collins and I set this briefly out elsewhere, see Collins & De Haan (Forthcoming). A similar suggestion can be found in Preda (2012).

address various collective problems.³ For example, can humanity have a collective obligation to combat climate change? As I discuss later, we can identify four obligation-counterparts analogous to the types of responsibility: shared, collective, distributed, and corporate duties. This allows the isolation of the most controversial claim: that non-agential groups can have (non-distributive) collective duties. I will not defend this claim.

Instead I investigate the relation between these collective variants of responsibility and obligations. If moral responsibility is connected to the moral demands that fall on agents, then, when considering the collective realm, one might think the following principle is true.

Type-Symmetry: A type of responsibility necessarily entails the violation or unfulfillment of the same type of *all-things-considered* duty.

For example, collective responsibility necessarily entails the violation or unfulfillment of an *all-things-considered* collective duty. My main aim in this paper is to show that Type-Symmetry is false, because non-agential groups can be collectively responsible without violating a collective obligation.

To show that non-agential groups can be collectively responsible, I will pry apart the agency condition from moral responsibility and attack the assumption that (collective) responsibility necessarily requires agency. I suspect this will raise some eyebrows. To some, 'blameworthy non-agential groups' may sound like an oxymoron. But I ask the reader to reserve their skepticism until the end of the paper. Most have trouble even considering this because of strong convictions about individual responsibility, but I will explain why this need not upset any work on individual responsibility.

The structure is as follows. First, I discuss the intuitiveness and relevance of Type-Symmetry. Second, I set out a case involving a blameworthy collective action and assume that collective actions are non-reductive. In order to show that this case is a counterexample to Type-Symmetry, I discuss the different types of both responsibility and obligations and show that not the same types apply to the case. Third, to answer a possible rebuttal that collective actions are in fact reductive, I set out a different counterexample to Type-Symmetry involving non-divisible consequences. Fourth, I show why the possibility that non-agential groups can be collectively responsible is not a problem for a comprehensive ethical theory. Finally, I show how both cases can be generalized.

³ For a helpful overview, see Schwenkenbecher (2018).

2. The Intuitiveness and Relevance of Type-Symmetry

The intuitiveness of Type-Symmetry is derivative of a very plausible principle at the individual level. In order to justifiably hold someone morally blameworthy for an act, as in actively blaming that person, one must hold that person *to be* morally culpable for that act, otherwise one's reactive attitudes would be out of place.⁴ Reactive attitudes such as resentment and indignation have a cognitive dimension in the sense that they implicitly address demands (Wallace, 1994, pp. 25–33). These negative reactive attitudes 'invariably involve "an expectation of, and demand for" certain conduct from one another' (Darwall, 2006, p. 16).⁵ If someone fails to behave in a certain way, they fail to show reasonable regard for someone's moral status. To demand certain behavior of an agent is 'to lay it down that unless the agent so behaves she will be liable to certain adverse or unwelcome treatment' (Watson, 1996, p. 236). Holding accountable involves the idea of liability to sanctions (Shoemaker, 2011; Watson, 1996). In order for it to be fair to hold a person morally accountable for an individual act (or omission) that agent must have had a corresponding moral obligation that he or she failed to follow (Wallace, 1994, p. 19).⁶ Moral accountability is therefore inextricably linked with moral obligation (Darwall, 2006, p. 94).⁷ One can do something morally wrong in a morally blameworthy way only if one is morally required and expected to do what is morally right.

This makes a lot of sense in ordinary cases. Suppose agent A is morally blameworthy for act ϕ , say stealing your wallet. But A is morally blameworthy for ϕ only if ϕ constitutes moral wrongdoing, and ϕ constitutes moral wrongdoing only if A violated (or failed to fulfill) an *all-things-considered* moral obligation concerning ϕ .

⁴ See also Smith (2007). Smith remarks that the reactive attitude need not be expressed in order to be considered active. Note that this can be true without making further claims about attributability as Smith does.

⁵ Darwall quotes P.F. Strawson (1962).

⁶ In Wallace's view, to adopt the stance of holding a person to a demand just is to be susceptible to the negative reactive attitudes when that person violates the relevant moral requirement.

⁷ Darwall argues that moral obligation is conceptually related to standards of minimally decent conduct that moral agents are accountable for complying with. The forms of moral accountability, blame, guilt, indignation, punishment, and so on, imply that agents have (conclusive) reasons to do what they are morally obligated and accountable for doing. See also Skorupski (2010, Chapter 12).

Thus, if A is morally blameworthy for stealing your wallet, this entails that A had a duty to refrain from said act *all-things-considered*, namely a duty not to steal. Of course, the control and the epistemic condition (however formulated) must be satisfied, otherwise the agent has a legitimate excuse and cannot be blameworthy for this instance of behavior.

In my view, moral blameworthiness is connected to subjective wrongdoing.⁸ Morally conscientious agents are concerned with evidence-subjective oughts in their moral deliberation, that is, what one ought to do based on one's available evidence. A typical blameworthy act can be described as follows: If, at t_0 , moral agent A has a (evidence-subjective) duty not to ϕ *all-things-considered*, and next at t_1 , A ϕ 's without a legitimate excuse, then A is morally blameworthy for ϕ -ing. So, the backward-looking practice of holding others morally responsible is situated within the distinctive nexus of subjective moral obligations, rightness and wrongness, as reactive attitudes and moral sanctions are only justified when an agent is morally accountable, and this implies a connection to a set of moral requirements (Hindriks, 2014, p. 1566; Wallace, 1994, p. 63).

This is captured by the following principle:

Act-Evaluation: If agent A is morally blameworthy for act ϕ , then ϕ -ing is morally wrong, and A violated or failed to fulfill an *all-things-considered* moral duty concerning ϕ .⁹

Given the clarity of how individual blameworthiness functions, one might think the same applies at the collective level. If a group is blameworthy, then it did something wrong without an excuse, and it must have violated an obligation.¹⁰ So, Act-Evaluation lends some initial intuitive support to Type-Symmetry.

⁸ So, moral blameworthiness presupposes subjective wrongness and the violation (or non-fulfillment) of a subjective 'ought'. See also Haji (1993, p. 46), Jackson (1986, p. 352); Parfit (1984, p. 25)

⁹ Support for variants of this principle is widespread, albeit often implicit, see among others: Cane (2002, Ch. 3); Hindriks (2014, p. 1566); Gilbert (2006, p. 98); Isaacs (2011, p. 130); Lawford-Smith (2015, pp. 239–43); Ross (1939, pp. 163–64); Smith (2007, p. 476); Shoemaker (2011, p. 623); Wallace (1994, pp. 62–63); Widerker (1991, p. 223). For criticism, see: Graham (2010, pp. 93–94); Scanlon (2015); and Zimmerman (2008, Ch. 4).

¹⁰ Especially the idea that blameworthiness necessarily requires the violation of a moral norm is prevalent within the debate on collective responsibility. See for example Gilbert and Priest (2020, p. 24); Tuomela and Mäkelä (2020, p. 65); or Hindriks (2014, p. 1566).

On Collective Responsibility and Collective Duties

If Type-Symmetry is correct, then collective responsibility influences what one can say about collective obligations.¹¹ In particular, this opens the door for two particular arguments. Take any case where a non-agential group appears to be collectively responsible. If Type-Symmetry is correct, one can deny that a non-agential group is collectively responsible via *modus tollens*, namely: If there is collective responsibility, then the non-agential collective must have violated a collective obligation. The non-agential collective did not (and cannot) have a collective obligation, because it is not a collective moral agent. Therefore, it cannot be collectively responsible (see also Collins, 2019a, p. 951; Lawford-Smith, 2015, p. 241).

At the same time, for those who reject collective duties, precisely this reasoning forces them to deny that there is group-level responsibility in certain cases even though this is counterintuitive.¹² Subsequently, one's *modus tollens* being another's *modus ponens*, one could argue that if there is good reason to think the non-agential collective must be collectively responsible, then there must have been a collective obligation. So: If a non-agential group is collectively responsible, then it must have had a corresponding moral obligation. The non-agential collective is collectively responsible. Therefore, non-agential groups can have collective obligations.

Thus, Type-Symmetry either forces one to accept that non-agential groups can have collective duties, which comes with certain costs, or it limits what one can say about responsibility in collective contexts. Rejecting Type-Symmetry is important, then, because it creates room for an ethical theory that is partly or entirely individualistic on the obligation-side, but that at the same time can invoke collective responsibility to avoid shortfalls in collective contexts on the responsibility-side. For example, one might accept that group agents can have corporate duties, reject that non-agential groups can have collective duties, but accept that non-agential groups can be collectively responsible. Or one could reject group agency, collective and corporate duties altogether but accept the possibility of collective responsibility.¹³ Such theories

¹¹ Three authors explicitly accept Type-Symmetry, see Lawford-Smith (2015, p. 241); Björnsson (2014, pp. 111–14), Schwenkenbecher (2014, p. 64; 2019, fn. 22). In Collins (2019a, p. 951), Type-Symmetry is discussed, but no definite position is taken by Collins, see fn. 8. Moreover, Virginia Held's (1970) seminal work can be read in this way. For example, see Schwenkenbecher (2013, p. 317).

¹² For such a case, see Chant (2015) or Jackson (1987).

¹³ This is mistaken in my view, but the argumentation here does not show this.

would have the same explanatory power in the sense that responsibility voids are avoided and the question of how to explain the various 'oughts' on the obligation-side can be treated in a largely isolated manner.

In what follows, I will show that Type-Symmetry is false. To do so, I will not focus on collective omissions, because in such cases there is a strong intuition that the agents should have acted together, which precisely gives rise to the aforementioned *modus tollens* – *modus ponens* standoff (see Chant, 2015; Feinberg, 1968; Held, 1970). Instead I focus on blameworthy collective actions, because, as I will argue, in these cases it makes less sense to think there is a collective duty.

3. Purposive Groups, Collective Actions and Irreducible Wrongdoing

In this section I will present a counterexample to Type-Symmetry and discuss the relevant types of duties and responsibility.

a) *Starvation*

Not every instance of acting together involves a group agent. It is possible for agents do something together intentionally, acting on the basis of a collective intention, without those agents constituting a group agent (Pettit and Schweikard, 2006, p. 33; List and Pettit, 2011, p. 34). We can distinguish between two types of aggregates: random collections of agents that are merely bound together through coincidence and groups that can be defined in terms of a common feature and 'by reference to the solidarity which allows the members of the group to engage in joint purposive behavior', such as teams or mobs (May, 1989, pp. 22–23). These *purposive groups* are capable of joint purposive behavior but lack any organizational structure and decision-making procedures. With that in mind, consider the following case.

Starvation: An incited mob of twenty farmers is dissatisfied with their standards of living. When storming through the countryside, the mob encounters a prosperous castle that has four heavy doors that connect to four bridges. The mob intends to starve the citizens to death for monetary gain by blocking their entrances and exits. It takes three people to effectively block each bridge. The mob surrounds the castle. The mob has no leader and no single member can calm down the entire mob. After a number of days, the citizens die of starvation.

On Collective Responsibility and Collective Duties

This is clearly a case of collective intentional behavior. In everyday thought and talk we are ready to ascribe a single action to the mob, to *them*. The mob surrounds the castle and forces the citizens into starvation. Surrounding the castle is a predicate that is ‘purely collective’ (French, 1984, pp. 14–18). The purposive mob is not a group agent, because the mob has no decision-making procedure for assigning roles whose fulfillment will enable the pursuit of group’s goals.

This is a paradigm case of a *collective action* that features a collective intention. How can we make sense of the following statement of a member of the purposive mob: we intend to surround the castle and block the four bridges together? There are several accounts, but for our current purposes it is best to classify them as follows: the we-intention is either *reductive* or *non-reductive*. Content-accounts hold that what is collective about collective intentionality is the *content*: to surround the castle and block the four bridges.¹⁴ A key feature of Michael Bratman’s account, for example, is that the shared intention that we surround the castle reduces to a complex structure of attitudes of individuals.¹⁵ But not all accounts are reductive. Mode-accounts hold that what is collective is the mode of intending, the agents must intend *collectively* to surround the castle and block the four bridges (Searle, 1990; Tuomela, 2013). Subject-accounts hold that what is collective is the *subject*, the agents form a plural subject that has the intention to surround the castle (Gilbert, 2014; Schmid, 2014). (Most) mode- and subject-accounts are non-reductive, meaning they hold that the collective intention is not reducible to individual intentions and/or attitudes.

Let us assume that the collective intentional action is not reducible to the summation of individual intentional acts (Searle, 1990). Not much is riding on this but let’s say that the collective act supervenes on, but is irreducible to, the individual contributory acts (Tuomela, 2017, pp. 25–33). Let me point out that the minimum required base on which the collective action must supervene is twelve, because the actions of three people is sufficient per bridge. But given that the twenty agents do act together, it is nonetheless correct to say that the collective act in Starvation supervenes on the twenty individual acts of the members.

Two questions arise concerning *Starvation*: What type of ought did the plurality of agents violate and in what sense is the non-agential

¹⁴ See Schweikard and Schmid (2013) for setting out the debate in terms content, mode and subject accounts.

¹⁵ See Bratman (2014, Ch. 2 & 3). For other reductive accounts, see Kutz (2000); Ludwig (2016).

group responsible for their collective action? I will start with the type of responsibility.

b) What Type of Responsibility?

In Starvation the mob is blameworthy for surrounding the castle and the subsequent deaths of the citizens. But what type of responsibility is at play here? I will argue that the agents in Starvation do not merely share responsibility, but that the group must be collectively responsible. Remember that if the agents merely share responsibility, then the content of the scope of the group's responsibility is fully distributed over the scope of the members' responsibility.

Elsewhere, Stephanie Collins and I (Forthcoming) distinguish between three principles of distribution by focusing on the scope of responsibility:

The Portion Principle: Each individual is blameworthy for a different portion of what the group is blameworthy for, with no overlap between individuals' scopes.

The Overlap Principle: Each individual is blameworthy for a portion of what the group is blameworthy for, with some overlap between individuals' scopes.

The Full Scope Principle: Each individual is blameworthy for the totality of what the group is blameworthy for. There is complete overlap between individuals' scopes.

Collins and I argue that the interconnectedness within the group determines which principle of distribution applies. This level of interconnectedness is determined by three factors: shared intentionality, interpersonal influence, and common knowledge. At the strongest level of interconnectedness, the Full Scope Principle is applicable. I will use our account to analyze the group's responsibility in Starvation.

Let's start with the Portion Principle. In our view, the Portion Principle is only applicable to very loose groups with very low interconnectedness. For example, a group where the agents have no shared intentions, no interpersonal influence and merely population common knowledge. Each agent is only responsible for their own contribution and nothing else.¹⁶ This does not appear to be the

¹⁶ The Portion Principle is still consistent with H.D. Lewis' (1948, p. 3) strict individualist principle that no one can be morally responsible for the conduct of another.

On Collective Responsibility and Collective Duties

case in Starvation, as all members we-intend to surround the castle and there is common knowledge.

The presence of shared intentionality and common knowledge suggests that there is at least medium interconnectedness within the group, meaning that minimally the Overlap Principle is applicable. Members are not only responsible for their own contribution, but potentially for other members' contributions as well depending on the role they played (May, 1992, p. 39).¹⁷ The harm would be divided into parts and each participant would be assigned some but not all responsibility for the harm (May, 1992, p.42). The degrees of responsibility may vary per agent and differential shares are not determined in proportion to the number of people who constituted the group (May, 1990, p. 273). Even if some agents hold a greater share of responsibility, this does not diminish the responsibility of others. Responsibility is distributed, either evenly or unevenly, depending on the exact roles members played and their (possible) contributions (Feinberg, 1968, p. 685). A member may inherit a share of responsibility without (full) control over what is in their scope. For example, suppose a member has influenced other agents to contribute. Because of this indirect interpersonal influence, the agent is also blameworthy for the other members contributions. Importantly, this includes counterfactuals: if a member could reasonably speaking have influenced others not to contribute, but she fails to do so, then their actions are within her scope. In Starvation there are no leaders with greater influence than others, meaning interpersonal influence is symmetric rather than asymmetric. Note that although the scopes of members may overlap, the scope is still determined by members' contribution and interpersonal influence. Given that no member had sufficient influence to stop the collective act, each member is only *partially* responsible based on the role one played.

This means the Overlap Principle cannot account for the collective wrongdoing in Starvation. Surrounding the castle with four bridges is not something any individual can do alone. We cannot understand the moral significance of the individual actions without looking at the collective action. The collective intentional act has a distinctive moral wrongness that colors the individual actions it supervenes on. The moral weight of the collective act, forcing people to starve to death, defines the moral dimensions of the contributory acts of members (Isaacs, 2011, p. 58). Each contributory act is morally wrong, because the collective act is wrong.

¹⁷ The Overlap Principle is largely in line with May's view.

In *Starvation* the collective act of starving the citizens to death is something that supervenes on but is irreducible to the individual actions. This means there is irreducible collective wrongdoing. The conjunction of member-responsibility for their contributory acts does not equal the responsibility for the entire action, because the collective act is something more than just the conjunction of these individual acts. After identifying all relevant agents and distributed responsibility, it is evident there is a remainder. No agent is responsible for the collective act of starving the citizens to death. The distinctive moral wrongness of the joint action is simply not captured by isolating individual contributions.

What about the Full Scope Principle? If each member is responsible for the entire collective action, then the irreducible collective wrongdoing has been accounted for. However, this principle is not applicable to *Starvation*. Christopher Kutz (2000, p. 141) argues that complicit members are properly held accountable for group actions in which they participate, because these actions represent a member's own conception of their agency and projects, it expresses what they desire, tolerate and believe. Collins and I agree with Kutz that applying the Full Scope Principle may be justified in some cases, but we argue that it cannot apply to all cases. Collins and I argue that Kutz's justification fails to take (non-moral) luck into account. When acting together in groups, I cannot control everything that others do. But the group action or outcome is taken to be a reflection of my agency. This is a deep tension. Collins and I argue that the application of the Full Scope Principle is justified only if the group is highly interconnected, meaning all three factors (shared intentions, interpersonal influence, and common knowledge) are maximally present. When these three factors are present, then luck is brought down to a minimum.

One such example is a joint killing. Suppose three agents we-intend to murder someone in a public bathroom. One holds the victim down, the other stabs the victim, and one is the look-out. Assume that each participatory act is necessary for the joint action to succeed. First, consider the intentionality. The agents have a joint intention to kill the person, this means their joint killing implies a strong reliance on the others' actions. Second, consider the interpersonal influence. Each agent has a strong influence over the other involved agent's choices and actions. But it equally includes counterfactuals: each agent could have influenced the other not to perform their part. Third, there is common knowledge in the strongest sense. Each agent knows exactly what the others intend and know that the others know that he or she knows. In such cases, Collins and

On Collective Responsibility and Collective Duties

I agree with Kutz that the collective act reflects their agency, hence the Full Scope Principle is applicable.

However, one should be wary of simply scaling-up from small-scale cases to large-scale joint actions and apply the same principle. If I do not have any interpersonal influence over the actions of some members, then the means these other members take towards the satisfaction of our joint intention may be completely different from what I had in mind. We may act for the same group reason, and have the same we-intention, but the means we take may be very different. Hence, it is a matter of luck if the others do the same as I do. In order to avoid the possibility of luck, such that the group action is an accurate reflection of *my agency*, I must have sufficient interpersonal influence over the actions of the others.

In Starvation no individual could have prevented or stopped the collective action. Luck re-enters the picture. Perhaps a subset of members over which I have no control take it too far. This of course does not justify my action, but in such groups the collective act is no longer an accurate reflection of my agency. Only when the members have significant interpersonal influence can we justifiably claim that the action is an accurate reflection of their agency. If this is correct, then the Full Scope Principle is not applicable to Starvation. The members do not have the right level of interpersonal influence. Note that if this is not yet convincing in Starvation, we can add n members to the case, at some point there will be a severe lack of interpersonal influence.

Let me emphasize that this does not necessarily imply a lower *degree* of responsibility. I am merely saying that the justification for a particular principle of distribution concerning the *scope* of responsibility does not apply. Because of the lack of interpersonal influence, the possibility of luck makes it inappropriate to say that the entire collective action is an accurate reflection of an agent's agency when the interconnectedness within the group is not characterized by high interpersonal influence. Therefore, the Full Scope Principle does not apply in Starvation. The members are, of course, still blameworthy for doing their part and given the severity of the action to a great extent.

Next, the following argument holds:

Premise 1. If the Portion or Overlap Principle applies in Starvation, then no member is morally blameworthy for the irreducible collective wrongdoing.

Premise 2. The application of the Full Scope Principle is not justified in Starvation.

Inference 1. No member is morally blameworthy for the irreducible collective wrongdoing in *Starvation*.

Premise 3. If an act constitutes moral wrongdoing and there are no legitimate excuses, then, necessarily, whoever performed that act is morally blameworthy for that act.

Premise 4. The mob's collective act of starving the citizens in *Starvation* is non-reductive, constitutes moral wrongdoing and there is no excuse for doing so.

Conclusion. The mob is collectively blameworthy for the collective act of starving the citizens in *Starvation*.

The argument is valid, so if the mob is not collectively responsible, there must be something wrong with the premises. I have already defended Premise 1 and Premise 2. Premise 4 is largely based on the assumption that collective actions are non-reductive.

Premise 3 spells out the underlying idea of filling responsibility voids. Remember that corporate responsibility is invoked in order to avoid deficits in the accountability books. The idea is that unexcused moral wrongdoing must not go unmet, otherwise the practice of holding wrongdoers accountable fails to deliver what is expected. It must not be possible to get away with unexcused moral wrongdoing simply by incorporating as a group agent or, in this case, by acting together. In *Starvation* we must invoke collective responsibility in order to avoid such a deficit.

Against Premise 3, the *Type-Symmetry* proponent might argue that the nature of an action need not determine the blameworthiness of it. A young child can perform the same intentional act as an adult, say stealing money, but will not be blameworthy for it. The child is exempted from moral responsibility. So, there must be something wrong with Premise 3. The child is not responsible, I agree, but there is a step missing that shows that the two cases are disanalogous. Exemptions show that an agent is not capable of adult interpersonal relationships and therefore is (temporarily) exempt from the community of agents whose conduct one could reasonably require to comply with the demands of the moral community (McKenna, 2005, p. 167; Strawson, 1962; Wallace, 1994). A young child is incapable of (fully) understanding moral reasons, which is why the child lacks normative competence and is not (yet) a moral agent. Therefore, the child cannot have a moral imperative to refrain from that action, which is why the act does not constitute *moral* wrongdoing and the child is not blameworthy.

On Collective Responsibility and Collective Duties

But the members of the mob in *Starvation* certainly do qualify as moral agents. We have not yet determined the precise type of duty, but clearly, they violated at least some type of moral duty. Whether this violated duty is shared or collective, the acts they together commit do constitute moral wrongdoing. I understand moral wrongdoing here as shorthand for the kind of act that would be considered a moral wrong were it to have been committed by one or more moral agents.¹⁸ This is undeniably true in *Starvation*, whereas this is undeniably false in the case of the child. To be clear, the assumption I am attacking is that, at the collective level, collective responsibility necessarily requires agency. If the collective act is morally wrong, why would this act not be blameworthy? There are no exempting factors other than the purposive group itself not being an agent, but the individuals certainly do have normative competence.¹⁹ And surely all agents knew or should have known that starving people is morally wrong. None of them were forced to do so either. They could have refrained from partaking in the collective action. If there are also no excusing factors, there is no reason to think that no one is responsible for the non-reductive collective wrongdoing. Instead we have overwhelming reason to think that *they* are blameworthy. The mob is non-distributively collectively responsible for jointly starving the people to death.

If the group is collectively responsible, then those who subscribe to Type-Symmetry are committed to saying that the duty is collective in *Starvation*. Let's see if this is correct.

c) *What Type of Duty?*

How can we understand the locution: 'The group ought to ϕ *all-things-considered*'? I am not concerned here with *pro tanto* duties, because only the violation of *all-things-considered* duties can generate moral blameworthiness. The scope of an obligation concerns *what* is required, that is, the content of the ought. We can distinguish again between four types.

¹⁸ For such a formulation in singular terms, see Sepinwall (2016, p. 5).

¹⁹ One might reply that on Tuomela's view the agents do qualify as a we-mode group agent. Note that unless whenever we perform a joint action, we thereby constitute a collective *moral* agent that has normative competence in its own right, this response is going to be problematic.

	Non-Agential Group	Group Agent
Distributive	Shared Duties	Distributed Duties
Non-Distributive	Collective Duties	Corporate Duties

Figure 2. Types of Duties in Collective Contexts.

Corporate duties are duties of a collective moral agent. These duties are non-distributive when what is morally required *all-things-considered* from the collective agent differs from what is morally required *all-things-considered* from its members. In some instances although the collective moral agent has a corporate duty to ϕ *all-things-considered*, members may only have a relevantly related *pro tanto* duty to do their part, because the overriding reasons for members may be different from those of the corporate agent (Copp, 2007). When the collective ought does result in *all-things-considered* duties on part of the members, then the group agent has *distributed duties*.

What about non-agential groups? When a plurality of agents faces a collective action problem, they can be morally required to act together in order to address a morally pressing matter. If the group’s ‘ought’ is non-distributive, then the content of what is *all-things-considered* required of the group is not fully explicably in terms of what is *all-things-considered* required of the relevant individuals. To claim there are such *collective duties* comes with a cost. Non-agential groups do not qualify as collective moral agents, therefore the question arises how such a group can exactly be morally required to ϕ when no individual alone can be morally required to ϕ . If one thinks the group is the bearer of this duty, then this forces one to rethink the (extent of) the Agency Principle: only moral agents can have moral duties (see Aas, 2015; Isaacs, 2011; May, 1987; Schmid, 2018; Wringer, 2014). If one thinks the agents somehow have this duty jointly, then this forces one to rethink the Ability Principle: A duty-bearer can have a duty only if the duty-bearer is able to do what the content of that duty specifies (see Björnsson 2014; Pinkert, 2014; Schwenkenbecher, 2014). Either way, collective duties require us to amend an important ethical principle.

I do not wish to make any claims here about whether or not we should accept collective duties into our typology.²⁰ This largely

²⁰ Some philosophers accept that there are corporate duties but think there are no collective duties. See Collins (2019b) or Lawford-Smith (2015).

On Collective Responsibility and Collective Duties

hinges on whether it is possible to account for all collective action problems in terms of *shared duties*, which are aggregated individual duties (that are possibly conditional). But even if there are collective duties, not every locution about what a non-agential group ought to do is non-distributive. Some nuance is in order.²¹ This will sharpen our understanding of the relation between duties and responsibility. In fact, there are (at least) two sub-types of shared duties: unconditional and conditional individual oughts.

Let me start with a simple example. Suppose a teacher says: 'You two ought to be quiet'. The scope of the ought here is not a collective action. We can explain this locution by saying that John ought to be quiet by himself and Billy ought to be quiet by himself.²² The 'ought' is fully explicable in terms of individual duties.²³ Moreover, the content of the collective *all-things-considered* 'ought' is fully distributable over the content of the conjunction of the individual *all-things-considered* duties.²⁴ Suppose Billy notices that a fire breaks out and wants to alert the class, he now has an overriding reason why he ought not to remain quiet. If so, then Billy and John no longer ought to be quiet *all-things-considered*.²⁵ But Billy still ought to be quiet when John starts shouting. The individual duties are unconditional in the sense that they do not rely on the actions of others.

With this in mind, let us consider Starvation again. Clearly, the group ought not to have surrounded the castle and starved the people to death for their own monetary gain. But the question is whether the group has violated a shared or collective duty. If the group has a collective duty, then what is morally required of the group must not be fully explicable in terms of what is required of

²¹ To the best of my knowledge, defenders of collective duties typically do not consider the possibility of various types of 'oughts'.

²² Unless perhaps when the teacher makes each responsible for the other's silence.

²³ This would simply mean that this particular collective 'ought' is analytically reducible. In general, of course analytic reductionism may still be false: Analytic reductionists hold that claims about collective obligations have the same meaning as more complex conjunctions of claims about individual obligations (see Wringer 2016, p. 486).

²⁴ This means this collective 'ought' is ontically reducible. Any claim about this type of ought is made true by facts about individual duties.

²⁵ It is hard to see how in such cases the collective ought could possibly supervene on the individual duties in a more complex way such that multi-realizability is possible. Wringer (2016), for example, uses multi-realizability in order to argue for the ontic irreducibility (in my terms non-distributiveness) of collective obligations.

the individuals. It may appear at first sight as if this is the case, because there is a collective action that constitutes wrongdoing. But this does not entail that the violated 'ought' is non-distributive.

Note that '*not* surrounding a castle' is not a purely collective predicate. In a non-trivial sense, every member is able to refrain from surrounding the castle. The members need not we-intend not to surround the castle in order to all intend not to surround the castle. If each member has an individual intention to not surround the castle, then, they, in a reductive sense, are all intending to not surround the castle and are plausibly not going to surround the castle. Similarly, if I ought *all-things-considered* to refrain from surrounding the castle, and you ought *all-things-considered* to refrain from surrounding the castle, and we do not violate our obligations, then we refrain from surrounding the castle. So, what the group ought to do *all-things-considered* is fully explicable in terms of what the agents individually ought to do *all-things-considered*, because refraining from partaking in a collective action need not be a collective action in itself. The content of the 'plural ought' does not contain any purely collective predicates. This means that each individual can be morally required *all-things-considered* to do so without this potentially upsetting the ability principle. The content of the 'plural ought' is distributive.

Some may object that in Starvation the obligations of each member are dependent on what the others do, because their actions are only harmful in light of what the other members do, hence the individual obligations are not unconditional. This leads us to an interesting point. Because if intentions can affect the permissibility of an action, the contributory action is wrong regardless of one's causal role. In order for us to surround the castle together each member must we-intend to surround a castle to starve people to death. Clearly, this is not a permissible action according to any ethical theory. Hence, each agent ought not to partake in surrounding the castle (irrespective of what others do). However, I take it that not all share this view, so suppose we accept that one contributory action by itself is not wrong, because it does not cause harm by itself.

An agent can only tell whether or not he ought to perform his part, if he considers the group perspective. One might think that all collective duties are primitive in the sense that they explain why individuals have certain individual duties (Wringe, 2016; Schwenkenbecher, 2019). It is only in conjunction with the actions of others that one ought not to stand at a bridge. But the primitiveness of a 'plural ought' does not tell us necessarily anything about whether this

On Collective Responsibility and Collective Duties

'plural ought' is distributive or non-distributive. Shared duties can have this 'collective flavor' as well.

Consider a situation in which agents must reason from the group's perspective in order to find out what they individually ought to do. Suppose I am pouring a substance in a lake and you are pouring a substance in the same lake, by itself the contributions are harmless. After some time, we both find out that together the substances cause the fish population in the lake to die. Only by considering our actions from the group perspective do we find out that we ought not to pour the substances in the lake together. We only understand our individual requirements by looking at what we together ought not to do. But this does not make the requirements necessarily non-distributive, only interdependent.²⁶ The content of what is required of the group is distributable over the individual obligations. What is essentially required of the group is this: one of you two ought to refrain from pouring your substance in the lake. This means that I have an *all-things-considered* obligation to refrain from pouring my substance in the lake only if you pour your substance in the lake. And you have an *all-things-considered* obligation to refrain from pouring your substance in the lake only if I pour my substance in the lake.

What is essential here is that my ability to fulfill my obligation is not conditional on you doing your part. Instead it is my *all-things-considered* obligation that is conditional on your action, or better yet, on you *continuing* your action. Once I am in a position to recognize the harm we are doing together, at every single instance that you continue to pour your substance in the lake, I am violating my obligation when I continue pouring my substance in the lake.²⁷ This is importantly different from cases where what is required of the group is a shared action, because here the unwillingness of others may (possibly) excuse my inaction, as I cannot perform a shared action alone.

We can provide a similar but somewhat more complicated answer to Starvation: Each member ought not to perform their contributory action towards blocking a specific bridge only if the three other bridges are likely to be blocked and there are at least two agents likely performing contributory actions to blocking this particular

²⁶ Elsewhere, I argue that if what is morally required is a shared action, then agents' reasons (and therefore the evidence-subjective duties generated by those reasons) are conditional on each other.

²⁷ As mentioned, on my view moral obligations are evidence relative.

bridge.²⁸ Again, the satisfaction conditions for the discharge of the obligation are completely under an agent's control.²⁹

A collective ought is non-distributive if and only if the content of what is required *all-things-considered* is not explicable in terms of what is required of individuals *all-things-considered*. I am not making any claims about whether this is possible or not. But even if some oughts are non-distributive, this does not entail that all oughts are. One good reason to think so is that refraining from partaking in collective actions does not require any shared abilities. This suggests that there is a relevant difference between wrongful collective omissions and wrongful collective actions.

Thus, in Starvation the mob is collectively responsible for the irreducible wrongdoing, but the agents 'only' violated a conjunction of individual duties. Therefore, Starvation shows Type-Symmetry is false.

4. Collective Responsibility and Non-Divisible Consequences

If the collective act is reductive, then all wrongdoing has been accounted for when we distribute responsibility via the Overlap Principle in Starvation.³⁰ To show that even on a reductive account

²⁸ See Kagan (2011). Starvation is a triggering case. When a threshold of some number of acts of a certain type is reached, this act triggers the morally relevant outcome. Because of this, even if one's act does not necessarily make a difference, there is a *chance* it makes a substantial difference. Note that even if Kagan's solution of focusing on expected utility does not solve all collective harm or benefit problems, it is still applicable in at least some cases. Starvation is such a case.

²⁹ Some may think there is no good reason for a member to stop contributing, because it does not make a difference. But I am not concerned here with what a member ought to do once the collective action commenced, we are concerned with which oughts they are violating. This is only relevant when the agents did not know about the potential harm, otherwise it is ad hoc to take the middle of the harmful activity as the starting point. And even if one's contribution would not make a discernable difference, if the act is clearly wrong, one still ought not to do so. I still ought to refrain from joining in on a shared bank robbery even if it makes no discernable difference to the bank being robbed. The point I am making is that we need not further say that the bank robbers and I have a collective obligation to refrain from robbing the bank to explain what is required of each of us.

³⁰ For example, List and Pettit (2011, p. 33) adopt Bratman's content-account that is reductive in spirit. They think that if there is no group agent,

On Collective Responsibility and Collective Duties

of collective action non-agential groups can be non-distributively collectively responsible without violating a collective duty, I will focus on blameworthy consequences.³¹

Consider the following case.

Lake: Eve, James, and Mark live around Forest Lake, at some distance from each other. Each is using a solvent to regularly clean their boats. Each knows that this solvent harms the lake's ecosystem a little but certainly not enough to cause the fish population to die as the ecosystem sufficiently regenerates. The pollution of two of them is sufficient for the fish to die. If any of them would know that more than one person is regularly cleaning their boats with this solvent, they could infer that this would be detrimental to the fish population. During the monthly assembly at Forest Hall, the question is raised whether anyone has something relevant to share concerning the lake's ecosystem. Eve, James and Mark all remain silent about their pollution out of fear for a fine. None of the three knows about the others' polluting activities. Each continues to clean their boat with the solvent. After some time, the entire fish population dies.³²

No agent intended killing the fish, yet it happened nonetheless as the consequence of Eve, James and Mark's actions. Note that each act is individually harmful, but that the combination of acts results into something much worse. This consequence is non-divisible in terms of individual actions, meaning aspects of the resulting states of affairs cannot be uniquely traced back to individual actions. The two questions arise again: what type of duties were violated and what type of responsibility is needed to fully account for the wrongdoing in Lake.

Consider the obligations first. Here it is important to note the difference between evidence-subjective and objective obligations. Objective duties concern what agents ought to do in light of all facts. Given that factually each act harms the lake's ecosystem a little, we can say that each agent has an objective individual obligation not to pollute the lake. But suppose for a moment that each contributory act by itself is harmless. Would there be an objective shared or

the 'we' in each of our intentions is easily reducible to an enumeration of individual pronouns.

³¹ This argument holds regardless of whether collective actions are reductive or non-reductive.

³² This is a variant of a case discussed in Björnsson (2011).

collective obligation to refrain from polluting the lake? Just as in Starvation, each agent is individually able to refrain from polluting the lake, hence again we can account for this simply in terms of conditional shared duties: Eve has an *all-things-considered* duty to refrain from pouring her solvent in the lake only if either Mark or James pours solvent into the lake. James and Mark have the same sort of conditional shared duty. The individual duties are interdependent, but the content of what is required of the group, that at least two of them should not pollute the lake, is fully explicable in terms of what is required of the individuals.

However, we are concerned with what is required of the agents given their available evidence, because moral responsibility is connected to subjective wrongdoing. One important factor in Lake is each person's ignorance of the others' pollution. Given that the available evidence of Eve, James and Mark does not include the fact that other agents are polluting the lake, there is no reason to think that they even have conditional shared duties not to pollute lake. Instead they can only have violated individual duties not to pollute the lake.³³

Next, who is responsible for what? It is clear that each individual is individually responsible for polluting the lake. Each individual's action showed a serious lack of concern for the lake's ecosystem. The question is whether the group is in some way or another responsible for the demise of the fish population.

Given that the killing of the fish population is a non-divisible consequence, we cannot simply 'chop it up' such that the scopes of individuals contain the parts, because there is no way of assigning any specific part to each contributing agent.³⁴ No individual's action was the sole reason for the extinction of the fish population. Moreover, no agent knew about the others' involvement, hence

³³ Again, for those who reject the relevance of intentions for permissibility, if no individual contributory act is harmful, then we can explain these trigger cases in terms of negative expected utility along the lines of Kagan's proposal.

³⁴ Björnsson (2011) points out for a similar case: The relation of the individuals' actions to the consequence is not one of necessity or difference-making, but they are non-redundant parts of nomically sufficient conditions for effects (see Mackie, 1974). In Lake the contributory actions are pair-wise sufficient for the consequence, each action being a non-redundant part of that. One might be tempted to explain responsibility in these terms (see Braham and van Hees, 2012). However, any such attempt will run into problems with cases of what David Lewis (1986) calls 'causal preemption'. I cannot discuss this in detail here, but I follow Björnsson here.

On Collective Responsibility and Collective Duties

they could not have prevented them from polluting.³⁵ Instead the only plausible causal story to tell about what caused the extinction of the fish involves *them*. So, if the distribution is according to each agent's contribution and role, as with the Portion or Overlap Principle, no individual is responsible for killing the fish.

So, if the agents share responsibility for killing the fish, then it must be according to the Full Scope Principle. Typically, an agent is responsible for the consequences of her actions only if she reasonably speaking could have foreseen and prevented the consequence.

However, each agent has a legitimate excuse. They did not know about the other agents' pollution and none of their individual actions was sufficient for causing the fish population to die. Now, ignorance excuses someone from their subsequent act or consequence only if the agent is not culpable for her ignorance. An agent is culpable for her ignorance of X if and only if the agent could reasonably speaking have taken steps to avoid or remedy her ignorance of X (Fitzpatrick, 2008). However, in this case Eve could reasonably speaking not have avoided *her* ignorance of the fact that her action would contribute to killing the fish. The other agents did not share any information about their pollution. The same goes for James and Mark. It would be unfair to hold any individual accountable for the killing the fish. The individuals are surely individually blameworthy for polluting the lake, but each has a legitimate excuse for killing the fish.

This becomes clearer when we consider the following possibility: Suppose Eve finds out about James' action and knows that her contribution will lead to the demise of the fish population, but she continues to pollute anyway. Eve is not merely blameworthy to a larger degree than James. The absence of an excuse typically affects *what* one is responsible for, because this determines whether the conditions for moral responsibility have been met. So, Eve is not just more to blame than James, she is blameworthy for something else entirely, namely killing the fish population. This strongly suggests that none of the individuals is blameworthy for the non-divisible consequence.

This corresponds with our earlier discussion about the interconnectedness of a group, because, after all, the Full Scope Principle is

³⁵ Some may think that speaking up at the assembly still counts as having sufficient interpersonal influence in order to be responsible for the entire group harm. But this can easily be avoided by adding n members to the case and increasing the required number of contributory acts for the fish population to die. Speaking up would no longer make a difference and no agent would have sufficient interpersonal influence.

justified only when the group action or outcome reflects our agency. But here nothing for the relevant sort of interconnectedness seems to be in place. There was no collective intention or common knowledge. No individual action was sufficient for the outcome. None of the agents' individual motivational states would explain the consequence in a normal way. And the agents did not have any sufficient interpersonal influence to prevent the outcome. The relevance of this becomes clearer when we scale up the case and add a n-number of agents. The further the consequence is away from my contributory act, the less sensible it becomes to think an agent is responsible for the non-divisible consequence.

Let me emphasize again that I am not making any claims about the degree of responsibility. Michael Zimmerman (1985) has argued that simply adding culprits cannot decrease one's degree of responsibility. This remains true, although we may add the qualification that all else must remain equal in the relevant sense. All I claim here is that the individual scopes do not include the death of the fish population.

One might think that the presence of individual excuses must imply then that the group is also excused. But it is not clear that the individual excuses get *them* off the hook. Although, no individual could reasonably be expected to have prevented his or her ignorance, *they* could reasonably speaking have prevented *their* ignorance. They had ample opportunity to speak up about their polluting activities but remained quiet out of fear for a fine. Therefore, *they* are culpable for *their* ignorance. Suppose a park ranger expresses his indignation towards the three for killing the fish population. Is this indignation somehow misplaced or unjustified? No, this would be highly counterintuitive. After all, it was *their* combined lack of concern that led to the demise of the fish. They caused this, they could reasonably speaking have been expected to avoid their ignorance and therefore they could have prevented the extinction of the fish population. Therefore, *they* are accountable for the extinction of the fish population. The group's blameworthiness is purely collective.

Hence, the group is collectively responsible. They 'only' violated individual obligations, but this turned into something much worse than mere pollution. The non-divisible consequence of their actions makes the wrongdoing in a sense irreducible. This holds likewise true for similar overdetermination cases without this ignorance component. The members are not responsible for the overall group harm for two reasons. First, the Full Scope Principle is not applicable, and no individual had sufficient interpersonal influence to prevent the outcome nor were their actions sufficient to cause the harm. This holds especially true for large-scale unstructured

On Collective Responsibility and Collective Duties

overdetermined group harms. Second, they each have an individual excuse for their ignorance of the consequence of their action. Because of this ignorance component, the agents can only have individual obligations. Thus, a non-agential group can be collectively responsible without thereby violating a collective duty, meaning Type-Symmetry is false.

Before explaining how Starvation and Lake can be generalized, I consider one final objection.

5. Reconciling Collective Responsibility with Act-Evaluation

Perhaps a final way to defend Type-Symmetry is to argue that non-agential groups cannot be collectively responsible, because this has such unpalatable consequences that we must reject it even if it is counterintuitive. And removing the agency condition from responsibility seems to be precisely that, unpalatable. In the introduction I asked the reader to reserve their skepticism. In the preceding sections I have effectively pried apart responsibility from agency. Let me explain why this is not a problem for any comprehensive ethical theory that tries to account for individual responsibility as well.

Consider again a typical blameworthy act: If at t_0 , moral agent A has a duty not to ϕ , and next at t_1 , A ϕ 's without a legitimate excuse, then A is morally responsible for ϕ -ing. When a purposive or unstructured group is collectively responsible, note that the members of the non-agential group do satisfy the moral agency principle at t_0 . The members of the group have individual duties at t_0 , but the responsibility is transformed by the action or outcome that constitutes wrongdoing at t_1 . The idea is that the agency condition is not inextricably linked with responsibility but conceptually precedes even obligation. Thus, the general ability of normative competence is only *directly* connected with having moral obligations. Whereas moral responsibility is directly connected to wrongdoing, which presupposes moral obligations and agency, thus normative competence is conceptually only *indirectly* connected to moral responsibility. Because of this, collective actions or non-divisible consequences may transform wrongdoing into something irreducible that is only imputable to the group thereby transforming the blameworthiness into something irreducible as well. The group is responsible for the wrongdoing. The reactive attitude directed towards the group for the irreducible wrongdoing implicitly expresses a plurality of demands in order to signal the moral significance of the combined violation of these individual duties by individual moral agents.

This is how collective responsibility becomes disentangled from collective agency. The collective action or consequence in a sense distorts this connection between agency and responsibility.

Importantly, note that I have not removed the agency condition but merely showed where the agency condition conceptually resides: prior to wrongdoing and prior to moral obligations. So, I have disentangled responsibility and agency, but I have not removed the connection entirely. That, I agree, would be absurd.

We never notice this at the individual level because whenever there is responsibility, there is a moral agent. But moral agency is conceptually prior to wrongdoing and the subsequent blameworthiness, as the agent must have had a moral imperative, and only moral agents can have moral imperatives. Thus, responsibility presupposes moral wrongdoing, meaning performing an action thereby violating a moral duty, which entails one must have had a duty prior to action. Having a duty requires one to be a moral agent capable of understanding and processing moral reasons. At the individual level there is never a disconnect between responsibility, wrongdoing, duty and the moral agent. This makes it seem as if responsibility is inextricably linked with moral agency, but there is a step in between. Individual moral wrongdoing always presupposes an individual moral agent as only moral agents can have moral obligations. Therefore, individual blameworthiness necessarily entails individual moral agency. However, at the collective level, irreducible collective wrongdoing does not always entail a single (collective) agent, therefore collective blameworthiness does not necessarily entail collective moral agency. Thus, in the view I have put forward, there is no need to think that this would result in an unpalatable consequence, because at the individual level these four concepts never come apart. At the collective level, collective wrongdoing can sever that link.

6. Generalizing Starvation and Lake

Let us take stock. If actions are non-reductive, purposive groups can be collectively responsible for their irreducible wrongdoing. Consider various kinds of riots. Or far worse still, think of the Rwandan genocide perpetrated by the Hutu against the Tutsi in 1994. Over the course of three months, a large group of Hutus killed around 800,000 Tutsis (and moderate Hutus). Although certain organized collectives were involved, this purposive group cannot be said to be a collective agent, as it comprised many incited civilians as well (Isaacs, 2011, pp. 3–7). This group lacked a

On Collective Responsibility and Collective Duties

decision-making mechanism or a 'master plan' but was evidently capable of horrific collective action (May, 2010, p. 9). The cohesiveness of this purposive group strongly suggests that (most) members did not act for privately endorsed reasons, but collectively accepted a reprehensible and appalling collective goal as that of the group and acted in light of that shared goal. But should we think that each participating Hutu is blameworthy for the entire genocide? No, *they* are blameworthy for the genocide, whereas each individual is blameworthy for their part.

Next, there are numerous non-divisible outcomes in our world that constitute wrongdoing somehow. For example, think of pollution, deforestation, or overfishing. The aggregation of the individual 'contributing' acts do not capture the moral significance of the consequence of the combination of these acts. A plausible explanation is that the non-divisible consequence transforms the wrongdoing into something irreducible. This consequence as a whole cannot feature in the scope of individuals. The non-divisible consequence is only imputable to the group.

Consider Joel Feinberg's (1968) example of racism in the post-bellum South of the US. The majority of Southern white people at that time had the despicable racist belief that African Americans were inferior and less deserving of equal social status and benefits. This belief of racial superiority resulted in prejudice, bigotry, lynchings, and ultimately the oppression and domination of a social group. The racist takes the participatory intentions of others to be a contingent *de facto* ground for his own, which essentially functions to propel what we might call *weak collective action*. The racist still relies on others to have the same belief. Individuals perform actions in light of the group's belief, for example, by expressing racial slurs towards African Americans or amongst each other, excluding African-Americans from public goods or services open to all, or 'contribute' in different ways to the domination and oppression of African-Americans. The domination and oppression of a social group is a non-divisible outcome.

Does it make sense in such cases to think that the responsibility is shared? Note that if the Portion or Overlap Principle is applicable, no agent is responsible for the non-divisible outcome of the domination and oppression of a social group. Some may respond: Well, perhaps we need to allow for those who merely accepted their privileged status, material comfort and respectability within that social group to have a lesser share of responsibility (Feinberg, 1968, p. 686). And possibly those that did not share this belief are free from responsibility if they distanced themselves sufficiently or worked to combat

the effects of this belief while remaining within that group. But many Southern white people simply had this racist belief and acted upon it, therefore the Full Scope Principle applies.

However, this case lacks all the features of the joint killing case. There is no collective intention. There is not sufficient interpersonal influence among the members. It is practically impossible for members to stop enough others such that the outcome will be prevented. And there is only population common knowledge of the involvement of others. Luck enters the picture. Other members may take it too far (from the perspective of the perpetrator), and the outcome is not an exact reflection of one's agency. The group is not characterized by a strong interconnectedness. Therefore, at best the Overlap Principle applies. Although we can surely be blameworthy for foreseeable outcomes of our actions, here the distance between the contributory act and the outcome is too great. A racist is to be blamed for a lot. But to include the entire outcome of oppression and domination of a social group in an individual's scope is simply implausible. Hence, Southern whites are collectively responsible for the oppression and domination of African Americans in the post-bellum South.

To see this, consider that typically, when we express our resentment, there is not only an implicit address of a demand in the form of a (set of) standing moral requirement(s), but moreover a demand for a rectifying response from the norm-violator (Walker, 2006, p. 26). Suppose that some racists acknowledge their individual blameworthiness and try to make amends and recompense. They would apologize for their racist beliefs, but that is not what a victim (presumably) would be after. That would be missing the point. This shows us that the moral significance of the resulting state of affairs has not been accounted for. It is the group that is accountable for the non-divisible outcome.

7. Conclusion

I have argued that Type-Symmetry is false. The same type of responsibility does not necessarily presuppose the violation or unfulfillment of the same type of duty. Non-agential groups can be collectively responsible without thereby violating a collective duty. This is an intuitive conclusion as we can make sense of the practice of holding non-agential groups responsible. Typically, reactive attitudes do abate or disappear when we learn that their target was not a moral agent. But this is not the case with non-agential groups, we still

On Collective Responsibility and Collective Duties

think *they* are responsible. This bad act or consequence is *their* fault. The group is the only identifiable entity whose member-actions constituted the relevant wrongdoing. If a third-party would blame the group, because this is the only entity to be blamed, do we think this third-party is treating the group unfair? No, the group is made up by the agents whose actions constitute the wrongdoing. Blame is an appropriate way of expressing this. The agents that make up that group owe something in response, for example they ought to morally acknowledge legitimate criticism, apologize, and make amends. But it is not sufficient for agents to do so merely for their own actions, they must do so in the realization that their actions were part of something bigger. They must acknowledge collective wrongdoing and group fault. The fact that a group is collectively responsible certainly should result in a change of behavior of its members. Though the group itself is not an agent, its members certainly can act and change. I will say more about this in further work, but for now I hope I have made clear how non-agential groups can be responsible and why blaming such groups can be justified.³⁶

References

- Sean Aas, 'Distributing Collective Obligation', *Journal of Ethics & Social Philosophy*, 9 (2015), 1–23.
- Gunnar Björnsson, 'Joint Responsibility Without Individual Control: Applying the Explanation Hypothesis', in Jeroen van den Hoven,

³⁶ I want to thank first and foremost Hans Bernhard Schmid and David Schweikard for extensive discussions and comments on earlier drafts. I also want to thank an anonymous referee of *Philosophy* for very helpful comments that helped to further improve the paper. I am further indebted to Franz Altner, Gunnar Björnsson, Olle Blomberg, Stephanie Collins, Mattias Gunnemyr, Frank Hindriks, Carlos Núñez, Grace Patterson, Herlinde Pauer-Studer, Matt Rachar, Michael Schmitz, Anne Schwenkenbecher, Vitor Somavilla, Leo Townsend, Elisabeth Widmer and Bill Wringe for helpful comments. Finally, I also want to thank the audiences of conferences where I presented earlier versions of this paper: 'Shared Commitment in Crisis: Social Ontology, Engagement and Politics' at the University of Belgrade; 'Social Ontology 2018' at Tufts University; 'MANCEPT 2018: The Shades of Political Responsibility' at the University of Manchester; 'Group Agency and Collective Responsibility' at the University of Vienna; 'Social Ontology 2019' at Tampere University; and 'MANCEPT 2019: Collective and Shared Responsibility' at the University of Manchester.

Niels de Haan

- Ibo van de Poel & Nicole Vincent (eds.), *Moral Responsibility: Beyond Free Will and Determinism* (Dordrecht: Springer, 2011), 181–200.
- Gunnar Björnsson, 'Essentially Shared Obligations', *Midwest Studies in Philosophy*, 38 (2014), 103–120.
- Matthew Braham and Martin van Hees, 'An Anatomy of Moral Responsibility', *Mind*, 121 (2012), 601–634.
- Michael Bratman, *Shared Agency: A Planning Theory of Acting Together* (Oxford: Oxford University Press, 2014).
- Peter Cane, *Responsibility in Law and Morality* (Portland, Hart Publishing, 2002).
- Sara Rachel Chant, 'Collective Responsibility in a Hollywood Standoff', *Thought*, 4 (2015), 83–92.
- Stephanie Collins, 'Collective Responsibility Gaps', *Journal of Business Ethics*, 154 (2019a), 942–53.
- Stephanie Collins, *Group Duties: Their Existence and Their Implications for Individuals* (Oxford: Oxford University Press, 2019b).
- Stephanie Collins and Niels de Haan, 'Interconnected Blameworthiness', *The Monist*, Forthcoming.
- David Copp, 'On the Agency of Certain Collective Entities: An Argument from Normative Autonomy', *Midwest Studies in Philosophy*, 30 (2006), 194–221.
- David Copp, 'Collective Moral Autonomy Thesis', *Journal of Social Philosophy*, 38, 3 (2007), 369–88.
- Stephen Darwall, *The Second-Personal Standpoint* (Cambridge: Harvard University Press, 2006).
- Toni Erskine, 'Assigning Responsibilities to Institutional Moral Agents: The Case of States and Quasi-States', *Ethics and International Affairs*, 15 (2001), 67–85.
- Joel Feinberg, 'Collective Responsibility', *The Journal of Philosophy*, 65 (1968), 674–88.
- William J. FitzPatrick, 'Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge', *Ethics*, 118 (2008), 589–613.
- Peter French, *Collective and Corporate Responsibility* (New York: Columbia University Press, 1984).
- Margaret Gilbert, 'Who's to Blame? Collective Moral Responsibility and Its Implications for Group Members', *Midwest Studies in Philosophy*, 30 (2006), 94–114.
- Margaret Gilbert, *Joint Commitment: How We Make the Social World* (Oxford: Oxford University Press, 2014).
- Margaret Gilbert and Maura Priest, 'Collective Moral Responsibility and What Follows for Group Members', in Saba Bazargan-Forward and Deborah Tollefsen (eds.), *The Handbook of Collective Responsibility* (New York: Routledge, 2020), 23–37.
- Peter A. Graham, 'In Defense of Objectivism about Moral Obligation', *Ethics*, 121 (2010), 88–115.

On Collective Responsibility and Collective Duties

- Ishtiyaque Haji, 'Alternative Possibilities, Moral Obligation and Moral Responsibility', *Philosophical Papers*, 22 (1993), 41–50.
- Virginia Held, 'Can a Random Collection of Individuals Be Morally Responsible', *The Journal of Philosophy*, 67 (1970), 471–81.
- Kendy Hess, 'Because They Can: The Basis for the Moral Obligations of Certain Collectives', *Midwest Studies in Philosophy*, 30 (2006), 203–221.
- Frank Hindriks, 'How Autonomous Are Collective Agents? Corporate Rights and Normative Individualism', *Erkenntnis*, 79 (2014), 1565–85.
- Frank Hindriks, 'Collective Agency: Moral and Amoral,' *Dialectica*, 72 (2018), 3–23.
- Tracey Isaacs, *Moral Responsibility in Collective Contexts* (Oxford: Oxford University Press, 2011).
- Frank Jackson, 'A Probabilistic Approach to Moral Responsibility', in Ruth Barcan-Marcus, Paul Weingartner, G.J.W. Dorn (eds.), *Logic, Methodology and Philosophy of Science VII* (Amsterdam: Elsevier 1986), 351–65.
- Frank Jackson, 'Group Morality', in Philip Pettit, Richard Sylvan and Jean Norman (eds.), *Metaphysics and Morality: Essays in Honour of J.J.C. Smart* (Oxford: Blackwell, 1987), 92–110.
- Shelly Kagan, 'Do I Make a Difference?', *Philosophy & Public Affairs*, 39 (2011), 105–41.
- Christopher Kutz, *Complicity: Ethics and Law for a Collective Age* (Cambridge: Cambridge University Press, 2000).
- Holly Lawford-Smith, 'What We?', *Journal of Social Ontology*, 1 (2015): 225–249.
- David Lewis, 'Postscript to "Causation"', *Philosophical Papers*, Vol.II (New York: Oxford University Press, 1986), 172–213.
- H.D. Lewis, 'Collective Responsibility', *Philosophy* 23 (1948): 3–18.
- Christian List and Philip Pettit, *Group Agency: The Possibility, Design, and Status of Corporate Agents* (Oxford: Oxford University Press, 2011).
- Kirk Ludwig, *From Individual to Plural Agency* (Oxford: Oxford University Press, 2016).
- John Mackie, *The Cement of the Universe* (Gloucestershire: Clarendon Press, 1974)
- Larry May, *The Morality of Groups* (Notre Dame: University of Notre Dame Press, 1989).
- Larry May, 'Collective Inaction and Shared Responsibility', *Nous*, 24 (1990): 273.
- Larry May, *Sharing Responsibility* (Chicago: University of Chicago Press, 1992).
- Larry May, *Genocide: A Normative Account* (Cambridge: Cambridge University Press, 2010).
- Michael McKenna, 'Where Frankfurt and Strawson Meet', *Midwest Studies in Philosophy*, 29 (2005): 163–180.
- Derek Parfit, *Reasons and Persons* (Oxford: Clarendon Press 1984).
- Philip Pettit, 'Responsibility Incorporated', *Ethics*, 117 (2007), 171–201.

Niels de Haan

- Philip Pettit, 'Corporate Responsibility Revisited', *Netherlands Journal of Legal Philosophy*, 38 (2009), 159–76.
- Philip Pettit and David Schweikard, 'Joint Actions and Group Agents', *Philosophy of the Social Sciences* 36 (2006), 18–39.
- Adina Preda, 'Group Rights and Group Agency', *Journal of Moral Philosophy*, 9 (2012), 229–54.
- Felix Pinkert, 'What We Together Can Be Required to Do', *Midwest Studies in Philosophy*, 38 (2014), 187–202.
- W. David Ross, *Foundations of Ethics* (London: Oxford University Press, 1939):
- T.M. Scanlon, 'Forms and Conditions of Responsibility', in Robert Clarke, Michael McKenna and Angela M. Smith (eds.), *The Nature of Moral Responsibility: New Essays* (New York: OUP, 2015), 89–114.
- Hans Bernhard Schmid, 'Plural Self-Awareness', *Phenomenology and the Cognitive Sciences*, 13 (2014), 7–24.
- Hans Bernhard Schmid, 'Collective Responsibility of Random Collections: Plural Self-Awareness among Strangers', *Journal of Social Philosophy*, 49 (2018): 91–105.
- David Schweikard and Hans Bernhard Schmid, 'Collective Intentionality', in Edward Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2013 Edition), URL= <<https://plato.stanford.edu/archives/sum2013/entries/collective-intentionality/>>.
- Anne Schwenkenbecher, 'Joint Duties and Global Moral Obligations', *Ratio*, 26 (2013), 310–328.
- Anne Schwenkenbecher, 'Joint Moral Duties', *Midwest Studies in Philosophy*, 38 (2014), 58–74.
- Anne Schwenkenbecher, 'Making Sense of Collective Obligations', in (eds.) Kendy Hess, Violetta Igneski and Tracey Isaacs, *Collectivity: Ontology, Ethics and Social Justice* (London: Rowman & Littlefield 2018), 109–132.
- Anne Schwenkenbecher, 'Collective Moral Obligations: We-Reasoning and the Perspective of the Deliberating Agent', *The Monist*, 102 (2019), 151–71.
- John R. Searle, 'Collective Intentions and Actions', in Philip Cohen, Jerry Morgan, Martha Pollack (eds.), *Intentions in Communication* (Cambridge: The MIT Press, 1990), 401–16.
- Amy Sepinwall, 'Corporate Moral Responsibility', *Philosophy Compass*, 11 (2016), 3–13.
- David Shoemaker, 'Attributability, Answerability and Accountability: Toward a Wider Theory of Moral Responsibility', *Ethics*, 121 (2011), 602–632.
- John Skorupski, *The Domain of Reasons* (New York: OUP 2010).
- Angela M. Smith, 'On Being Responsible and Holding Responsible', *The Journal of Ethics*, 11 (2007), 465–484.
- Peter F. Strawson, 'Freedom and Resentment', *Proceedings of the British Academy*, 48 (1962), 187–211.

On Collective Responsibility and Collective Duties

- Raimo Tuomela, *Social Ontology: Collective Intentionality and Group Agents* (Oxford: Oxford University Press, 2013).
- Raimo Tuomela, 'Non-Reductive Views of Shared Intention', in Marija Jankovic and Kirk Ludwig (eds.), *The Routledge Handbook of Collective Intentionality* (New York: Routledge, 2017), 25–33.
- Raimo Tuomela and Pekka Mäkelä, 'A We-Mode Account of Group Action and Group Responsibility', in Saba Bazargan-Forward and Deborah Tollefsen (eds.), *The Handbook of Collective Responsibility* (New York: Routledge, 2020), 65–77.
- Margaret Urban Walker, *Moral Repair: Reconstructing Moral Relations after Wrongdoing* (New York: Cambridge University Press, 2006).
- R. Jay Wallace, *Responsibility and the Moral Sentiments* (Cambridge: Harvard University Press, 1994).
- Gary Watson, 'Two Faces of Responsibility', *Philosophical Topics*, 24 (1996), 227–48.
- David Widerker, 'Frankfurt on 'Ought Implies Can' and Alternative Possibilities', *Analysis*, 51 (1991): 222–4.
- Bill Wringe, 'From Global Collective Obligations to Institutional Obligations', *Midwest Studies in Philosophy*, 38 (2014), 171–186.
- Bill Wringe, 'Collective Obligations: Their Existence, Their Explanatory Power, and Their Supervenience on the Obligations of Individuals', *European Journal of Philosophy*, 24 (2016), 472–497.
- Michael Zimmerman, 'Sharing Responsibility', *American Philosophical Quarterly*, 22 (1985), 115–122.
- Michael Zimmerman, 'Taking Luck Seriously', *The Journal of Philosophy*, 99 (2002), 553–76.
- Michael Zimmerman, *Living With Uncertainty: The Moral Significance of Ignorance* (Cambridge: Cambridge University Press, 2008).
- NIELS DE HAAN (ndehaanphilosophy@gmail.com / niels.de.haan@univie.ac.at) currently works as a prae doc at the University of Vienna. His research interests are ethics, philosophy of action, political philosophy, and social ontology.