

Interconnected Blameworthiness

Stephanie Collins and Niels de Haan

The final and definitive version of this paper will be published in The Monist, special issue on 'Vicarious Responsibility and Circumstantial Luck'.

Abstract

This paper investigates agents' blameworthiness when they are part of a group that does harm. We analyse three factors that affect the scope of an agent's blameworthiness in these cases: shared intentionality, interpersonal influence, and common knowledge. Each factor involves circumstantial (and some resultant) luck. The more each factor is present, the greater is the scope of each agent's vicarious blameworthiness for the other agents' contributions to the harm. We then consider an agent's degree of blameworthiness, as distinct from her scope of blameworthiness. We suggest that an agent mostly controls her degree of blameworthiness—but even here, luck constrains what possible degrees of blameworthiness are open to her.

1. Introduction

It is commonplace to view groups as blameworthy: corporations are blameworthy for disasters caused by reckless company cultures; rioting mobs are blameworthy for property damage; humanity is blameworthy for global warming. Often, group blameworthiness has implications for members: individual CEOs, rioters, and humans are blameworthy for (their role in) group harms. In this paper, we conceptualise the factors that influence members' blameworthiness over wrongful group outcomes, focusing on the role of circumstantial luck.

Section 2 distinguishes corporate, distributed, collective, and shared group blameworthiness. Corporate and collective blameworthiness attach to the group non-distributively. Corporate blameworthiness attaches to structured groups, while collective blameworthiness attaches to unstructured groups. Distributed and shared blameworthiness attach to group members. Distributed blameworthiness attaches to members of structured groups; shared blameworthiness attaches to members of unstructured groups.

We ask how to apportion distributed and shared blameworthiness. The trend in the collective responsibility literature has been to sharply distinguish structured from unstructured groups (e.g., Held 1970; French 1984; List and Pettit 2011; Collins 2019). Thus, one of our contributions is to reveal that some underlying luck-related factors apply to *both* distributed blameworthiness (structured groups) *and* shared blameworthiness (unstructured groups).¹

How should we apportion distributed and shared blameworthiness? Section 3 distinguishes three principles: each agent might be responsible only for her own contribution, or each might also be responsible for others' contributions (such that there is 'overlap' between members' responsibilities: both you and I might be responsible for what I do), or each might be responsible for the entire group outcome. Both the second and third principles involve 'vicarious' responsibility: I am responsible for (at least part of) what you cause, and vice versa. We examine Christopher Kutz's (2000) argument for the third principle, arguing that it sometimes makes agents blameworthy for too much.

¹ There may nonetheless be differences between structured and unstructured groups. Herlinde Pauer-Studer (2018) argues that an agent's complicity might be due to the dependence of his professional role on the normative principles that make up the organisation in whose practices he partakes.

In Section 4, we argue that the level of ‘interconnectedness’ within the group determines which principle applies. Interconnectedness has three elements: shared intentionality, interpersonal influence, and common knowledge. We analyse each element, demonstrating the role of circumstantial (and some resultant) luck.

This all concerns scope of blameworthiness—*what* one is blameworthy for. Following Michael Zimmermann (2002, 560), we distinguish scope from degree—*how much* one is blameworthy for what one is blameworthy for. Section 5 turns to degree. One might think an agent’s degree of blameworthiness is unaffected by luck—even if scope is affected by luck, as Section 4 argues. Yet luck affects one’s degree of blameworthiness in group harm cases—though only indirectly and unproblematically, as we’ll explicate.

Before we begin, a caveat. Blameworthiness is not the only species of responsibility. Rather than blaming individuals involved in group-level harms, perhaps individuals simply have obligations because of those harms, or perhaps individuals should acknowledge their agential involvement without receiving blame (Pickard and Ward 2013; Lacey and Pickard 2013). This is appropriate in some—perhaps many—cases. But in this paper, we restrict our focus to blameworthiness traditionally conceived.

2. Varieties of Group Blameworthiness

Our argument covers structured and unstructured groups. Structured groups satisfy conditions for group agency. Following Christian List and Philip Pettit’s (2011) influential account,² group agents have an organizational structure that, *inter alia*, includes a decision-making procedure enabling the group robustly to hold (functional equivalents of) representational and

² For alternatives, see French (1984), Rovane (1998), Hindriks (2018).

motivational states while satisfying desiderata of rationality. These representational and motivational states are modelled as attitudes towards propositions.

For example, BP holds (the functional equivalent of) the representational state ‘Lowering safety standards will increase profits’ just in case BP behaves as if this is the case and is disposed to assent to propositions entailed by that proposition. BP holds (the functional equivalent of) the motivational state ‘BP having high profits’ just in case BP behaves in ways likely to raise BP’s profits, given its representational states. BP satisfies desiderata of rationality if its representational and motivational states interlock in a mutually supportive fashion. Finally, for BP to be apt for blame, it must (minimally) have the capability to understand and process moral reasons and act accordingly (List and Pettit 2011, ch. 7; Hindriks 2018; Collins 2019, ch. 6). We’ll use ‘structured group’ to refer to group agents with this capability.

What about unstructured groups? They are any collection of agents that don’t together constitute a group agent. There are many sorts of unstructured groups, from mobs to social movements to millions of people unintentionally causing climate change. For many such groups, there are predicates that are true of the group without being true of any member. Such predicates are “purely collective” (French 1984, 21; Pinkert 2014). Consider ‘they surrounded the building.’ ‘Surrounded the building’ is a predicate that cannot belong to any individual, because no individual can surround a building. The predicate can belong only to the group. If the building is on fire, then it may be morally significant that the unstructured group surrounds the building, namely when they prevent the occupants from escaping.

Suppose BP lowers its safety standards and the mob surrounds the burning building. This does not imply that any member of BP has lowered their individual safety standards, or that any mob member has surrounded the building. Those are purely collective predicates. It is at least conceptually possible that blameworthiness is similar: perhaps, the group is

blameworthy, without any member being blameworthy. Philosophers have argued that unstructured groups can be blameworthy without members being blameworthy (e.g., Feinberg 1968; Chant 2015; cf. Lawford-Smith 2015). Others have argued likewise for structured groups (French 1984, 5-13; Pettit 2007; Copp 2007). We take these conclusions to be sound. We assume group-level blameworthiness is a different level from individual-level blameworthiness, where the levels don't entail or preclude each other (Isaacs 2011).

This produces four categories of blameworthiness, presented in Table 1.³

	Unstructured Group	Structured Group
Individual-level Blameworthiness	Shared Blameworthiness	Distributed Blameworthiness
Group-level Blameworthiness	Collective Blameworthiness	Corporate Blameworthiness

Table 1: Varieties of Group Blameworthiness

We emphasise the possibility of group-level blameworthiness—both collective and corporate—because, in Section 4, we will constrain individual-level blameworthiness along three dimensions. This may appear to produce too little blameworthiness. Importantly, then, collective and corporate blameworthiness are still in the picture, though we will not discuss them further. We ask: what determines the appropriate allocation of individual-level blameworthiness? What role does luck play in that allocation?

³ Preda (2012) makes a similar suggestion.

3. Allocative Principles and Kutz's Complicity

As mentioned, following Zimmerman (2002, 560), we distinguish *scope* from *degree* of blameworthiness. To illustrate, suppose I break a promise to you. My scope of blameworthiness might be 'breaking the promise.' While keeping this scope fixed, my degree of blameworthiness varies based on whether I could have foreseen (when making the promise) that I would break it, whether I intended to break it, and so on. Zimmerman uses this distinction to constrain luck: one's scope might (un)luckily be beyond one's control, but one's degree is always within one's control (in emphasizing control, we follow Nelkin's (2019) characterization of moral luck). We discuss degree in Section 5; for now, we focus on scope: *what* are individuals blameworthy for, regarding group harms, and what role does luck play?

Our discussion subsumes circumstantial and resultant luck. Circumstantial luck concerns the circumstances in which an agent makes decisions and acts. Consider Sophie in *Sophie's Choice*, who is told by a Nazi guard that she must give one of her two children to the guard: she is unlucky in the choice she faces.⁴ Resultant luck is luck in the consequences of an agent's choices. Consider Bernard Williams' (1981) lorry driver, who faultlessly hits a small child: he is unlucky in the consequences of his choice. In both cases, the agent feels remorse and guilt: the hallmarks of first-person blame. We accommodate this by saying the deaths fall within the *scope* of what the agents can be assessed as blameworthy for, although

⁴ In Lillehammer's (forthcoming) terms, we discuss cases of "redemptive" and "choice-dependent" circumstantial moral luck: cases where the agent can redeem themselves partly through their choices, even though they are unlucky in facing a morally revealing choice.

these agents have zero *degree* of blameworthiness. Does luck similarly affect agents' scope in group harm cases?

Let's consider the options. In group harm cases, there are three possible principles for agents' scope of blameworthiness:

Portion Principle: Each agent is blameworthy for a different portion of what the group is blameworthy for, with no overlap between agents' portions. (Lewis 1948)

Overlap Principle: Each agent is blameworthy for a portion of what the group is blameworthy for, with some overlap between agents' portions. (May 1992)

Full Scope Principle: Each agent is blameworthy for the totality of what the group is blameworthy for. There is complete overlap between agents' portions. (Kutz 2000)

The Overlap and Full Scope Principles imply *vicarious* responsibility: each is blameworthy for (at least part of) what others cause.⁵ In the rest of this section, we critique Christopher Kutz's well-known defence of the Full Scope Principle. This will pave the way for our more nuanced account.

Kutz's argument centres on 'complicity.' One is complicit when one is involved in morally wrongful activity with or of others. Kutz proposes the following 'Complicity Principle', the second half of which is equivalent to our Full Scope Principle:

⁵ Yet unlike some discussions of vicarious responsibility, we discuss only cases where agents produce harm *together*, not where one agent is the fiduciary, proxy, or principal of the other.

(Basis) I am accountable for what others do when I intentionally participate in the wrong they do or the harm they cause. (Object) I am accountable for the harm or wrong we do together, independently of the actual difference I make. (2000, 122)

According to Kutz, when individuals “intentionally participate” in a wrong or harm, they each intentionally do their part of what they conceive as a collective project, where their conceptions of that project sufficiently overlap. Kutz thereby distinguishes ‘exclusive’ and ‘inclusive’ authorship. When I act alone, I am the exclusive author of that action and events caused by it. Exclusive authorship involves a causal relation between intention and outcome. But when I and others act collectively, I am an inclusive author of the group’s actions, because I am among several who can truly say ‘we did it’ (Kutz 2000, 139). The causal relation between my intention and the outcome doesn’t matter; what matters is the teleological relation: I aim at the collective outcome, even if I do not cause it. We can ascribe such outcomes to both the group (in our typology, producing collective or corporate blameworthiness) and to the members (producing shared or distributed blameworthiness).

Kutz places two caveats on the Complicity Principle. First, members’ actions that are outside any plausible refinement of the shared goal cannot be ascribed to the (structured) group or other members. Second, the same holds for unintended consequences that flow from behaviour “beyond the pale of any reasonable collective expectation” (Kutz 2000, 155). Kutz further acknowledges that agents can be accountable to different degrees, especially when participants are unaware of the nature of the harm, or when some would prefer that aspects of the shared project go unrealized (2000, 165).

In our view, luck shows us something important about the appropriateness of the allocative principles. Although Kutz does not discuss luck, the Complicity Principle implies that luck plays an important role in agents’ accountability. From my position, it is a matter of

circumstantial luck that you aimed at the collective end, and that you did so with your particular dispositions (regarding recklessness and so on). Yet Kutz's *defence* of the Complicity Principle curiously eschews any role for luck:

We are properly held accountable for the actions of groups (and of individual group members) in which we participate, because these actions represent our own conception of our agency and our projects. This conception, embedded in our participatory action, is thoroughly normative: it expresses what we desire, what we will tolerate and what we believe. If a set of agents' participatory intentions overlap, then the will of each is represented in what each other does qua group member, as well as what they do together. The logical overlap permits us to say that they manifest their attitudes through one another's actions. (2000, 141)

Thus, Kutz justifies the luck-imbued Full Scope Principle via a non-luck (agency-based) justification. This is a deep tension. It produces three problems for Kutz's account.

First, if my blameworthiness is grounded in what I desire, tolerate, believe, accept, and so on, then (surely) for the Full Scope Principle to apply, my attitudes must be fully in sync with the general attitude in the group. Otherwise, the group goal does not fully represent my agency. But what are the chances that any member's desires, tolerances, etc, match perfectly with those of *any* other member? If there is no such matching, then we cannot even say what the general attitudes in the group are (List 2014). Kutz says the attitudes of members must be 'sufficiently' similar, but mere sufficiency of similarity threatens to undermine the connection between my blameworthiness and my agency, which Kutz insists upon. Such an extreme principle as the Full Scope Principle should apply only when the

general attitude of the group is no distortion of each agent's mental states. This can be so only when the group has strong interconnectedness, as we will explain.

Second, for Kutz, my blameworthiness is independent of the difference I make. My intention is what matters. But Kutz ultimately cares about my agency, and there is more to my agency than my intentions. There is also the means I take to pursue those intentions. For example, suppose I have a participatory intention to participate in a mob's looting of the street. In one scenario, I shyly sidle up to the other looters and steal from shops they have already broken into. In another scenario, I run down the street before anyone else, yelling that the street is ripe for looting. The latter is more causally efficacious, by encouraging others' contributions. Yet Kutz gives each agent the same scope of blameworthiness in each case.

Third, for Kutz, a person has the same scope regardless of their knowledge about other participants. Suppose three young-adult siblings intend to protect their youngest sibling from the local adolescent bully. Two siblings reproach the bully verbally. However, the oldest sibling beats up the bully. The other siblings did not know or expect this would happen nor knew that their sibling was capable of this. But beating up the bully is a plausible refinement of protecting one's sibling and a foreseeable outcome (Kutz 2000, 155), just not one the other two siblings would endorse. It's even possible the two siblings had no idea their oldest sibling would get involved. Yet, according to Kutz, the scope of each would be the same (although the degree would not be). Kutz claims that when we act together, we must expect that the group act may have "moral agency costs", that is, aspects we do not know about but with which we will have to reckon (2000, 156). However, a lack of knowledge concerning others' dispositions and actions allows for circumstantial luck to play too big of a role in what agents have to reckon with, again creating a tension with Kutz's agency-based justification for the Full Scope Principle (see also Lawford-Smith 2018, 327-8).

To be clear, luck sometimes rightly affects the scope of agents' blameworthiness. We simply suggest that Kutz pushes luck too far with his widespread application of the Full Scope Principle. We think a more nuanced way of allocating blameworthiness is appropriate. Our three objections correlate with three components of interconnectedness, analysed in the next section. Thus, we agree with Kutz that interpersonal phenomena are important for understanding agents' blameworthiness in collective contexts. However, unlike Kutz, we think more interpersonal factors matter for the allocation of blameworthiness than only participatory intentions. We also agree with Kutz's conclusion that luck partly determines one's scope of blameworthiness (even though he implicitly denies this in his justification). However, we don't think agents are so unlucky as to face the Full Scope Principle as often as Kutz suggests.

4. How Interconnectedness Determines Scope

Overall our proposal is this: the stronger the interconnectedness between agents' intentionality, influence, and knowledge over a group harm, the more each is responsible for the others' contributions as well as their own, such that each has a larger scope and there is more overlap between their scopes. In group harm cases with no interconnectedness, the Portion Principle applies: each is blameworthy for their individual portion, share, or contribution towards the outcome, where these portions do not overlap. In group harm cases with some (but non-maximal) interconnectedness, the Overlap Principle applies. In these cases, there is more overlap in individuals' portions—and, therefore, larger individual portions—if there is more interconnectedness. Finally, in cases with maximal possible interconnectedness, the Full Scope Principle applies. In cases with maximal interconnectedness, then, our account produces the same result as Kutz's: each agent is blameworthy for the entire group harm. Interconnectedness tends to correlate with the

number of agents: usually, the more agents there are, the less interconnectedness there is. But we do not believe the number of agents matters per se, for an individual's scope of blameworthiness.

In this section, we motivate our view by explaining three factors of 'interconnectedness': shared intentionality, interpersonal influence, and common knowledge. These are central factors for interconnectedness, though there may be other factors too. Each factor comes in degrees, and each is social, as the language of 'shared,' 'interpersonal,' and 'common' indicates. This sociality produces vicarious blameworthiness of each agent, for some aspect of the other agents (specifically, those agents' intentions, influence, and knowledge). Circumstantial and resultant luck play a role in determining the extent of each factor. But none of the three factors can exist without each individual exercising *some* agency. And individuals have some control over the extent of each factor. Only if each scope-factor is maximally present, meaning maximal interconnectedness within the group, does the Full Scope Principle apply, because only then is the group action an exact reflection of my agency. We now explain the three factors in turn.

i. Shared Intentionality

Raimo Tuomela's (2013) account of collective intentionality provides a helpful starting point for our first factor, although space prevents us from discussing the intricacies of Tuomela's view.

According to Tuomela, the strongest sense of acting together is "we-mode joint action," which involves members acting for group reasons rather than private reasons. In such cases, there is a *joint intention*: an intention ascribed to individuals collectively and consisting of interdependent "we-mode" "we-intentions" of members. A *we-intention* is a member's intention to participate in the group members' doing X (for example, committing a murder)

together, and to do his part of X (as his part). The *we-mode* implies thinking and acting *as a group member*, that is, from the group's perspective.

In Tuomela's view, we-mode joint action is uniquely able to satisfy three unity-creating criteria: (i) an ideal "we-moder" thinks and acts only for authoritative, group-centered, uniformly motivating reasons, regardless of whether they conflict with individual motives; (ii) a we-mode we-intention is only satisfied when it is satisfied for all members, because each member intends to participate in *the members' doing X together*; and (iii) members are collectively committed to the satisfaction of the group's ethos and performing their parts (Tuomela 2013, 38-46). These features serve to unify the group around its ethos: its constitutive goals, interests, values, beliefs, practices, and/or norms (Tuomela 2007, 32).⁶

Plausibly, sharedness of intentions comes in degrees. The more fully agents approximate strongly shared intentions (as in Tuomela's we-mode joint action with joint intentions), the more of the others' contributions each is blameworthy for. Why? Because each person's agency is more reflected in, and is more a reflection of, the others' agency (at least vis-à-vis the group's ethos). True, others' intentions are partly a matter of circumstantial luck for me. But when I adopt the we-mode, I embrace this role for luck in determining my scope of blameworthiness. This suggests that minimally, the Overlap Principle applies when agents engage in we-mode joint action, and possibly, depending on the other factors, the Full Scope Principle applies.

Next, consider Michael Bratman's influential account of shared intentions:

(1)(a) I intend that we J and (b) you intend that we J and (2) I intend that we J in accordance with and because of (1)(a), (1)(b) and meshing sub-plans of (1)(a), (1)(b),

⁶ In Tuomela's view, we-mode joint action is sufficient for group agency (2013, 112). We use a (somewhat) different functionalist conception of group agency.

and you intend likewise and (3) (1) and (2) are common knowledge between us.

(1999, 121)

Although there is no full-blown group reason or “we-thinking” involved, Bratman’s shared intentions still require significantly more of agents than Kutz’s participatory intentions. The latter do not require meshing sub-plans, or that we have the intentions because each other does, and so on. More so than participatory intentions, Bratmanian shared intentions *reflect the other’s agency*. After all, if I intend that we commit a murder partly because of, and partly via, your sub-plans that serve as the roadmap of your contribution to the murder, then plausibly your plans and sub-plans partly reflect my agency, for the reasons Kutz gives in the earlier quote. Therefore, the more the scope of our blameworthiness overlaps—I am partly to blame for your intention and participatory action, and vice versa. Bratman’s account shows that even if the members are not “we-moders”, there can nonetheless be strong interconnection within a group (but not as strong as in Tuomela’s we-mode joint action).

Bratman’s account is highly demanding (Blomberg 2016). It’s just one of many instances of (what Tuomela would call) “pro-group I-mode” action. Bratman’s account suggests that almost all members of companies do not have shared intentions that their companies lower safety standards, that most of a mob does not share the intention to loot the street, and that members of humanity lack intentions over global warming. As Hess (2018) points out, members of large companies generally just intend to earn a paycheck—they do not intend any of their company’s ends. Similarly, Tuomela (2013, 115) points out that even we-mode joint actions do not always require a joint intention involving all members. A group (e.g. a corporation) can perform a joint action via a subset of members, through which a group action can be attributed to the group.

Even weaker cases involve mere “mutual behavioural dependence” (Tuomela 2013, 103). These collective actions consist of individual actions performed for a shared mutually believed reason. This may involve a shared divided goal and (I-mode) intentions roughly equivalent to Kutz’s “participatory intentions”. Or this may involve strategic action where agents respond optimally to other agents based on a mutual dependence belief (e.g. stockbrokers speculating for profit). The weakest category is “commonality” (Tuomela 2013, 103). Here agents simply have the same shared undivided goal. For example, consider car drivers. Members of this group prefer to get around by car. But the aggregation of such individual actions may have major harmful consequences, for example air pollution. The agents are aware that others act similarly and the shared mutually believed reason is known to imply a lack of concern for the environment.⁷ Because the agents act for the same broad reason, their agency is partly reflected in each other’s actions—but only partly, and only broadly. So, at best the Overlap Principle is appropriate (pending discussion of the other two interconnectedness factors)—but not the Full Scope Principle, as Kutz suggests.⁸

ii. Interpersonal Influence

A second factor is the influence agents try to, and do, have over one another. One might think influence was implicit in, for example, Bratman’s account of shared intentions, which says each agent intends the goal partly *because of* the others sharing it and the others’ sub-plans.

⁷ Plausibly, various collective actions involving different reasons can lead to the same harmful (non-divisible) outcome.

⁸ Some cases of shared intentionality are even more minimal. See Knowles’ (2019) discussion of women’s complicity in their own oppression. Here, the Overlap Principle shades into the Portion Principle.

But interpersonal influence is not built into all accounts of shared intentions. We think influence is separate from intentions, and worth analysing in its own right. That said, we don't adopt any particular account of influence or causation. We use 'influence' as a common-sense notion and encourage the reader to insert their preferred account.

We are concerned not with an agent's *direct* influence over an outcome, but instead the *indirect* influence she has over that outcome *via* her influence over other agents who contribute to that outcome. Only this indirect—or, as we will say, *interpersonal*—influence can explain why the first agent should be partly responsible for what the second agent does, that is, explain why the Overlap Principle should apply (rather than the Portion Principle). To motivate the thought, suppose I want my competitor dead. I persuade you to help me. Here, I have influenced your contribution to the outcome. Yet that contribution is also your own. Both you and I are blameworthy for your contribution. That is, the Overlap Principle applies. By contrast, suppose I did not wield any influence over you, and did not even try to. Yet you performed half of the murder. In that case, your contribution should not be included in my scope (unless we have shared intentions and common knowledge).

Why does interpersonal influence affect scope of blameworthiness? For the same reason that all unexcused contributions to harm influence scope of blameworthiness: unexcused contributions to harm are wrongful. The twist we add is that, in cases where that influence is interpersonal (i.e., exercised via influence over a second agent's contribution), then *both* agents are blameworthy for the second agent's contribution. This brings in some luck: the first agent faces resultant luck in whether the second agent is influenceable; the second agent faces circumstantial luck in whether the first agent tries to exert influence. But each has agency over the influence (over its attempt and its effect, respectively), so the Overlap Principle is appropriate—and the Overlap Principle is more appropriate, the more such influence agents exercise.

Sometimes influence is asymmetric. Consider riot leaders. Their scope of blameworthiness is larger than that of more passive members, because they exercise more influence. This was illustrated in our discussion of Kutz, where we compared the mob follower with the mob initiator. Yet often, influence runs in both directions. It can even be perfectly symmetrical: you and I might equally influence the other to continue in our murderous scheme. This gives us equal responsibility over each other's actions. Importantly, influence includes counterfactuals: if I could, reasonably speaking, have influenced someone *not* to perform their behaviour that contributes to the harm, but I do not exercise that influence, then his actions are in my scope. In this way, a particularly influential individual's scope may include (a large chunk of) the collective action.

iii. Common Knowledge

The final factor is common knowledge. Common knowledge of what? Of each other's shared intentions and interpersonal influence: common knowledge is 'meta' to the other two factors. Like the other factors, there are weaker and stronger variants of common knowledge. Consider David Lewis's (1969) classic conditions: I know that you know I know that you know (etc) that you and I have shared intentions to contribute to the harm and have interpersonal influence over each other's contributions (see also Schiffler 1972 and Tuomela 2007, 74-82; this kind of common knowledge is included in the we-mode). This high level of common knowledge might hold if you and I together plot and execute a murder, with face-to-face contact in the planning and physical assistance of one another in the execution.

Weaker forms of common knowledge include Margaret Gilbert's view about knowledge dispersed throughout a society. Gilbert examines knowledge of joint commitments, saying "It is possible, indeed, that the members of a large population can be jointly committed in some way *without each one knowing of each as an individual.*" (2014,

66, emphasis added) Plausibly, the overall goals of a large company are the objects of population common knowledge in Gilbert's weaker sense. I might not recognize my fellow company-members if I fell over them in the corridor. So it's false that 'I know that *you* know.' Instead, it's 'I know that, in general, people around here know'— 'in general' because there might be non-members in my company's building.

There are even weaker forms of common knowledge, leading to even less overlap between members' scopes of blameworthiness. For example, consider Natalie Gold and Robert Sugden's (2007, 31) definition of 'T-conditional' common knowledge: each member of a group T endorses goal G; each member of T knows there is more than one member of T; each member of T knows that, if some agent A is a member of T, then A endorses G; each member of T knows that, if some agent A is a member of T, then A knows that, if some agent A is a member of T, then A endorses G—and so on. But the members might not know who each other are, or even how many members there are. This is not Lewis's 'I know that you know.' It's not even Gilbert's 'I know that, in general, people around here know.' Instead, it's 'I know that someone, somewhere, knows.'

Why think that common knowledge—of whatever strength—leads to overlapping scopes of blameworthiness? Recall that shared intentions (to participate in harm) give us overlapping scopes because our shared intentions reflect each other's agency. Interpersonal influence gives us overlapping scopes because it means that we make a difference to what each other does. Common knowledge implicates both these rationales: the stronger our common knowledge, the more the object of that knowledge (that is, each other's intentions and influence) reflect one another's agency. And the stronger our common knowledge, the more chance we have of influencing one another. Yet we think it is worth mentioning common knowledge in its own right, because of its 'meta' status: the knowledge is knowledge of shared intentions and interpersonal influence. It operates on the other two

factors, working upon and through them to further implicate the individuals in one another's mental states (via shared intentions) and behaviours (via interpersonal influence). Like the other factors, each agent faces circumstantial luck when it comes to what the other agents (can) know.

5. How Scope Influences Degree

We have proposed three factors that comprise 'interconnectedness' and that evoke the Overlap Principle. When these factors exist, the Portion Principle is insufficient. Yet the factors are often not instantiated to their fullest possible extent, meaning we should resist the Full Scope Principle. This resistance follows from our three objections to Kutz: the Full Scope Principle should not apply if our intentions do not fully mesh, if we do not exert all the influence we can on each other in pursuing the group harm, and if we cannot know what others are willing and doing. The Overlap Principle is attractive in many cases, because it admits of a wide range—from only a small amount of overlap between individuals' scopes, to so much overlap that we approach the Full Scope Principle.

If interconnectedness thus determines individuals' scope of blameworthiness, what about degree of blameworthiness? Following Zimmermann, the scope/degree distinction enables us to say one's scope is partly a matter of luck, while one's degree is under one's control.⁹ We have proposed factors that bring an individual's scope partly under her control and partly a matter of luck. A lack of interconnectedness can lead to a small(er) scope. But this does not necessarily imply a small(er) degree of blameworthiness.

⁹ Degree can be explicated in a coarse-grained (full, partial) or fine-grained (scalar) manner.

We think coarse-grained terms can be understood as particular points on a scale.

To see this, consider an individual case. The degree of an individual's blameworthiness is determined by her action and mental states, including her intentions, goals, beliefs, and knowledge. This may include counterfactuals such as what she could have known. Luck affects the scope without affecting the degree. For example, the scope of a drunk driver who kills someone is different from a drunk driver who does not, though they are blameworthy to the same degree. Similarly, one might be blameworthy to an extremely high degree for one's contribution to a genocide, even though interconnectedness is so low that one is blameworthy only for one's own contribution, and not for any components of others' contributions, such that the Portion Principle applies. Scope and degree are separate.

Yet there is an interesting wrinkle. Scope and degree are not *completely* separate. Contrary to our intuitions about drunk-drivers, sometimes scope must influence degree. Consider Willing Willy. Willing Willy will go along with whatever the group does. Willy does not have a good excuse, but Willy simply wants to belong to the group. In one case, Willy contributes to a wallet theft with the highest possible interconnectedness. The Full Scope Principle applies. In another case, he contributes to a genocide with extremely low interconnectedness. The Portion Principle applies. It is bizarre to say that Willy could possibly be as blameworthy in the former case as in the latter. Even if the agent is to the fullest possible extent blameworthy for stealing a wallet, he still is not as blameworthy as any agent who contributed to genocide. This might seem obvious, but it introduces an important question.

On the one hand, the wallet-versus-genocide comparison demonstrates that the overall group harm must somehow affect the degree of blameworthiness. On the other hand, if we want to remove luck from degree, the overall group harm might be included in an individual's scope (if interconnectedness is as high as possible), but it should never affect an individual's degree, because the overall group harm is partly a matter of (circumstantial and

resultant) luck. Both claims seem correct. The question is: how can scope both affect, and not affect, our degree of blameworthiness?

Our answer is that scope determines the range of degree. In group cases, the overall group harm determines the upper and lower limits of an individual's possible degree of blameworthiness. Yet only facts about his agency determine where he sits within that range (that is, in which percentile he sits). Of course, it's hard to exactly quantify varying degrees of blameworthiness. The following values are pure concoctions that we mention merely to get the point across. Suppose that when two agents produce a wallet theft, the possible degree of blameworthiness for each ranges from 1 to 20, where 1 is the lowest and 20 is the highest.¹⁰ Suppose that when someone contributes to genocide, the possible degree of blameworthiness ranges from 2,000 to 3,000. The degree within that particular range is determined by factors under the agent's control. The degree is not determined by interconnectedness: recall, interconnectedness determines scope, not degree; after all, interconnectedness is partly a matter of luck. The ranges-of-degrees also explain how two agents can have different scopes due to luck (via the luckiness of shared intentions, interpersonal influence, and common knowledge), and yet can be responsible to the same degree: there is an overlap in the ranges of degree that correlate to each agents' group harms.

Thus, luck enters in two places in determining an agent's blameworthiness for group harms. First, it partly determines the agent's extent of interconnectedness with other agents that contribute to the outcome. Interconnectedness determines the agent's scope of blameworthiness. In this way, luck affects scope—though not to the extent Kutz allows. Second, luck partly determines the lower and upper bounds of the range of degrees of

¹⁰ Assuming this range of degrees, if you are more blameworthy than "20" this would imply that you did something worse than steal a wallet.

blameworthiness the individual might have. Luck has this effect because these boundaries are determined by the overall group outcome, which depends partly on other people (and is thus partly a matter of luck). After all, depending on which principle applies, what I am responsible for is (my contribution to (and your contribution to)) *what we did together*. However, luck does not fully determine exactly where an agent falls within the range of degrees of blameworthiness. The agent controls whether they sit at the bottom, middle, or top of the range. Nor does luck entirely determine the agent's scope of blameworthiness, because some aspects of interconnectedness are within an agent's control.

6. Conclusion

To conclude, we return to the distinctions with which we started: between corporate, collective, distributed, and shared blameworthiness. As we mentioned, collective and corporate blameworthiness—non-reductive blameworthiness of structured and unstructured groups—are useful, though we haven't analysed them. They allow us to assign the full scope of blameworthiness to the group as such, in those cases where the Portion Principle or Overlap Principle apply at the individual level. To any readers who fear we have diminished individuals' blameworthiness too much (either in scope or degree), we remind them that group-level blameworthiness exists alongside individual-level blameworthiness. The group might have a large scope and degree.

Our discussion has concerned individuals' blameworthiness. It has covered both individuals in structured groups (distributed blameworthiness) and individuals in unstructured groups (shared blameworthiness). It has elucidated some factors that are common across these blame-types, despite these types often being treated separately in the literature. But the distributed/shared distinction has importance. All else equal, interconnectedness is more

likely to occur in structured groups. Of course, all else is not always equal: many large corporations aim to keep their low-level employees ignorant of the corporation's harmful ends and side-effects, preventing the possibility of shared intentions, interpersonal influence, and common knowledge. This explains why those low-level employees have small scopes of blameworthiness. Yet in structured groups with norms of transparency, accountability, and democracy, members are likely to have large and overlapping scopes, via high levels of shared intentions, interpersonal influence, and common knowledge. Our view gives a good explanation of why these individuals have overlapping scopes. In unstructured and diffuse groups, scopes are again likely to be small and non-overlapping. And, in all cases, an individual's degree of blameworthiness might be high indeed. In delineating individuals' scope and degree of blameworthiness, we have not gotten anyone off the hook.

Acknowledgements

For helpful comments on written drafts, we thank Franz Altner, Olle Blomberg, Anca Gheaus, Tim Fowler, Chris Mills, Carlos Núñez, Tom Parr, Grace Patterson, Herlinde Pauer-Studer, Matthew Rachar, Michael Schmitz, and Adam Slavny. For comments on a presented version, we thank Gothenburg University's Practical Philosophy Seminar.

References

- Blomberg, O. 2016. Common Knowledge and Reductionism about Shared Agency. *Australasian Journal of Philosophy* 94(2): 315-326.
- Bratman, M. 1999. *Faces of Intention*. Cambridge University Press.
- Chant, S.R. 2015. Collective Responsibility in a Hollywood Standoff. *Thought*, 4: 83-92.
- Collins, S. 2019. *Group Duties: Their Existence and Their Implications for Individuals*. Oxford University Press.

- Copp, D. 2007. The Collective Moral Autonomy Thesis. *Journal of Social Philosophy* 38(3): 369-388.
- Feinberg, J. 1968. Collective Responsibility. *The Journal of Philosophy*, 65(21): 674-88.
- French, P. 1984. *Collective and Corporate Responsibility*. Columbia University Press.
- Gilbert, M. 2014. *Joint Commitment: How We Make The Social World*. Oxford University Press.
- Gold, R. and N. Sugden. 2007. Collective Intentions and Team Agency. *Journal of Philosophy* 104(3): 109-137.
- Held, V. 1970. Can a Random Collection of Individuals Be Morally Responsible? *Journal of Philosophy*, 67(14): 471-81.
- Hess, K. 2018. "The Peculiar Unity of Corporate Agents." Pp. 35-60 in K. Hess, V. Ignieski, and T. Isaacs (eds), *Collectivity*. Rowman and Littlefield.
- Hindriks, F. 2018. Corporate Agency: Moral and Amoral. *Dialectica* 72(1): 3-23.
- Isaacs, T. 2011. *Moral Responsibility in Collective Contexts*. Oxford University Press.
- Knowles, C. 2019. Beauvoir on Women's Complicity in Their Own Unfreedom. *Hypathia*, 34(2): 242-265.
- Kutz, C. 2000. *Complicity: Ethics and Law for a Collective Age*. Cambridge University Press.
- Lacey, N. and H. Pickard. 2013. 'From the Consulting Room to the Court Room? Taking the Clinical Model of Responsibility Without Blame into the Legal Realm' *Oxford Journal of Legal Studies*, 33:1.
- Lawford-Smith, H. 2015. What 'We'? *Journal of Social Ontology* 1(2): 225-250.
- Lawford-Smith, H. 2018. Does Purchasing Make Consumers Complicit in Global Labour Injustice? *Res Publica*.
- Lewis, H.D. 1948. Collective Responsibility. *Philosophy*, 23(84): 3-18.
- Lewis, D. 1969. *Convention: A Philosophical Study*. Harvard University Press.

- Lillehammer, H. Forthcoming. "Moral Luck and Moral Performance." *European Journal of Philosophy*.
- List, C. 2014. "Three Kinds of Collective Attitudes." *Erkenntnis* 79(S9): 1601-1622.
- List, C. and P. Pettit. 2011. *Group Agency: The Possibility, Design and Status of Corporate Agents*. Oxford University Press.
- May, L. 1992. *Sharing Responsibility*. University of Chicago Press.
- Nelkin, D. 2019. Moral Luck. In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. URL=<<https://plato.stanford.edu/archives/sum2019/entries/moral-luck/>>.
- Pauer-Studer, H. 2018. Complicity and Conditions of Agency. *Journal of Applied Philosophy*, 35(4), 643-60.
- Pickard, H. and L. Ward. 2013. Responsibility Without Blame: Philosophical Reflections on Clinical Practice. In KWM Fulford, Martin Davis, Richard GT Hipps, George Graham, John Z. Sadler, Giovanni Stanghellini, and Tim Thornton (eds), *The Oxford Handbook of Philosophy and Psychiatry*. Oxford University Press.
- Pinkert, F. 2014. What We Together Can (Be Required To) Do. *Midwest Studies in Philosophy* 38(1), 187-202.
- Preda, A. 2012. Group Rights and Group Agency. *Journal of Moral Philosophy*, 9: 229-54.
- Schiffler, S. 1972. *Meaning*. Oxford: Oxford University Press.
- Tuomela, R. 2007. *The Philosophy of Sociality: The Shared Point of View*. Oxford: Oxford University Press.
- Tuomela, R. 2013. *Social Ontology: Collective Intentionality and Group Agents*. Oxford: Oxford University Press.
- Rovane, C. 1998. *Bounds of Agency*. Princeton University Press.
- Williams, B. 1985. *Moral Luck*. Cambridge University Press.
- Zimmerman, M. 2002. Taking Luck Seriously. *The Journal of Philosophy*, 99(11): 553-76.