

Blomberg, Olle, and Björn Petersson. Forthcoming. "Team Reasoning and Collective Moral Obligation." *Social Theory and Practice*. <https://doi.org/10.5840/soctheorpract2023120177>

Please quote from the official version of the paper which will be available at the Philosophy Documentation Center at <https://pdcnet.org/soctheory>

(A few minor changes were made to this document on January 25, 2023.)

Team reasoning and collective moral obligation

Olle Blomberg (olle.blomberg@gu.se) &

Björn Petersson (bjorn.petersson@fil.lu.se)

Abstract: We propose a new account of collective moral obligation. We argue that several agents have a moral obligation together only if they each have (i) a context-specific capacity to view their situation from the group's perspective, and (ii) at least a general capacity to deliberate about what they ought to do together. Such an obligation is irreducibly collective, in that it does not imply that the individuals have any obligations to contribute to what is required of the group. We highlight various distinctive features of our account. One such feature is that moral obligations are always relative to an agential perspective.

Keywords: collective moral obligation, group identification, we-framing, team reasoning, joint ability, agent-relativity

1. Introduction

It is often apt to ascribe a moral obligation to several moral agents collectively, even if they do not together constitute a collective moral agent. Consider the following case:

Burning Building: Three children are trapped in a burning building. One of them is in one room, and the other two are in a second room some distance away. The neighbours Agnetha and Benny see each other approaching the building from opposite sides. Agnetha breaks in and has enough time to do her part of rescuing either the child in the first room or the two children in the second room. The rescue can succeed only if Benny heads straight for the same room with his fire extinguisher. If both go to the first room, they will only rescue the first child. If both go to the second room, they will only rescue the two other children. If each goes to a different room, no child will be rescued. Suppose that Agnetha and Benny can make these choices without any significant risk to their own or each other's life or health. All this is common knowledge between them, but they do not have any opportunity to communicate with each other—each must choose which room to head for independently of the other. (Adapted from Colman et al. 2014: 36)

It seems to us that Agnetha and Benny are obliged to save the two children.¹ It would make sense to address them with a moral imperative: “You must save the two children!” Call this *the basic intuition* (following Schwenkenbecher 2020: 28). It is to *them*—collectively rather than distributively—that we would direct a moral demand to save the two children. Only together do they have the ability to rescue the two children, and rescuing the two children is the best they can do, morally speaking. Since they know that they can do this without significant risk to themselves, rescuing the two children is not only morally best, but arguably morally required. Such

¹ Throughout the paper, “obliged” and “obligation” are elliptical for “morally obliged” and “moral obligation,” unless otherwise noted.

considerations have convinced many philosophers that not only an individual, but also a group can have an obligation (e.g. McKinsey 1981; Copp 1991; Wringer 2010; Cripps 2013; Pinkert 2014; Aas 2015; Schwenkenbecher 2019, 2021; Björnsson 2014, 2020).

But one might think that only a single agent can be the subject of an obligation (see Lawford-Smith 2015; Collins 2019). The basic intuition arguably concerns a *deliberative* obligation; the obligation is connected to the question: “What ought I do?” (Lord 2015: 28). This may seem to imply that Agnetha and Benny could have the obligation only if they together constituted an entity capable of asking and settling on an answer to this deliberative question (cf. Wringer 2010: 224–225).² But the only relevant agents on the scene are Agnetha and Benny.

In this paper, we explain and vindicate the basic intuition by presenting our own novel account of collective obligation, an account that we believe has distinct advantages over others. We argue that demanding of Agnetha and Benny that they save the two children makes sense, because it is possible for them to each ask: “What ought we do?” Agnetha and Benny can therefore jointly be subjects of an obligation in *Burning Building*.³ They need not constitute a collective moral agent that can deliberate; it suffices that they each have the capacity to deliberate from a we-perspective.⁴ When we ascribe an obligation to ϕ to several agents considered collectively, we do so on the implicit assumption that each agent has the context-specific capacity to identify with the group and view the decision situation from the group’s point of view—to “we-frame” the situation—as well as at least a general capacity to deliberate about what the group ought to do. We develop this proposal by drawing on decision-theoretic work on “team reasoning,” sometimes also called “we-reasoning” (for a recent review, see Colman and Gold 2018).

² Holly Lawford-Smith (2015: 234–235) holds that only a single (individual or group) agent can have a moral obligation to Φ , because only a single agent can *try* to Φ , and being able to try to Φ is necessary for having the kind of ability relevant for “ought implies can.” However, several agents can *together* try to Φ : each of them can try to do his or her part of their joint Φ -ing (see Aas 2015: 15–17; cf. Collins 2019: section 3.2).

³ We are neutral on whether the plurality of Agnetha and Benny is nothing but Agnetha and Benny, or whether the plurality is some group-like object that isn’t entirely ontologically innocent (see Linnebo 2022).

⁴ Thomas Smith (2009: 48–49) makes some brief remarks roughly in line with this proposal.

Others have claimed that capacities for we-framing and team reasoning are important for obligations in the context of social dilemmas. Anne Schwenkenbecher (2019, 2021) would hold that Agnetha and Benny have a collective obligation to save the two children, and like us, she would require that they each have capacities for we-framing and “we-reasoning.”⁵ But furthermore, Schwenkenbecher assumes that each of them would then have an individual epistemic obligation to we-frame and we-reason their way to the conclusion that they must save the children. If either of them lacked such an obligation, then they would not together have a collective obligation on Schwenkenbecher’s account. In section 4, we argue that this is mistaken. Agnetha and Benny can have the collective obligation to save the two children even if neither of them has any obligation to we-frame and we-reason. The strong connection that Schwenkenbecher draws between we-reasoning and collective obligation needs to be relaxed to a mere capacity-requirement. As a result, it makes sense on our account to ascribe an obligation to a group of agents not only in social dilemmas such as *Burning Building*, where we-framing and team reasoning is required for agents to rationally act in accordance with their obligation, but also in a wider range of cases where we-framing and team reasoning are likely to occur, but are not necessary for such rational action.

Stephanie Collins (2019) would also argue that Agnetha and Benny each have something like individual obligations to we-frame and we-reason, although only if they share the objective of saving more rather than fewer children. They then have “duties to we-frame and coalition-reason” (2019: 140): each ought to consider the options available to them if they act together, and then do their part in the pattern of actions that is best in light of their shared objective. However, unlike Schwenkenbecher and us, Collins would deny that Agnetha and Benny could jointly be subjects of a collective obligation.⁶ Furthermore, in contrast to the team reasoning or we-

⁵ While we emphasize the ways in which we disagree with Schwenkenbecher, we are indebted to her work. We briefly sketched a team-reasoning-based account of collective moral obligation in Blomberg and Petersson 2018, but it differs in many ways from the one presented here. The details of our current view have to a large extent been worked out by engaging with Schwenkenbecher’s account.

⁶ Collins (2019) argues that only a group that is united under a group-level decision procedure—a single “group agent”—can have a “group duty.”

reasoning that we and Schwenkenbecher appeal to, Collins’s “coalition-reasoning” is fundamentally an individualistic form of practical reasoning. This is explained further in section 2, where we also criticize individualistic accounts such as Collins’s for not being able to make sense of the basic intuition.

In the next section, we explain why agents who can only practically reason about what “I” ought to do cannot be subjects of a collective obligation. We argue that this holds true not only in social dilemmas such as *Burning Building*—a collective obligation could never be aptly ascribed to such agents. In section 3, we introduce the notions of group identification and team reasoning before stating our positive proposal. In section 4, we argue that the concept of a collective obligation is irreducibly collective, in the sense that the subject of such an obligation is a group of individuals, while the concept cannot be reductively analysed as a constellation of ascriptions of individual obligations.⁷ Section 5 explores a consequence of our proposal, namely that collective obligations are always relative to an assumed agential perspective, where these different agential perspectives are incommensurable. This highlights a difference between our account and the accounts of Schwenkenbecher and Collins. On each of their accounts, there is a fundamental deliberative perspective from which the agent can determine whether it is objectively best to we-frame or I-frame a decision situation. According to Schwenkenbecher, a collective moral obligation is therefore not relative to any assumed agential perspective in the way that we argue it to be.

Our aim is to give an account that makes the idea of collective obligation intelligible. It is a further question exactly in what circumstances agents can have the capacities that we take to

⁷ As we note in footnote 16, we assume a version of a perspectival understanding of group identification (see Petersson 2017). Such an understanding is often taken to exemplify a non-reductive approach to shared or collective intentionality, but this is an issue that is distinct from the question of whether or not the concept of a collective moral obligation is irreducibly collective. Moreover, given a realist view of moral obligations, it might be fair to ask whether our conceptual irreducibility claim involves or presupposes an ontological irreducibility claim about the property of being collectively obliged—that is, of “having” a collective obligation. However, here we are not committed to any specific metaethical position concerning the ontology of moral obligations, but merely to a claim about the conditions for applying the concept properly in the collective context. While we make use of realist-sounding discourse about agents “having” a collective obligation, we assume that both metaethical realists and metaethical antirealists can make sense of such discourse.

be necessary for having a collective obligation. *Burning Building* is an idealized small-scale case that perspicuously illustrates the need for individuals to be able to view and reason about decision situations from the group's point of view. It is obvious to each of them that they can make a significant moral difference together, and that the contribution of each is necessary for success. Each of them can also identify the relevant others, and there are no other salient emergencies, aside from the burning building, which could give rise to competing obligations. The whole situation is also common knowledge between the two agents. Finally, once they have chosen which room to run to, it is obvious to both of them what needs to be done. All of these features are lacking in many large-scale problems that many agents could solve if they were to cooperate, like a state's citizens who together could overthrow their dictator, or the world's affluent people who together could slow down global warming. It is often assumed that if an account of collective obligation applies to idealized small-scale cases, then it should be applicable to messy large-scale cases as well (see e.g. Held 1970: 480–481; Wringer 2016: 480; Collins 2019: 129–130, 142–143; Schwenkenbecher 2021: ch. 7). We are somewhat sceptical about that move, but this is not an issue that we will go into here.

2. *I-reasoning and the basic intuition*

We will now explain why Agnetha and Benny would not be able to have a collective deliberative obligation if they could only ask and answer the question: "What ought I do?" We claim that if Agnetha starts by asking what she ought to do in *Burning Building*, then she cannot rationally settle on any determinate answer. The same is true of Benny. They could therefore not have any collective deliberative obligation to save the two children. But why is this so?

Burning Building can be seen as a Hi-Lo game where the sizes of the payoffs are determined by moral value.⁸ We assume that these also represent the preferences of the players. The payoff matrix will look like this to Agnetha and Benny.

		Benny	
		Room 1	Room 2
Agnetha	Room 1	1, 1	0, 0
	Room 2	0, 0	2, 2

Figure 1 *Burning Building*

Here, what Agnetha ought to do depends on what she reasonably expects Benny to do, and vice versa. Rationality in the game-theoretical sense merely dictates that each of them should pick the first room given that the other picks the first room, and that each of them should pick the second room if the other picks the second room. Given their common knowledge of the rationality of each of them, neither of them can rationally expect the other to go to the first room, and neither of them can rationally expect the other to go to the second room. This illustrates that in a Hi-Lo game there are two sets of strategies in which no player could raise her payoff by unilaterally changing strategy—the game has two Nash equilibria—and both players prefer the outcome of one of these sets to the outcome of the other. In this game, the individualistic account of practical rationality embedded in game theory cannot recommend a choice for either player, even though it is intuitively obvious that each of them ought to run to the second room (for detailed discussion of the challenge Hi-Lo poses to orthodox game theory, see Bacharach 2006: ch. 1).

Most discussions of Hi-Lo are focused on the question of what rationality requires in light of the players' preferences, not on what morality requires in light of what is morally valuable.

⁸ For other cases that can be seen as moral Hi-Lo games, see Collins 2019: 127–129; Schwenkenbecher 2019: 152, 155–156; 2021: 82. For a mundane non-moral example of a Hi-Lo game, imagine two friends who are trying to find each other at a festival and have to decide between their two standard meeting places, of which one is slightly better than the other.

When players have self-interested preferences and confront certain other types of potential social dilemmas, like a potential Prisoner's Dilemma, the situation may become unproblematic if the individual players each change their preferences and start caring about the best outcome for the group—if they become “team benefactors,” in Bacharach's terminology. However, as Bacharach stresses, such preference transformations will not suffice to resolve a potential Hi-Lo game. Individual rationality provides no determinate answer to the question “What should *I* do to promote the satisfaction of our preferences?” in a choice situation with a Hi-Lo structure. But Hi-Lo is not only a challenge to an individualistic account of prudential rationality, but also to an individualistic account of what morality requires of moral agents. Individual rationality provides no determinate answer to the question: “What should *I* do to promote the morally best outcomes?”

A defender of orthodox game theory might respond to the Hi-Lo challenge as follows. In the absence of evidence regarding which probabilities the other agent will choose among her alternative strategies, it is rational to distribute one's credences equally among these strategies. So, given that it is rational for Agnetha and Benny to each distribute their credences according to such a *principle of insufficient reason*, as well as to maximize expected utility, it would be rational for each of them to run to the second room ($EU = 0.5 \times 2 + 0.5 \times 0 = 1$) rather than to the first room ($EU = 0.5 \times 1 + 0.5 \times 0 = 0.5$).⁹ For each of them, running to the second room would be the choice that maximizes expected utility. But why ought an agent distribute her credences in this way if the probabilities are unknown?¹⁰ In some Stag Hunt games, the principle of insufficient reason will, problematically, require a rational agent to hunt hare rather than hunt stag.¹¹ Furthermore, given the standard game-theoretic assumption of common knowledge of

⁹ Bacharach refers to this principle as a “principle of equiprobability” (2006: 52). It is also known as the *principle of indifference*.

¹⁰ For a brief and accessible critical discussion of the principle of insufficient reason, see Peterson 2017: 54–56.

¹¹ Consider the following non-moral Stag Hunt: If each player gets a payoff of 2 for hunting hare whether or not the other does, each gets 0 for hunting stag on their own, and they get 3 each for hunting stag together, then the principle of insufficient reason tells each player that their only rational choice is to hunt hare: EU for hunting hare ($0.5 \times 2 + 0.5 \times 2 = 2$) > EU for hunting stag ($0.5 \times 3 + 0.5 \times 0 = 1.5$). For a moral Stag Hunt game, which depends on morality being agent-relative in a peculiar way, see section 5.

rationality, this principle would lead agents to have inconsistent beliefs about probabilities. Since Agnetha knows that Benny adheres to the principle of insufficient reason and is maximizing expected utility, she would have to assign a higher-than-50% probability to Benny running to the second room. After all, she knows that Benny's expected utility of running to the second room is twice as high as his expected utility of running to the first room. Since Benny knows that Agnetha knows this, and Agnetha knows that Benny knows that, and so on and so forth, the probability that Agnetha would have to assign to Benny running to the second room would have to be iteratively ratcheted up higher and higher. Given their symmetrical positions, the same would be true regarding Benny's estimate of the probability that Agnetha would run to the second room. At first glance, this may seem like a virtue rather than a bug, since it would deliver the desired result—it would become obvious for each of them that they ought to run to the second room. However, they cannot each rationally infer that the probability of the other running to the second room is higher than 50% on the basis of an estimate that this probability is 50%! They would then hold contradictory beliefs. Hence, either the principle of insufficient reason or the assumption of common knowledge of rationality must be abandoned (see Bacharach 2006: 52–53, 66 n. 8). But without the principle, the proposed solution would not get off the ground, and without the standard assumption, it would be incompatible with the “I-reasoning” theory of rational interaction that is part of game theory.

Collins (2019) attempts to explain why agents ought to play Hi in Hi-Lo by appealing to what she calls “coalition-reasoning,” which is similar to the reasoning involved in this mistaken solution. She argues that in Hi-Lo, agents can rationally transition from an initial 50%/50% credence distribution regarding what the other will do (2019: 131) to “a strong and justified belief that all the others will do their part in the optimal pattern” (2019: 145) via “a ‘ratcheting up’ effect” (2019: 144). Due to this iterative effect, the agents are supposed to become “entitled to make presumptions about what the other agents will do” (2019: 149). If Collins is here describing how a practical conclusion can be rationally inferred from certain premises, then her account of coalition-reasoning arguably requires agents to have contradictory beliefs. However,

it may be that Collins intends the ratcheting-up effect to be understood as the result of a temporally extended psychological process that agents facing Hi-Lo problems tend to go through in parallel. So conceived, coalition-reasoning may not require reasoners to have contradictory beliefs. But it is at best unclear whether such parallel psychological processes would be rational, even if we were to accept the principle of insufficient reason.¹²

A different defence of individualistic game theory appeals to the idea that the Hi equilibrium is *salient* (cf. Pinkert 2014: 194 n. 22; Collins 2019: 131). Given that this salience is common knowledge between Agnetha and Benny, it seems to allow each of them to rationally predict that the other will run to the second room, or allow each of them to believe that the other will predict that they themselves will run to the second room, or allow each of them to believe that the other will believe that they themselves will predict that the other will run to the second room, and so on and so forth. However, it is not clear why the Lo equilibrium isn't just as salient as the Hi equilibrium. After all, it is *the worst* Nash equilibrium. While an appeal to a psychological tendency to choose a particular kind of salient option may explain how we behave in situations such as the one represented by Hi-Lo, it cannot explain the rationality of each agent to do their part in the Hi equilibrium (see Gilbert 1989; Sugden 1993: 77–84). Without any explanation of why it would be *rational* for Agnetha and Benny to save the two children, it is not clear how they could have an obligation to save them.

Another explanation of why a collective obligation cannot be reduced merely to several parallel answers to the question “What ought I do?” concerns the agents’ normative reasons.¹³ Plausibly, the subject of an obligation to Φ must not only have the ability to Φ , but also the ability

¹² In response to a similar objection to coalition-reasoning (Blomberg 2020: 112–114), Collins (2020: 149) argues that the objection could just as well be levelled against an account of team reasoning that depends on a non-rational psychological mechanism of group identification (see section 3). But team reasoning prompted by group-identification can make sense of playing Hi as a rational choice from the adopted group perspective. It is not clear to us that coalition-reasoning could make sense of playing Hi as a rational choice from any perspective that the agent could occupy.

¹³ Of course, there will often be several parallel answers to the question “What ought I do?” when a group has a collective moral obligation. Given that the other does their part, Agnetha and Benny will each answer the question “What ought I do?” with the answer: “Run to the second room.” But for each of them to arrive at this answer in a determinate and non-accidental way, they each first need to answer the question: “What ought we do?”

to Φ for the normative reasons that make Φ -ing morally obligatory (see Lord 2015; cf. Collins 2019: ch. 3).¹⁴ Agnetha and Benny would lack this ability if each of them were restricted to asking and answering the question: “What ought *I* do?” In its first-person singular form, the deliberative question concerns what to do solely on the basis of the person’s own agentic abilities. Hence, if Agnetha were limited to “I-reasoning,” then she would not have a normative reason to do her part in saving the two children. Nor would Benny if he were limited in the same way. Only together can they have a normative reason to save the two children (cf. Dietz 2016: 960–963); and only if they together have this normative reason will each of them have a normative reason to do their part.

Call the versions of Agnetha and Benny who can only I-reason *Agnetha_I* and *Benny_I*. The obligation of each of *Agnetha_I* and *Benny_I* will be constrained by their own personal abilities, and they can therefore at most each have an obligation to do their part in a collective endeavour such as a joint rescue. Hence, if individuals were restricted to I-reasoning, then all the deontic facts would be captured by ascriptions of individual obligations (including quite complex conditional obligations). We take this to not only be true in Hi-Lo-like situations such as *Burning Building*, but also in other types of social situations. Here is an analogy that should help to bring out this point. Consider an individual agent who is limited to asking: “What ought I do *now*?” What could she be morally obligated to do? She could not be morally obligated to perform a temporally extended composite action, since she would not be able to grasp the normative reasons that would make

¹⁴ As an anonymous reviewer pointed out, this premise is controversial. According to Collins, “bearing a duty entails having the ability to fulfil *at least one duty*” (2019: 92), but not “having the ability to fulfil *that very duty*” (ibid.). Here, “fulfilling” a duty (or obligation) means acting in accordance with it for the normative reasons that make the duty-according behaviour obligatory. The reason why Collins does not embrace the stronger thesis is that there are cases where one has (or “bears”) an obligation even though there is no time for the deliberation or decision-making required for *fulfilling* it. Wringer goes further and suggests that it may be sufficient if an obligation-bearer can act in accordance with obligations in “a non-fluky and relatively reliable manner” (2020: 123): the obligation-bearer need not also be able to act for the right reasons. We believe, though, that even in the cases considered by Collins where there is no time for deliberation, the agent needs at least a *general* capacity to do what is obligatory for the normative reasons that make it obligatory. Without such a capacity, the agent wouldn’t be able to learn and rationally modify their behaviour in similar future situations in response to being confronted with blame for having violated an obligation in a situation (see section 3). Since we think it is plausible that there is a tight connection between obligations and the possibility of appropriate blame in response to unexcused violations of obligations, we find both Collins’s and Wringer’s views too relaxed here.

that composite action obligatory. At most, she could have, at different stages, obligations to perform individual components of such an action. (Consider the obligation of a lone parent to raise and take care of a child.) Similarly, agents who lack reasoning capacities that go beyond I-reasoning cannot have a collective obligation. Having a collective obligation together requires each agent to have the capacity to deliberate about and make plans for what *we* ought to do. It would be pointless to posit or ascribe a collective obligation to a group of I-reasoners; just as there is no need to posit or ascribe an obligation to perform temporally extended composite actions to an individual agent who only has the capacity to deliberate in response to the question: “What ought I do *now*?” We take this to be a challenge for accounts of collective obligation that do not require agents to have capacities for we-framing and team reasoning (e.g. Björnsson 2014, 2020; Pinkert 2014).

Agnetha₁ and Benny₁ could clearly *end up* successfully saving the two children, and in doing so, each of them could also fulfil their obligations. Suppose Benny₁ plumps for running to the second room, hoping that Agnetha₁ will do the same. Suppose that Agnetha₁ catches sight of him and sees where he is running to. She then acquires an individual obligation to do her part in saving the children. When she acts to fulfil this obligation, suppose that Benny₁ in turn sees her running to the second room. He then in turn acquires an individual obligation to follow through on running to the second room and do his part in saving the children. However, the fact that at this point they would each have an individual obligation to do their part in saving the two children does not give us any reason to posit a collective obligation. The deontic features of their situation would be exhaustively captured by their individual obligations. Furthermore, Benny₁ never had an individual obligation to plump for running to the second room in the first place. In sum, given that Agnetha and Benny were each restricted to asking the deliberative question “What ought *I* do?”, they could not have any collective obligation to save the two children in *Burning Building*.¹⁵

¹⁵ Would the situation be different if Agnetha₁ and Benny₁ were able to openly communicate and agree to save the two children together? It is not clear that it would. Agnetha₁ only ought to act according to the agreement if she is justified in assuming that Benny₁ will. And Benny₁ only ought to act according to the agreement if he is justified in assuming that Agnetha₁ will. After all, if Agnetha₁ acts contrary to the agreement and runs to the first room, then Benny₁ ought to run to the first room too, and vice versa.

Of course, each of them could have an obligation to, say, plump for running to either room unless they see the other running for a room, in which case they ought to run to the same room. But these obligations wouldn't account for the basic intuition.

At this point, one might be tempted to give up the idea that Agnetha and Benny have a collective obligation. When we demand of them that they must save the two children, one might think that we are actually ascribing a complex conditional individual obligation to each of them. In this vein, Collins (2019: 117) would claim that Agnetha has an obligation to do what she can (within reasonable cost) to make it reasonable for Benny to believe the following: that she will take responsive steps with a view to the two children being saved if she believes that he will do likewise (and she is obliged to also take these steps if she reasonably believes that he will do likewise). Benny would have a symmetrical obligation.

We find this type of complex individualistic account unsatisfying for two reasons. First, what would ground Agnetha's complex conditional obligation? Intuitively, it would be the fact that she and Benny *ought* to save the two children, but this would presuppose the collective obligation that Collins is trying to explain away (see Schwenkenbecher 2021: 31). Secondly, what if it were common knowledge between Agnetha and Benny that both of them would *not* do likewise? It would then be pointless for each of them to e.g. signal their own willingness, and each of them would thus be excused. This is surely an unhappy result. Not only do we want to say that Agnetha and Benny would be violating an obligation to save the two children here, but blaming them for this would also seem to be justified. Collins (2019: 118–120) suggests that in real-world cases, no individual agent can rule out the possibility that she could convince the unwilling other(s) to contribute. It may be right that there are few real-life cases where the subjective probability of succeeding when trying to make others cooperate is zero. Less rare though, are cases where there is a clear discrepancy between what I should do, as part of what we should do unconditionally, and what I should do, given what, with some degree of confidence, I believe that you will do after attempted persuasion.

It should now be rather obvious what is required to save the basic intuition. If it is possible for each of Agnetha and Benny to ask and answer a collective version of the deliberative question—“What ought *we* do?”—then their capacities to do this can explain and vindicate the basic intuition. We turn next to this idea, and to our own positive account.

3. *Moral team reasoning and collective obligation*

The notion of “group identification” has been central to social psychology since the 1970s. In the thinnest sense, “*A* identifies with group *G*” may simply refer to an individual’s categorization of herself as belonging to a particular group. In a slightly stronger sense, it might mean that *A* sympathizes with *G*, shares the group’s values, etc. Here we are talking about an individual’s identification with her group in a different sense, implying not only that the individual can see herself as part of the group and care about this, but also that she is able to view the choice situation from the group’s perspective. Put differently, an individual agent can have attitudes that are held from her group’s viewpoint. This would mean that there is a sense in which a kind of identification between the plural subject and the individual subject takes place, albeit “in the head of an individual,” to use John Searle’s phrase (1997: 27–28). This view of group identification is related to a particular approach in the current collective intentionality debate.¹⁶

Experimental evidence suggests that group identification leads individuals in social dilemmas to be concerned about and act on the group’s interests rather than their own private interests. In these experiments, group identity has been induced through pre-existing shared community membership (Kramer and Brewer 1984), common fate (Brewer and Kramer 1986)

¹⁶ The philosophical literature that explicitly discusses “group identification” in those very terms is relatively recent and limited (among the few instances of this are Pacherie 2011; Salice and Montes Sánchez 2016, 2019; and Salice and Miyazono 2020). However, much work related to the understanding of collectivity discusses closely related concepts. Approaches to collective intentionality are usually grouped into three categories (see Schmid and Schweikard 2021). Some authors claim that collective intentionality is intentionality with a collective *content* (Ludwig 2016), others seem to invoke a special *mode* or *perspective* (Searle 1997), while still others claim that what’s collective about collective intentionality has to be the *subject* (Helm 2008). We assume that a version of the perspectival understanding of group identification is the most promising candidate for capturing the notion that would fulfil the role assigned to it in the “team reasoning” framework (Petersson 2017).

or through face-to-face discussion (Dawes, Van De Kragt, and Orbell 1990). The hypothesis that group identification tends to lead to cooperation with in-group members is a plausible live hypothesis.

According to Michael Bacharach (2006), group identification can also be triggered by the strong interdependence that exists between individuals' interests in a social dilemma such as the Prisoner's Dilemma or Hi-Lo.¹⁷ Group identification—viewing the decision situation from the collective's perspective—is an essential element of team reasoning in Bacharach's sense. Group identification does not guarantee team reasoning, though. Team reasoning is a particular model of decision-making, and one could imagine situations where individuals evaluate options from their group's perspective but simply refrain from making a decision about what to do, or where “various subpersonal perception-action routines” guide each individual's behaviour (see Björnsson 2014: 115). The point is that if individuals group-identify in the process of deciding what to do, then they will (in normal non-akratic cases) choose their part in the optimal set of strategies.

A team reasoner views the decision situation from the point of view of the group—she asks herself what *we* should do, rather than what *I* should do for us. She evaluates the courses of action available to the group, and infers which component of the collective action she should perform. This last step means that each individual asks herself: “What should I do as part of what we should do?” In many cases, it will not be necessary for her to know in detail what the various components are, or how they fit together. As long as she knows that her own contribution will fit together with the contributions of the others, she may in the limiting case represent those contributions only in a coarse-grained way, as merely “the others' contributions.”

Even if group identification typically goes hand in hand with concern for the group's interests, the essential point for us is that mere concern for the group, without group identification, is sometimes insufficient to guarantee the collectively optimal set of individual strategies. Furthermore, agents who group-identify need not act together in the group's interests.

¹⁷ As far as we know, this hypothesis has not been experimentally tested (see Colman and Gold 2018).

They can arguably also act together to do what they morally ought to do, whether or not this aligns with the group's interests.

The team reasoner identifies with the group in a practical sense—the shift from I-reasoning to we-reasoning is what Bacharach calls an “agency transformation.” Given that the situation in *Burning Building* triggers Agnetha and Benny to group-identify, it becomes rational for each of them to do their part in saving the two children: each identifies the outcome that they save the two children as the best option available to the group (given the moral preferences assumed in the payoff matrix in section 2), and each of them knows what their own part in bringing this outcome about consists in. If they act rationally, each of them will then intend to run to the second room and do their part in saving the two children.¹⁸

While individuals can rationally solve social dilemmas thanks to team reasoning, they cannot, on Bacharach's model, rationally choose to adopt team reasoning rather than I-reasoning. Group identification is *triggered*; it is a-rationally prompted by the character of the choice situation. This doesn't mean that an agent who has group-identified is not reasoning rationally when they engage in team reasoning. Tendencies to group-identify in certain types of situations are likely to be universal or shared among community members, and a team reasoner may therefore be warranted in acting on the presupposition that the others will do their part (for complications, see Roth 2014).

Is Bacharach right to think that group identification isn't a matter of choice? If identifying with a group involves believing that one belongs to a group, then one arguably cannot *directly* choose to identify with it: beliefs are acquired rather than chosen. However, within the current framework, conceiving of a decision situation from a group perspective is not supposed to be

¹⁸ While we rely on Bacharach's (2006) model of team reasoning, there are various variations of the core idea (see Sugden 1992; Postema 1995; Gold and Sugden 2007; Hakli, Miller, and Tuomela 2010; Colman and Gold 2018; Gold and Colman 2018; Bermúdez 2021). We take the result of an individual's team reasoning to be a choice for the group, a “we-intention,” the content of which is a profile, i.e. a set of strategies for the group. The team reasoner then infers what her part in that collective endeavour is, and she will typically form an individual intention to perform that part. This does not exclude the possibility that the team reasoner for some reason fails to form the latter kind of intention, due to weakness of will, for example.

equivalent to believing that one is a member of the relevant group. The perspectival understanding of conditions for collectivity is arguably a distinct alternative to accounts requiring a conception of the group in the content of participants' attitudes (see Schweikard and Schmid 2021; Petersson 2017). But involuntarism seems to apply to perspectival features of cognition too: when I recall the shameful thing I did at that party twenty years ago, I cannot choose whether or not the remembered event seems temporally close or distant. Such involuntarism might extend to the agential perspective from which I conceive a decision situation.

We claim that the ascription of a collective obligation to a group presupposes that its members have capacities to group-identify and to team reason. When we address the group collectively with a moral demand, rather than distributively with a set of demands, we implicitly appeal to these capacities. It may seem unreasonable to make a demand while presupposing that its primary addressees have the capacity to take one agential perspective rather than another if they cannot choose which agential perspective to take. But consider, for comparison, the obligation to fulfil a promise. This obligation presupposes that the subject of the obligation has the capacity to remember that they have made the promise, and to recall what to do when an opportunity to fulfil the promise arises. This kind of prospective remembering isn't something that the subject of the obligation has direct control over (see Grünbaum and Kyllingsbæk 2020). Nevertheless, we can have obligations to do things in the future, and we can be blameworthy for violating them as a result of forgetfulness. One can arguably have an obligation to do something later without also having an obligation to recall that one must do it. The normal functioning of the capacity to recall is presupposed by the obligation (see Clarke 2017). Similarly, the capacity to group-identify could be presupposed by the collective obligation. Alternatively, having an obligation to do something in the future requires indirect control of recall—a control one exercises by e.g. putting a reminder in one's calendar to do what is morally required at the right time. Such indirect control of one's capacity to recall may suffice for enabling obligations that presuppose the possession of that capacity. We can arguably have such indirect control over group identification. Consider an analogy with visual perspectives. Once you are aware of the two

ways in which you can perceive the duck-rabbit, it becomes easy to switch between the two perceptions. When this procedure becomes sufficiently automatic, the direct/indirect distinction does not matter much. If your life depended on me perceiving the duck, it would not seem implausible to say that I had an obligation to do so, regardless of whether I could do so directly by choice or if I had to proceed via some kind of (possibly highly automatized) self-nudging.

The question of whether an individual has the capacity to group-identify might concern her general capacity to view things from the group perspective in any situation where external conditions for group identification are favourable. It might also be a question about her capacity to do this in a specific choice context. It is the latter context-specific capacity that is presumed by the ascription of a collective obligation. Now, capacities and abilities can be general and specific to different degrees and in various ways.¹⁹ The context-specific capacity that the individual is presumed to possess is not a capacity “in view of the totality of facts obtaining at some specific time” (Jaster 2020: 117). If it were that specific, then the individual would only have the capacity for group identification if they actually group-identified.

In standard examples of non-agential groups that intuitively have obligations, like Virginia Held’s subway passengers who fail to intervene and stop an assault (Held 1970), external conditions for group identification are typically favourable: in Held’s case the group is confined to a limited common space, there are no obstacles to communication or signalling preparedness to coordinate, etc. Little should be needed to prompt group identification, and consequently our moral intuition about the collective obligation seems very stable. As we have described *Burning Building*, the external conditions are perhaps not quite as favourable, since opportunities for signalling and talking are absent, but group identification is nevertheless to be expected. There is the strong interdependence between the agents’ (moral) interests that tends to trigger group identification according to Bacharach. It is out in the open in *Burning Building* that this is what

¹⁹ We generally use “capacity” to refer to the ability to group-identify or the ability to team reason, and we use “ability” to refer to the ability of agents to do what is required by a collective or individual obligation. While all of these are abilities, the exercise of the capacities to group-identify and team reason is at least typically not under the agent’s direct voluntary control, in contrast to the exercise of the ability to e.g. do their part in saving the two children.

Schwenkenbecher (2019) calls a “strict joint necessity” case. Each agent must contribute—there is no room for some to take up the slack of others. Given this, it is plausible to assume that in *Burning Building*, Agnetha and Benny each have a context-specific capacity to group-identify.

It seems that the strength of our intuition about their collective obligation will vary according to the extent to which these conditions are likely to prompt group identification. We assume that we each have an implicit grasp of whether a given situation is likely to lead to group identification, that is, of whether the agents have the context-specific capacity to group-identify in that situation. In some cases, the conditions for prompting group identification make it inappropriate to ascribe a collective obligation to a group: a set of individuals randomly selected from the world’s population will not be the subject of a collective obligation just because they each have the general capacity to group-identify. We cannot here specify when the external conditions are favourable enough for the individuals to have a context-specific ability to group-identify. All we can do is point toward some factors that are likely to be relevant.²⁰ It has been a recurring thought that some sort of “structure” that falls short of unified (group) agency may suffice for a group to have a collective moral obligation or be collectively blameworthy (see e.g. Manning 1985: 100, 102; Copp 1991: 75–76; Lawford-Smith 2015: 237). We suspect that this thought reflects the fact that such structures coincide with conditions under which agents have the context-specific capacity to group-identify.

Our proposal, then, is that several moral agents have a collective (deliberative) obligation to ϕ if and only if:

- (1) ϕ -ing realizes the morally best available outcome;
- (2) they have the context-specific joint ability to ϕ ;

²⁰ When we say that a group has an obligation to ϕ , we imply that the group ought to ϕ and that it would be meaningful to address the group with a moral imperative, ϕ ! Such a moral imperative can contribute to group identification as an external condition in the choice context. By demanding collective action, we stress the collective features of the case, and we may thereby tip them over into taking the group perspective.

(3) they each have a context-specific capacity to group-identify and team reason about what they ought to do, and to identify ϕ -ing as an option for collective action in their deliberation.

In *Burning Building*, saving the two children is a salient option for collective action that is open to the group. It is the action that realizes the morally best result, and intuitively ought to be performed. What makes it true that condition (1) holds is certain moral considerations that a correct normative moral theory should help us to pick out. However, note that our account does not involve any substantive moral commitments about what makes an outcome morally best. What if one or more of the agents has the wrong view of the relevant moral considerations, thereby preventing the group from fulfilling its obligation? This does not cancel the collective obligation; just as having the wrong moral view does not cancel a single agent's obligations. Nevertheless, since they are moral agents, we take it that they have a general capacity to understand the moral considerations that make (1) true.

The agents also need to have the joint ability to ϕ in the context at hand. Agnetha and Benny only have the collective obligation to save the two children if they have the ability to do so, that is, if they together have the joint ability to bring about the outcome that the two children survive. Hence condition (2). Furthermore, they need to know or at least be able to figure out that they have this joint ability. This is implicit in the next condition (3): they must be able to identify ϕ -ing as an available option for collective action in their deliberation.

Condition (3) is the key condition that we have focused on in this section. Note that it is typically not part of the content of a collective obligation ("to ϕ ") that the group members *exercise* the capacities required by condition (3). If Agnetha and Benny somehow accidentally managed to coordinate on running to the second room and ended up saving the two children, then they would have acted in accordance with their collective obligation. However, it would be inappropriate to ascribe that collective obligation to them if they lacked capacities to group-identify and team reason.

Must it be *possible* for the agents to fulfil their collective obligation by exercising their capacities for team reasoning? Suppose that a large person suddenly stumbles and is about to fall between us as we are walking in parallel in the street. For either of us, stopping the fall and holding the person upright on our own would be difficult and cumbersome, but doing so together would be easy. Given the tight time constraints of the situation, might we here have a collective obligation to do this? There is no reason to think that group identification and we-framing are more cognitively demanding than I-framing the situation. So we should not rule out the possibility that we could team reason in this situation. Team reasoning can be a process that, from the subject's point of view, seems to involve no more than perceiving and acting on an affordance for joint action (cf. Gold 2018). But could we be morally required to do something together, such as catching the stumbling person, even if the joint solution to the task facing us couldn't even be computed on-the-fly by processes of team reasoning?²¹ If so, then agents can have a collective obligation without having context-specific capacities for team reasoning. However, individuals would need to at least have a general capacity to team reason their way to the belief that they ought not to have done what they did, and to do so on the basis of the normative reasons grounding the collective obligation. If we had a collective obligation to catch the stumbling large person but violated this obligation with no excuse, then it would be appropriate to blame us for this violation. But it would arguably be pointless and hence inappropriate to blame us for this if we lacked general capacities for team reasoning. Without such general reasoning capacities, and without the context-specific capacity to group-identify, we could not be expected to change our ways to act differently in a similar future situation as a result of being confronted with blame. Hence, without at least the general team reasoning capacity, we would not have the collective obligation in the first place.²²

There is no reason to think that conditions (1)-(3) could only be satisfied in Hi-Lo cases like *Burning Building* or in other social dilemmas. Hi-Lo cases are illuminating because they

²¹ In such a case, a joint solution could perhaps only be computed by virtue of "an interagential structure of motor representations." (Sinigaglia and Butterfill 2022: 8)

²² This paragraph benefited from discussion with Gunnar Björnsson.

highlight the functional difference between I-reasoning and team reasoning, and the role of group identification for transforming potential social dilemmas. But there is no reason to restrict the possibility of group identification and team reasoning to collective choice situations with that structure. Group identification and team reasoning may be prompted by factors like those mentioned above—common fate, face-to-face interaction, etc.—in many kinds of situations, which needn't be social dilemmas. We can meaningfully ascribe collective obligations in cases where the morally required set of actions could also be reached by means other than group identification and team reasoning. In such cases, it could also be meaningful to ascribe a set of individual obligations to perform those same actions. For example, a gang of football hooligans may be ascribed a collective obligation to avoid harassing other citizens, and each hooligan may simultaneously be ascribed an individual obligation to avoid harassing other citizens.²³

²³ It may be useful to compare our necessary and sufficient conditions with the sufficient conditions presented by Schwenkenbecher:

"Agents *a*, *b* and *c* have a collective moral obligation if

- (i) There exists a specific morally significant joint-necessity problem *P*, such that agents *a*, *b* and *c* can collectively, but not individually, address *P* [joint necessity + joint ability].
- (ii) Conscientious moral deliberation leads all of them (or a sufficiently large subset of them) to believe that some collectively available option *O* is morally optimal with regard to *P* [they have reason to we-frame *P* and to consider *O*].
- (iii) *A*, *b* and *c* (or a sufficiently large subset of them) are in a position to determine individual (or joint) strategies to realize *O* and to achieve *P*." (2021: 93, brackets in the original)

These are not necessary conditions, because Schwenkenbecher thinks that *a*, *b* and *c* could have a collective obligation even if one of them could address *P* on their own, but it would be unfairly burdensome for one to do it without the others' help (2021: 24 n. 24, 101). Examples include collective obligations of "co-parenting or sharing housework" (2021: 101). In such cases, condition (i) is not satisfied—they are not joint necessity cases. It is unclear whether Schwenkenbecher thinks that the agents must each have a reason to we-frame and we-reason in these cases too. But if so, then her view would be that something like conditions (ii) and (iii) would be necessary conditions. If condition (i) were reformulated into a disjunctive condition that required either joint-necessity or a fairness-related problem that could be addressed collectively, then this would yield a generalized version of Schwenkenbecher's account, set out in terms of necessary and sufficient conditions.

How does our account differ from such a generalized version of Schwenkenbecher's account? We suspect that neither joint necessity nor fairness-related circumstances are necessary for agents to have a collective obligation, so the disjunctive version of condition (i) is not necessary. The football hooligans example is a case in point. First, it is not clear that this is a joint necessity case, since the collective outcome is likely to be achieved without any coordination or cooperation, by each of the hooligans simply avoiding any harassment of citizens (although each of them can unilaterally ensure that the collective obligation is violated). Second, any reasons for the hooligans to fairly share the "burdens" of restraint seems to be irrelevant for the ascription of the collective obligation here. But the hooligans' avoidance of citizen-harassment is nevertheless (part of) what realizes the morally best outcome. More importantly, something like condition (ii) is not necessary on our account. Agents who have the capacity to group-identify and team reason may nevertheless not group-identify and team-reason, even though they deliberate conscientiously. They may have no reason to group-identify and team reason if they know that the others *will* do their part,

4. *The irreducibly collective nature of collective obligations*

The capacities to group-identify and team reason required in condition (3) are individual capacities, but a collective obligation is not thereby reducible to individual obligations. Agnetha and Benny's obligation to save the two children is not reducible to—nor does it entail, require or somehow generate—an individual obligation for each of them to we-frame, team reason, and do their part of the required joint action.²⁴ As we have mentioned, this sets our account apart from Schwenkenbecher's account. According to Schwenkenbecher (2019, 2021), several individuals together have an obligation to ϕ in a strict joint necessity case such as *Burning Building*, if and only if, roughly: they have a joint ability to ϕ , and they each have an epistemic obligation to “we-frame” and “we-reason” their way to the conclusion that the best they can do (or what they must do) is to ϕ .²⁵

However, even if it would be possible for each of them to have the individual obligation required by Schwenkenbecher in a case such as *Burning Building*, no individual has to have it. Suppose that each of Agnetha and Benny has the context-specific capacities to group-identify, we-frame, team reason and do their part. However, it is common knowledge between them that neither will do their part in saving the two children: each of them knows that the other is a weak-willed team reasoner who, even if they were to group-identify and judge that they ought to save the two children, would definitely not bring themselves to intend to do their part. In this situation,

or if they know that the others *will not* do their part (see section 4). Nevertheless, they may still have a collective moral obligation on our account.

²⁴ We find plausible Wringer's principle that “all-out collective obligations entail, but are not equivalent to, certain pro tanto conditional obligations on the parts of the individuals who make up the collective” (2016: 488). However, these conditional obligations require each individual to perform their part of the required joint action if the others do theirs; they do not also require each individual to we-frame and team reason if the others do so. When others will not do their part, an agent's pro tanto conditional obligation to do their part does not require them to do anything. But the agent may also have ordinary individual obligations, such as an obligation to try to convince the others to do their part (at least if there is a reasonable chance that they may be convinced).

²⁵ Schwenkenbecher (2021) prefers to talk of “reasons” to we-frame and we-reason (see footnote 23), but she assumes that these reasons generate epistemic obligations (on the relationship between reasons and obligations, see pp. 27, 88; on epistemic obligations to we-frame and we-reason in particular, see e.g. pp. 66–67, 83, 95 n. 3; see also her 2019: 152).

given the other's incontinence, neither of them would have an individual moral or epistemic obligation to team reason and do their part of the required collective action. Neither of them would have such an obligation, because complying with such a putative obligation would be pointless in the situation at hand. Alternatively, suppose that it is common knowledge between them that they are racists and that the two children in the second room in *Burning Building* have what they consider to be the wrong skin colour (cf. Björnsson 2020: 128), and each of them knows that the other would refuse to help if they themselves were to run to the second room, or if they tried to convince the other to run there. Schwenkenbecher (2021: 87, 106) herself suggests that we-framing would be pointless in such situations. It would also be pointless for Agnetha to we-frame and we-reason if she knew both that Benny we-framed and we-reasoned in the situation, and also that he presupposed that she did as well. She could then predict that Benny would run to the second room, and through ordinary I-reasoning she could conclude that she ought to run to the second room. Hence she arguably has no obligation to we-frame and we-reason in this situation.

It follows from Schwenkenbecher's account that Agnetha and Benny would not have a collective obligation in any of these situations. But intuitively, Agnetha and Benny would together have a collective obligation to save the two children in all three situations (cf. Dietz 2016: 962; Björnsson 2020: 133). Were they not to fulfil this obligation due to incontinence or unwillingness, then they would be collectively blameworthy for the deaths of the children. This shows that the strong connection that Schwenkenbecher draws between we-reasoning and collective obligation needs to be relaxed to a mere capacity-requirement.

One might object that it is unclear, at least in the cases involving incontinence or unwillingness, whether Agnetha and Benny would really have the joint ability to save the two children (the ability required by condition (2)). If this were the case, then our claim that the concept of collective obligation is irreducibly collective could be disputed.

Typically, the realization of a joint ability is just a fitting combination of individual abilities in the right context, including individual abilities to respond appropriately to the exercise of

others' individual abilities.²⁶ Since individual unwillingness to Φ does not deprive an individual of her ability to Φ , this realization relation provides some reason for thinking that a group member's unwillingness to do her part in ϕ -ing does not deprive a group of their joint ability to ϕ . To illustrate, consider the following claim by Barbara Vetter (2015: 105): "If Hannah has the ability to play the piano and Jane has the ability to play the flute, then Hannah and Jane together have the ability to play a duet for flute and piano." This is essentially correct, even if somewhat simplified. Vetter is taking for granted that Hannah and Jane each have individual abilities to respond and adapt appropriately to the other's exercise of her abilities. This is required for them to play a duet. But Vetter need not also take for granted that Hannah and Jane are each willing to contribute to the duet for them to have the ability to play it. Each might be steadfastly unwilling to do so, but clearly, they nevertheless together have the joint ability to do it. This would be accepted by many in the collective obligations literature (see McKinsey 1981: 316, 323 n. 5; Cripps 2013: 45–46; Pinkert 2014: 196, 201; Aas 2015: 15–16; Björnsson 2020: 128–129; Schwenkenbecher 2021: 106–107). However, Collins (2019: 70–72) and Hindriks (2019: 213) argue that the individual unwillingness *of others* can undermine an agent's individual ability to contribute, so that that agent and the others are thereby deprived of their joint ability. Since the agent's individual ability is part of what realizes the joint ability, if this ability is undermined, it would follow that the joint ability is undermined too.

Hindriks (2019) argues that a group's joint ability depends on not just any relations between individual abilities, but on relations of *reliance* between the bearers of those abilities. Agents such as Agnetha and Benny must be able to "rely on each other to contribute to a coordinated effort" in order to have a collective obligation (Hindriks 2019: 214). Importantly, they must be able to rely not only on each other's relevant abilities, but also on each other's willingness to exercise them. Suppose that Agnetha is unwilling to contribute. Hindriks thinks that this makes Benny unable to contribute to a coordinated effort. Benny is deprived of his specific ability to contribute, as he lacks something that is required for him to make his

²⁶ Joint abilities are multiply realizable (see Vetter 2015: 114; Collins 2019: 73, 76–77).

contribution, namely Agnetha's complementary contribution. This would further mean that Agnetha's "unwillingness also incapacitates the collective as a whole" (Hindriks 2019: 213). Benny's specific ability to contribute is a necessary part of the realization of the joint ability to save the two children, and according to Hindriks, Benny is deprived of this specific ability due to Agnetha's unwillingness.

Hindriks's construal of the situation might be correct when we consider the abilities of Agnetha_i and Benny_i. But this construal is misleading when it comes to Agnetha and Benny, who have capacities for group identification and team reasoning. When we ascribe a collective obligation to Agnetha and Benny, their relevant deliberative question is the moral team-reasoning question: "What ought *we* do?" It is therefore inappropriate for either of them to treat the other's willingness or unwillingness as a fixed parameter of their own environment. Rather, they should each treat it as something that ought to be determined by the result of the team reasoning. Agnetha's unwillingness to contribute removes Benny's *opportunity to exercise* his specific ability to do his part in saving the two children, but it does not *deprive him of the ability*; nor therefore, does it deprive them of their joint ability; nor, finally, does it remove the opportunity for their joint ability to be exercised. The fact that Agnetha's unwillingness removes Benny's opportunity to exercise his specific ability to contribute implies that he lacks the even more specific ability to contribute to saving-the-two-children-when-Agnetha-is-unwilling.²⁷ But given that we are concerned with collective obligations, this way of specifying Benny's ability to contribute involves specifying something (Agnetha's agency) which is part of the relevant unit of agency (Agnetha and Benny's agency) as part of a different and irrelevant unit of agency's (Benny's agency's) fixed circumstances.

Consider an analogy with self-reliance. In the individual case, reliance on one's own future willingness to follow through on earlier decisions or ongoing projects is typically taken for granted. When such reliance is lacking, Hindriks might say, this deprives the agent of her ability

²⁷ Abilities can be more or less general or specific, and which kind of specific ability is relevant depends on context (for a rich and useful discussion, see Jaster 2020: section 4.5).

to implement a decision or complete a project. Siding with “actualists” about obligation, Hindriks might then say that the agent could therefore not be the subject of an obligation to do what has been decided or initiated.²⁸ While we realize that there is room for reasonable disagreement here, we find this counterintuitive.²⁹ If Hindriks agrees, then he owes us an argument for why unwillingness would play a different role for joint ability than it plays for individual ability.

Collins (2019: 70–72) provides a somewhat different argument for why another individual’s unwillingness to contribute can deprive an agent of her own ability to contribute, and thereby deprive the group of their joint ability. Her starting point is that an agent lacks an ability to Φ if she is unable to try to Φ . If I have acrophobia, then I lack the ability to go up the Eiffel Tower, even though I would ascend it *were I to try*. But since I lack the ability to try to do this, I also lack the ability to do it. According to Collins, if Benny knew that Agnetha was unwilling, then he would be similarly unable to rationally *try* to contribute. He would therefore lack the ability to contribute to saving the two children. Collins thus thinks that in the context of joint ability, the other’s unwillingness deprives an agent of their ability to contribute, in a way that is relevantly similar to how their own phobia could deprive them of their ability to do something.

However, having a phobia and being in the company of someone who is unwilling to contribute are not relevantly similar. If Benny were unable to rationally try to contribute due to an actual phobia, then he would indeed be unable to contribute to saving the two children (see Pinkert 2014: 196 n. 26). One of the necessary realizers of what would otherwise be their joint ability to save the two children would then really be missing. Furthermore, given that Benny would be unable to contribute, it would be appropriate to count this inability as part of the fixed environment relative to which we specify Agnetha’s specific abilities to contribute. Thus, she would lack the specific ability to contribute to saving-the-two-children-without-Benny’s-help,

²⁸ For discussion of how the debate between actualism and possibilism relates to team reasoning, see Woodard 2011.

²⁹ The abilities that are relevant for our obligations are arguably “conative abilities,” where these are specific abilities that abstract away from agents’ motivational states (see Jaster 2020: 117, 120). As Romy Jaster points out, such abilities “seem to underlie much of our thinking about whether or not agents could have done otherwise” (2020: 120).

which would be the relevant specific ability in light of Benny's phobia. But the unwillingness of a potential cooperator does not have this effect on the appropriate specification of an agent's ability to contribute. Agnetha's unwillingness rids Benny of his opportunity to exercise his ability to contribute. It does not deprive him of that ability. With respect to trying, the crucial question is whether *they*, Agnetha and Benny, can try to save the two children, not whether each of them, holding the other's unwillingness fixed, is able to (rationally) try.

Given that Agnetha and Benny are each capable of group-identification and team reasoning, the fact that one or both of them is unwilling to contribute does not itself deprive them of their relevant individual abilities, nor does it deprive them of their joint ability to save the two children. We take this to be in line with common-sense thought and talk, and supported by the view that joint abilities are realized by individual abilities, as well as being supported by the analogy with temporally extended agency. If an agent is reluctant to try to initiate a project today because she thinks that she will be unwilling to follow through on it tomorrow, then this does not deprive her of her ability to do it.³⁰ With that, we hope to have presented a plausible case for the claim that individual unwillingness does not deprive a group of their joint ability. This helps to support the view that an ascription of a collective obligation cannot be reductively understood as a constellation of ascriptions of obligations to individuals.

Before ending this section, let us briefly address what may seem to be a problematic implication of the sort of non-reductive view of collective obligation that we are providing. Suppose that each of Agnetha and Benny has the capacity to group-identify and team reason. Agnetha identifies with the group, reasons as a team member, and runs to the second room, where she expects to meet Benny. However, Benny does not exercise his capacities to group-identify and team reason, or perhaps he does so but is unwilling to do his part of what team reasoning recommends. In this situation, Benny is clearly blameworthy. Agnetha, on the other hand, has done nothing wrong and is therefore not blameworthy. It may seem unclear in what sense, if any,

³⁰ Our discussion of individual unwillingness and joint ability has benefited from comments from Niels de Haan and Stephanie Collins.

Agnetha and Benny could be collectively blameworthy for not having done what was required of them. But if they were not collectively blameworthy, then this would suggest that there was never any collective obligation which they could have violated.³¹

While it is true that Agnetha has done nothing wrong, it is also true that Agnetha and Benny have done something wrong. After all, they could have saved the two children, but failed to do so. Benny is indirectly morally responsible and blameworthy for the death of the two children, but Agnetha and Benny are also together morally responsible and blameworthy for those deaths. Here it is important to realize that collective and individual blameworthiness can come apart in such a way that Agnetha is not substantially tainted as a private person by the collective blameworthiness she shares with Benny (see Schwenkenbecher 2021: 108–109). Nevertheless, given that she has group-identified, guilty feelings on her part, *qua* group member, understood as feelings held by the individual from the group perspective, would arguably be apt (see Petersson 2020).³²

5. *The perspective-relativity of moral obligations*

We have argued that ascriptions of collective obligations assume a group perspective. The ascribed obligation is relative to a perspective of collective rationality. Our proposal therefore implies a kind of agential perspective-relativity with respect to the appropriate ascription of obligations. The group perspective that the obligation is relative to does not have to actually be occupied by any group member; it is rather a perspective that all must be *capable* of occupying in the situation at hand. This means that it is possible for group members to have a collective obligation (relative to the group perspective) while simultaneously having individual permissions or obligations (relative to the I-perspective of each of them) such that what is permitted or obligatory from the I-perspective could not be done without violating what is

³¹ Jules Salomone-Sehr raised this issue with us.

³² We are grateful to Gunnar Björnsson and an anonymous reviewer for comments.

obligatory from the group perspective. With the help of a few example cases, we will now illustrate these possibilities.

Consider the following variation of *Burning Building*:

More Children: Five children are trapped in a burning building. Agnetha and Benny approach the building from opposite sides. One child is in room A, another is in room B, and the remaining three are in room C. The three rooms are all some distance away from each other. Only room C is surrounded by fire. Agnetha has enough time to either rescue the child in room A on her own, or do her part in rescuing the three children in room C together with Benny (who has a fire extinguisher). Benny has enough time to either rescue the child in room B on his own, or do his part in rescuing the three children in room C together with Agnetha. If Agnetha goes to room A, or if Benny goes to room B, then the three children in room C will all die. If Agnetha and Benny together rescue the children in room C, then the children in room A and room B will die. Suppose that Agnetha and Benny can make their choices without any significant risk to their own or each other's life or health. All this is common knowledge between them, but they do not have the opportunity to communicate—each must choose which room to head for independently of the other.

Suppose that Benny and Agnetha should each simply see to it that as many children as possible are saved, that is, that morality is agent-neutral in the sense that it does not matter for the moral worth of any of their actions whether they themselves do any actual saving. Suppose also that this is reflected in their moral preferences. If they reason from the I-perspective, then *More Children* will, like *Burning Building*, be a Hi-Lo game for them, where the question "Where should I go?" has no determinate answer (the Lo equilibrium is the outcome where two children are saved, the Hi equilibrium is the outcome where three children are saved, and the two non-equilibria are the outcomes where one child is saved). If they instead reason from the we-perspective, then each of them should conclude "We should go to C," and thus infer that they themselves should go to C. So,

when considered from the I-perspective, the obligation of each is conditional on the other's choice. But when considered from the we-perspective, given the presupposition that the other will do their part, they each have an unconditional obligation to do their part in saving the three children. In that sense, what morality prescribes is relative to the agential perspective.

The perspective-relativity of moral obligations will be more salient on other moral views. Assume that morality is agent-relative in the sense that it assigns for each agent a higher value to "I save a child" than to "A child is saved," and that this is reflected in Benny's and Agnetha's moral preferences.³³ As I-reasoners, they would then face a moral Stag Hunt game:

		Benny	
		room B	room C
Agnetha	room A	1,1	1,0
	room C	0,1	3,3

Figure 2 *More Children* (with agent-relative morality)

Each agent here knows that if they run to room A/B, they will save a child. If they instead run to room C, the outcome is uncertain (the probability that they will save the three children together with the other is unknown). In this situation, and given this assumption about morality, it would from the I-perspective be permissible for each agent to run to A/B in order to save a child themselves.³⁴

³³ Derek Parfit assumes that "M," or "Common Sense Morality," allows for agent-relativity by letting agents give higher weight to people related to them in special ways, like family members, but also by telling the agent that what *she does* matters, not only *what happens*. Both features can make M directly self-defeating in situations with a Prisoner's Dilemma structure. Parfit argues that M should be complemented with a proviso, telling agents that in such special cases "we should all ideally do what will cause the M-given aims of *each* to be better achieved," given that sufficiently many others do so as well (1986: section 38). This move, Parfit claims, brings a non-self-defeating version of M closer to an agent-neutral position, like C, Consequentialism (1986: section 42).

³⁴ On some views, going to room A/B would for two I-reasoners not only be permissible given an agent-relative morality, but obligatory. For example, this could be the case given what Parfit calls the "Share-of-the-Total View" which roughly says that when we produce benefits together, each can take credit for her share of the total benefit (1984: section 25). If Agnetha and Benny save three lives together, the moral worth of Agnetha's action would on this view equal 1.5 lives. If it is rational for Agnetha to take the

Derek Parfit points out that an agent acting on an agent-relative morality that allows her to assign a higher value to “I save a child” than to “A child is saved” would sometimes do better in terms of her own agent-relative morality if she allowed for substantive exceptions to agent-relativity in special circumstances. This could be the case if she finds herself in a situation with a moral Prisoner’s Dilemma structure (see Parfit 1987: sections 36–38). But as indicated above, in *More Children*, switching from agent-relative to agent-neutral I-reasoning will merely turn the Stag Hunt into a Hi-Lo game, bringing back the indeterminacy that we have already discussed.

On the other hand, even without allowing for substantive exceptions to agent-relativity, an adherent of such morality who identifies with her group and reasons from the group perspective would face no dilemma in *More Children*. Arguably, if such an agent group-identifies and sees the situation from the perspective of the group, then what becomes especially morally valuable would be for the children to be saved by the group that the agent identifies with—that they are saved by *us*, and not just that they are saved.

This means that the adherent of agent-relative morality in *More Children* faces a tension between what morality prescribes from the group perspective (you must do your part in our saving the three children in room C) and what it prescribes from the individual perspective (you are allowed to save the single child in room A/B; after all you are guaranteed to save at least one child, while it is uncertain which result running to room C will bring about). Nothing rules out the possibility that Agnetha and Benny could each have both a context-dependent capacity for group identification and team reasoning and a context-dependent capacity for “I-identification” and I-reasoning.

It is even possible for an agent to have an individual obligation relative to the I-perspective, which is in irresolvable practical conflict with a collective obligation that the agent

probability of Benny’s going to C to be 0.5 (suppose that the principle of insufficient reason should be accepted), then the morally optimal choice for Agnetha would be to go to room A, where with a probability of 1 she will save one life, rather than to go to C and produce what equals the worth of 1.5 lives with a probability of 0.5. But even if one accepts this admittedly dubious assumption about how to count individual contributions to benefits, the moral prescription from the we-perspective would still be that Agnetha should do her part in their saving of the three children.

has together with others. In a discussion of the agent-relativity of reasons for action, Alexander Dietz offers a potential example involving a special obligation to an agent's own child: "*Rescue Mission*: You and I are about to carry out a rescue mission to save the lives of two strangers in imminent danger. But I then learn that my child's life is also in danger. If I continue with our rescue mission, there will not be enough time to save my child" (2022: 58). Dietz suggests that for the agent, "I," the situation "will feel *paralyzing*: whenever we consider what we ought to do, our collective disposition to do what we ought will drive me to one conclusion about what to do, but whenever I remember what I ought to do, my individual disposition to do what I ought will drive me to another" (2022: 65). In terms of our group-identification framework, the paralysis would be the effect of the agent vacillating between the we- and the I-perspective.

Here we leave aside questions about possible conflicts between the first-order moral views held by agents who face moral dilemmas together, as well as questions about the plausibility of agent-neutral or agent-relative moral positions. In the cases under consideration, we assume that the values assigned to each option reflect the moral worth of that option as determined by some substantive moral view—an agent-neutral and an agent-relative view, respectively—and that these values correspond to both agents' moral preferences. The point of this digression about two substantively different normative positions is to highlight different ways in which moral obligations may diverge, solely depending on agential perspective.

On the assumption of agent-relative morality, a bystander could either demand of Agnetha and Benny, considered collectively, that they must save the three children, or else advise one of them, considered as an individual, to run to room A/B to save a child. It may seem that each of them will then face a moral dilemma. Each of them will have both an obligation to do their part in saving the three children and an ordinary individual obligation to go to room A/B to save a child. But there is, strictly speaking, no moral dilemma here. The collective obligation to save the three children and the individual permission to save a child are not morally conflicting, but rather incommensurable. Suppose that Agnetha runs to room A. Relative to the I-perspective, she will have acted permissibly, while relative to the we-perspective, she and Benny will have done

wrong. As a result of this incommensurability, agents may find themselves in practical conflicts that cannot be rationally solved. If Agnetha occupies the I-perspective, then she will judge that it is permissible for her to save the child in room A. But if she occupies the we-perspective, then she will judge that she ought to do her part in saving the three children together with Benny. If she is able to directly or indirectly switch between occupying these two perspectives, then there will be no “final” answer to the question of whether it is morally permissible for her to save the child in room A.

If there were a neutral perspective from which agents could objectively weigh the reasons for and against taking the I/we-perspective, then they could escape any such irresolvable practical conflicts. According to Schwenkenbecher, each individual agent can choose which perspective to adopt in light of each agent having a “reason to believe that the collectively available option (joint rescue) is morally best,” as well as a “reason to include that option in her deliberation about her obligations (we-framing the problem)” (2021: 19). As far as we can tell, Schwenkenbecher thinks that these reasons are perspective-neutral, grounded in the fact that the collectively available option is morally most valuable. But there is arguably no neutral perspective from which such reasons would be accessible to deliberating agents (see Pacherie 2011: 186–187; Gold and Colman 2018: 313–314; Dietz 2022: 68). Even if the we-perspective can be chosen—indirectly, or perhaps even directly—it will always be rationally chosen relative to an agential perspective.³⁵

One might think that the perspective-relativity of I/we-reasoning only makes requirements of rationality relative to the agential perspective, not requirements of morality. In this vein, Collins (2019: 141) acknowledges that “instrumental practical reasoning presupposes

³⁵ José Luis Bermúdez argues that fairness is a value that is “frame-neutral” (2021: 206), and that “considerations of fairness can be raised from within a perspective that includes both the ‘I’-frame and the ‘we’-frame and allows them to be compared to each other” (2021: 194). In particular, he argues that agents can be led from the I-perspective to the we-perspective through reflection on the unfair distributions of payoffs in the so-called Chicken game (2021: 206–208, 211–212). He also suggests that reflection on other values that lie behind agents’ preferences might similarly lead agents from the I-perspective to the we-perspective in a rational way (2021: 206). But fairness considerations are irrelevant in Hi-Lo, and we do not see how the value of human lives is frame-neutral in such a way that Agnetha or Benny could rationally choose to switch from the I-perspective to the we-perspective.

a unit of agency” (quoting Gold and Sugden 2007: 124), but she does not think that this has any impact on what moral reasoning can require: “Thankfully, I’m concerned with moral duties, not duties of rationality. [...] The agents should engage in whichever form of reasoning—I-reasoning or coalition-reasoning—will enable them to do their coordination duties in the morally optimal way” (Collins 2019: 141). Likewise, in a discussion of what Natalie Gold and Andrew Colman see as failed attempts to show that agents can rationally choose whether or not to group-identify, they seem to assume that it is somehow easier to show that agents could be morally right in group-identifying and team reasoning (2018: 313–314).

The morally best outcome of *More Children* is clearly the one where both Agnetha and Benny run to room C to save the three children there. But their collective obligation to save the three children is nevertheless relative to the group-perspective. Granted, if agents could solve Hi-Lo through what Collins calls “coalition-reasoning,” then Collins could claim that what grounds the individual obligation of each agent is not that they ought to save the three children (their collective obligation), but merely that saving the three children would be morally best. As we have indicated, we doubt that coalition-reasoning provides a satisfactory solution to Hi-Lo. If it doesn’t, then the fact that it would be morally best to save the three children doesn’t mean that what Agnetha or Benny morally ought to choose or do is to run to room C. If Benny (Agnetha) runs to room B (A), then it is better if Agnetha (Benny) runs to room A (B), rather than to room C. So, not only instrumental practical reasoning presupposes a unit of agency, but moral practical reasoning does so as well. This means that deliberative obligations presuppose a unit of agency. Thus, for several agents to have a deliberative obligation together, they must each have the capacity to reason about what they together, as one unit of agency, ought to do.

6. Conclusion

When we ascribe an obligation to several agents collectively, we presume that they each have a context-specific capacity to view their situation from their group’s perspective and at least a

general capacity to deliberate about what they ought to do—to “team reason.” If the agents’ external conditions are such that group identification is to be expected, then this will make the ascription of an obligation to act collectively seem more apt. A live hypothesis that we find plausible is that group identification can be triggered by the strong interdependence that exists between individuals’ interests in a social dilemma such as Hi-Lo. But there may also be other favourable external conditions, such as the group being confined to a limited common space, there being no obstacles to communication or signalling preparedness to coordinate, etc. It seems to us that the strength of our intuitions about several agents’ collective obligation will vary with the presence or absence of such features, in a way that supports our theory.

When we ascribe a collective obligation to a group, we conceive of the individuals as group members, and we assume a group-perspective when we consider what the members ought to do together. The “ought” is thus relative to a group perspective. It is also possible to single out the individual perspective of each group member, and consider what they themselves ought to do. We have shown that in some situations, the “I”-relative oughts and the “we”-relative oughts require or permit different actions. Such cases are examples of genuine moral incommensurability. This is a somewhat surprising result, but it follows from a combination of reasonable considerations about obligations and plausible socio-psychological theorizing about group identification.

Department of Philosophy, Linguistics and Theory of Science, University of Gothenburg
olle.blomberg@gu.se

Department of Philosophy, Lund University
bjorn.petersson@fil.lu.se

Acknowledgements

Earlier versions of this paper were presented at the “Collective and Shared Responsibility” workshop at MANCEPT 2019, the “Group Agency and Collective Responsibility” workshop, Flensburg, 2019, the Higher Seminar in Practical Philosophy in Lund, the Practical Philosophy and Political Theory research seminar in Gothenburg, and at the online Social Ontology 2020 conference. We are grateful for useful comments, questions, and criticisms from participants at all these events. We want to thank the following particular individuals (apologies to those we have forgotten): Gunnar Björnsson, Stephanie Collins, Hadi Fazeli, Mattias Gunnemyr, Niels de Haan, Frank Hindriks, Yuliya Kanygina, Arto Laitinen, Kirk Ludwig, Per-Erik Milam, Jules Salomone-Sehr, Michael Schmitz, David Schweikard, David Shoemaker, Caroline Touborg, and Robert Williams. In addition, we are very grateful to this journal’s two anonymous reviewers for two rounds of excellent comments, as well as to several anonymous reviewers of earlier versions of this paper that were submitted to other journals. Our research was funded by the Swedish Research Council (Vetenskapsrådet) project grant 421-2014-1025, and by the Lund Gothenburg Responsibility Project (PI: Prof. Paul Russell), which is in turn also funded by the Swedish Research Council.

References

- Aas, Sean. 2015. "Distributing Collective Obligation," *Journal of Ethics and Social Philosophy* 9(3): 1–23.
- Bacharach, Michael. 2006. *Beyond Individual Choice: Teams and Frames in Game Theory*, Princeton: Princeton University Press.
- Bermúdez, José Luis. 2021. *Frame It Again: New Tools for Rational Decision-Making*. Cambridge: Cambridge University Press.
- Björnsson, Gunnar. 2014. "Essentially Shared Obligations," *Midwest Studies in Philosophy* 38(1): 103–120.
- Björnsson, Gunnar. 2020. "Collective Responsibility and Collective Obligations Without Collective Moral Agents," in *The Routledge Handbook of Collective Responsibility*, ed. Saba Bazargan-Forward and Deborah Tollefsen. New York: Routledge, 127–141.
- Blomberg, Olle. 2020. "What We Ought to Do: The Decisions and Duties of Non-agential Groups," *Journal of Social Ontology* 6(1): 101–116.
- Blomberg, Olle, and Björn Petersson. 2018. "Plikt att kollektivera?," *Tidskrift för politisk filosofi* 2018(2): 36–46.
- Brewer, Marilynn B., and Roderick M. Kramer. 1986. "Choice Behavior in Social Dilemmas: Effects of Social Identity, Group Size, and Decision Framing," *Journal of Personality and Social Psychology* 50(3): 543–549.
- Clarke, Randolph. 2017. "Ignorance, Revision, and Commonsense," in *Responsibility: The Epistemic Condition*, ed. Philip Robichaud and Jan Willem Wieland. Oxford: Oxford University Press, 233–251.
- Collins, Stephanie. 2019. *Group Duties: Their Existence and Their Implications for Individuals*. Oxford University Press.
- Collins, Stephanie. 2020. "Response to Critics," *Journal of Social Ontology* 6(1): 141–157.
- Colman, Andrew M., Briony D. Pulford, and Catherine L. Lawrence. 2014. "Explaining Strategic Coordination: Cognitive Hierarchy Theory, Strong Stackelberg Reasoning, and Team Reasoning," *Decision* 1(1): 35–58.
- Colman, Andrew M., and Natalie Gold. 2018. "Team reasoning: Solving the puzzle of coordination," *Psychonomic Bulletin & Review* 25(5): 1770–1783.
- Copp, David. 1991. "Responsibility for Collective Inaction," *Journal of Social Philosophy* 22(2): 71–80.
- Cripps, Elizabeth. 2013. *Climate Change and the Moral Agent: Individual Duties in an Interdependent World*. Oxford: Oxford University Press.
- Dawes, Robyn M., Alphons J. C. Van De Kragt, and John M. Orbell. 1990. "Cooperation for the Benefit of Us—Not Me, or My Conscience," in *Beyond Self-Interest*, ed. Jane J. Mansbridge. Chicago: Chicago University Press, 97–110.
- Dietz, Alexander. 2016. "What We Together Ought to Do," *Ethics* 126(4): 955–982.
- Dietz, Alexander. 2022. "Collective Reasons and Agent-Relativity," *Utilitas* 34(1): 57–69.

- Gilbert, Margaret. 1989. "Rationality and salience," *Philosophical Studies* 57(1): 61–77.
- Gold, Natalie. 2018. "Team Reasoning and Spontaneous Collective Intentions," *Revue d'économie politique* 128(3): 333–353.
- Gold, Natalie, and Andrew M. Colman. 2018. "Team Reasoning and the Rational Choice of Payoff-Dominant Outcomes in Games," *Topoi* 39(2): 305–316.
- Gold, Natalie, and Robert Sugden. 2007. "Collective Intentions and Team Agency," *Journal of Philosophy* CIV(3): 109–137.
- Grünbaum, Thor, and Søren Kyllingsbæk. 2020. "Is Remembering to do a Special Kind of Memory?," *Review of Philosophy and Psychology* 11(2): 385–404.
- Hakli, Raul, Kaarlo Miller, and Raimo Tuomela. 2010. "Two Kinds of We-Reasoning," *Economics & Philosophy* 26(3): 291–320.
- Held, Virginia. 1970. "Can a Random Collection of Individuals be Morally Responsible?," *The Journal of Philosophy* 67(14): 471–481.
- Helm, Bennett W. 2008. "Plural Agents," *Noûs* 42(1): 17–49.
- Hindriks, Frank. 2019. "The Duty to Join Forces: When Individuals Lack Control," *The Monist* 102(2): 204–220.
- Jaster, Romy. 2020. *Agents' Abilities*. Berlin: De Gruyter.
- Kramer, Roderick M., and Marilyn B. Brewer. 1984. "Effects of Group Identity on Resource Use in a Simulated Commons Dilemma," *Journal of Personality and Social Psychology* 46(5): 1044–1057.
- Lawford-Smith, Holly. 2015. "What 'We'?", *Journal of Social Ontology* 1(2): 225–249.
- Linnebo, Øystein. 2022. "Plural Quantification," *The Stanford Encyclopedia of Philosophy (Spring 2022 Edition)*, ed. Edward N. Zalta, URL = <https://plato.stanford.edu/archives/spr2022/entries/plural-quant/>.
- Lord, Errol. 2015. "Acting for the Right Reasons, Abilities, and Obligation," in *Oxford Studies in Metaethics 10*, ed. Russ Shafer-Landau. Oxford: Oxford University Press, 26–52.
- Ludwig, Kirk. 2016. *From Individual to Plural Agency – Collective Action: Volume 1*. Oxford: Oxford University Press.
- Manning, Rita. 1985. "The Random Collective as a Moral Agent," *Social Theory and Practice*, 11(1): 97–105.
- McKinsey, Michael. 1981. "Obligations to the Starving," *Noûs* 15(3): 309–323.
- Pacherie, Elisabeth. 2011. "Framing Joint Action," *Review of Philosophy and Psychology* 2(2): 173–192.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Clarendon Press.
- Peterson, Martin. 2017. *An Introduction to Decision Theory: Second Edition*. Cambridge, UK: Cambridge University Press.
- Petersson, Björn. 2020. "Collective Guilt Feelings," in *The Routledge Handbook of Collective Responsibility*, ed. Saba Bazargan-Forward and Deborah Tollefsen. New York: Routledge, 228–242.

- Petersson, Björn. 2017. "Team Reasoning and Collective Intentionality," *Review of Philosophy and Psychology* 8(2): 199–218.
- Pinkert, Felix. 2014. "What We Together Can (Be Required to) Do," *Midwest Studies in Philosophy* 38(1): 187–202.
- Postema, Gerald J. 1995. "Morality in the First-Person Plural," *Law and Philosophy* 14(1): 35–64.
- Roth, Abraham Sesshu. 2014. "Team Reasoning and Shared Intention," in *Institutions, Emotions, and Group Agents: Contributions to Social Ontology*, ed. Anita Konzelmann Ziv and Hans Bernhard Schmid. Dordrecht, Netherlands: Springer, 279–295.
- Salice, Alessandro, and Kengo Miyazono. 2020. "Being one of us. Group Identification, Joint Actions, and Collective Intentionality," *Philosophical Psychology* 33(1): 42–63.
- Salice, Alessandro, and Alba Montes Sánchez. 2016. "Pride, Shame, and Group Identification," *Frontiers in Psychology*, 7: 557
- Salice, Alessandro, and Alba Montes Sánchez. 2019. "Envy and Us," *European Journal of Philosophy* 27(1): 227–242.
- Schweikard, David P., and Hans Bernhard Schmid. 2021. "Collective Intentionality," *The Stanford Encyclopedia of Philosophy (Fall 2021 Edition)*, ed. Edward N. Zalta, URL = <<https://plato.stanford.edu/archives/fall2021/entries/collective-intentionality/>>.
- Schwenkenbecher, Anne. 2019. "Collective Moral Obligations: 'We-Reasoning' and the Perspective of the Deliberating Agent," *The Monist* 102(2): 151–171.
- Schwenkenbecher, Anne. 2021. *Getting Our Act Together: A Theory of Collective Moral Obligations*. New York: Routledge.
- Searle, John. 1997. *The Construction of Social Reality*. New York: Free Press
- Sinigaglia, Corrado, and Stephen A. Butterfill. 2022. "Motor representation in acting together." *Synthese* 200(2): 1–16. <https://doi.org/10.1007/s11229-022-03539-8>
- Smith, Thomas H. 2009. "III-Non-Distributive Blameworthiness," *Proceedings of the Aristotelian Society* 109(1): 31–60.
- Sugden, Robert. 1993. "Thinking As a Team: Towards an Explanation of Nonselfish Behavior," *Social Philosophy and Policy* 10(1): 69–89.
- Vetter, Barbara. 2015. *Potentiality: From Dispositions to Modality*. Oxford: Oxford University Press.
- Woodard, Christopher. 2011. Rationality and the Unit of Action. *Review of Philosophy and Psychology* 2(2): 261–277.
- Wringe, Bill. 2010. "Global Obligations and the Agency Objection," *Ratio* 23(2): 217–231.
- Wringe, Bill. 2016. "Collective Obligations: Their Existence, Their Explanatory Power, and Their Supervenience on the Obligations of Individuals," *European Journal of Philosophy* 24(2): 472–497.
- Wringe, Bill. 2020. "The Duties of Non-Agential Groups: Some Comments on Stephanie Collins' Group Duties," *Journal of Social Ontology* 6(1): 117–125.