

Collective Moral Obligations: ‘We-Reasoning’ and the Perspective of the Deliberating Agent

Anne Schwenkenbecher*

ABSTRACT

Together we can achieve things that we could never do on our own. In fact, there are sheer endless opportunities for producing morally desirable outcomes *together with others*. Unsurprisingly, scholars have been finding the idea of collective moral obligations intriguing. Yet, there is little agreement among scholars on the nature of such obligations and on the extent to which their existence might force us to adjust existing theories of moral obligation. What interests me in this paper is the perspective of the moral deliberating agent who faces a collective-action problem, i.e., the type of reasoning she employs when deciding how to act. I hope to show that agents have collective obligations precisely when they are required to employ ‘we-reasoning’, a type of reasoning that differs from I-mode, best-response reasoning. More precisely, two (or more) individual agents have a collective moral obligation to do x if x is an option for action that is *only collectively available* and each has sufficient reason to rank x highest out of the options available to them.

INTRODUCTION

Together we can achieve things that we could never do on our own. In fact, there are sheer endless opportunities for producing morally desirable outcomes *together with others*. Unsurprisingly, scholars have been finding the idea of collective moral obligations intriguing. Yet, there is little agreement among scholars on the nature of such obligations and on the extent to which their existence might force us to adjust existing theories of moral obligation. More ‘revisionist’ scholars are of the view that individual agents who are not (yet) in any way organised can hold obligations *as a group* (Held 1970; May 1992; Wringer 2005; Wringer 2010; Isaacs 2011). Others argue that individual agents can *share* moral obligations or hold them *jointly* (Björnsson 2014; Pinkert 2014; Schwenkenbecher 2013; Schwenkenbecher 2014; Björnsson forthcoming). More ‘conservative’ approaches insist that there is nothing so special about collective obligations but that individual agents may simply be required to work towards collective outcomes sometimes (Collins 2013; Lawford-Smith 2015). There is great diversity amongst existing views and this taxonomy is a rough one. However, I will not concern myself with the differences between these

*Murdoch University

views now—I have done so elsewhere (Schwenkenbecher 2018) and it is not essential to my argument here.

What interests me in this paper is the perspective of the moral deliberating agent who faces the kind of collective-action problem that motivates the above-mentioned views, i.e., the type of reasoning she employs when deciding how to act. I hope to show that agents have collective obligations precisely when they are required to employ ‘we-reasoning’, a type of reasoning that differs from I-mode, best-response reasoning, as I shall explain below. More precisely, two (or more) individual agents have a collective moral obligation to do x if x is an option for action that is *only collectively available* (more on that later) and each has sufficient reason to rank x highest out of the options available to them.

Note that this enquiry is concerned with collective obligations of agents who are not (yet) organised as a group¹ and do not form what some scholars call group agents (List and Pettit 2011; Tollefsen 2015). As such, the focus is on relatively simple actions that require coordination, but not complex organisation.²

The paper will proceed as follows: section (1) introduces two motivating cases; section (2) explains the ideas of ‘we-reasoning’ and ‘collectively available options’; (3) shows how ‘we-reasoning’ can yield results that ‘I-mode reasoning’ does not yield and suggests that all of the above-mentioned accounts implicitly assume that moral agents should we-reason. Section (4) argues that two agents have a collective obligation if they have conclusive reason to pursue a collectively available option, and section (5) deals with potential objections.

1. TWO COLLECTIVE RESCUE CASES

The cases motivating most discussions of collective obligations are characterised by *joint necessity*.³ This is a feature of actions (and outcomes) that cannot be performed (or produced) by one person on their own, but require at least two people in order to be realized. Joint necessity is *analytic* where it is part of what it means to do x that x is done by at least two people, as in ‘getting married’. It is *circumstantial* where, as a matter of fact (but not as a matter of principle), an action (or outcome) cannot be performed (or produced) by one person alone, for instance, if it takes two or more people to lift a heavy object.

Further, the cases motivating accounts of collective obligations tend to be *collective rescue* scenarios where individual agents must spontaneously collaborate in order to assist someone in urgent need. I will be using two such cases to motivate my argument.

The first is a two-person *strict joint necessity* case (more about that term in a moment).⁴

Hikers: Two hikers encounter a third while hiking in a remote area. The third hiker is trapped underneath a fallen tree. In order to free the trapped hiker, the two others must collaborate and lift the tree together. None of them can do it individually, but together they can. If they do not lift the tree, the trapped person is likely to die. There is no one else to help.

Most people—including both revisionist and conservative scholars—agree that the two hikers have duties to do something about the morally dire situation; that they have duties to assist the person in need.⁵ How to properly account for this basic intuition is what the above-mentioned approaches disagree on.⁶

Hikers is a *strict* joint necessity case. This means that all available helpers are necessary for the joint action to succeed or the joint outcome to be produced. It takes a minimum of two people to lift the tree and there are exactly two people to help.

This differs from *wide* joint necessity cases, where there are more available helpers than minimally required, as in my second example:

Commuters: On a busy weekday morning a man gets trapped between a commuter train and the station's platform. He will be crushed should the train move. Dozens of people who happen to be on the platform witnessing his predicament join forces in pushing the train to tilt it away from the man. Together they manage to free him, therewith saving his life.⁷

Again, I am assuming that there is general agreement that the trapped man ought to be assisted by the people nearby (or else that we can easily fill in enough details for this case to generate such agreement). In each of the two cases, the morally best outcome is only collectively available. And it is this outcome that—ideally—the bystanders should pursue over and above individually available outcomes. In order for this to be the case, the collectively available outcome must be an option that would reasonably be considered by the individual deliberating agents.

In the following, I will explain how there is a special kind of reasoning involved in choosing collective options, which I call 'we-reasoning' and which competes with 'I-reasoning' in joint necessity scenarios.

2. WE-REASONING EXPLAINED

I am using the term 'we-reasoning' in a specific way here, which is inspired by but not identical to how this and related concepts (such as team-reasoning) are used in different theories of team agency (Gold and Sugden 2007; Sugden and Gold 2007; Hakli, Miller et al. 2010; Tuomela 2013) and in nonstandard game theory (Bacharach 1999; Butler, Burbank et al. 2011; Butler 2012).⁸ Let me explain what I mean by 'we-reasoning' (henceforth used without single quotation marks).

I will call *we-framing* the act of identifying collectively available options and including them when deliberating about which option is best.⁹ *We-reasoning* succeeds we-framing and consists in determining individual strategies or action choices. When an agent we-reasons, she decides what she needs to do in order to bring about the collective outcome she has identified as optimal.

In nonstandard game theory and theories of team agency, we-reasoning is seen as explaining both cooperative behaviour in situations that resemble prisoners' dilemma (PD) and optimal choices in coordination games like stag hunt and the 'hi-lo' game (Bacharach 1999; Hakli, Miller et al. 2010; Butler, Burbank et al. 2011).

The most general way of explaining the concept is this: agents regularly face choices where the outcome of their action will depend on how others choose. There

are two fundamentally different ways in which we can try to optimize the outcome of our choices while ignorant of how the other person chooses. We can think of our choices as best responses to the other's choice (I-reasoning) or we can think of our choices as contributions to the best possible collective outcome or pattern (we-reasoning). Let me illustrate the difference by using a simple coordination game: Hi-lo.¹⁰

Hi-lo is a coordination game with two equilibria,¹¹ whereas one is Pareto-optimal¹² and one is not. Here is the payoff-matrix for a two-player hi-lo game:

Table 1. Payoff-matrix for two-player hi-lo game.

Player 2	A	B
Player 1	Hi/Hi	0/0
B	0/0	Lo/Lo

$$\text{Hi} > \text{Lo} > 0$$

It is assumed that both players know the payoffs and that both players know that both players know the payoffs.¹³ So there is some shared knowledge and some (low) level of common (*de dicto*) knowledge between players, but their individual choices are made independently. It is also assumed that players (if rational) strive for the maximum payoff. If you are player 1, there are two different ways in which you can think about your options. If you think of them as best responses to the other's choices, then you will reason in the following way: "If I think you will choose B then the act it's best for me to choose is B, and if you think I will choose B then it's best for you to choose B . . ." (Bacharach 2006, 44). This approach, however, does not deliver any clear indication that you should select A over B. In fact, because the best-response approach produces a conditional recommendation, it does not give you any clear indication on how you should choose if you want to maximize your payoff (Hakli, Miller et al. 2010).

Alternatively, you can start by identifying the best outcome and reason backwards. For each of you, maximizing your payoff is only possible if the other player makes the corresponding choice. In that sense, it is an outcome you can only achieve together. Michael Bacharach argues that players faced with hi-lo type scenarios *in fact* regularly start their deliberation with the question "what should *we* do?" rather than "what should *I* do?" It is in this sense that, according to Bacharach, they we-reason about their choices: they do not approach their choices as best responses to other players' choices, but in terms of the best collectively achievable outcome.

Most, if not all of us would instantly choose option A over option B in the above game. The abovementioned scholars contend that this is best explained by assuming that players use this different mode of thinking: framing the problem as one for the group and selecting one's individual strategy (or option) accordingly.¹⁴

One might object by pointing to two alternative explanations¹⁵ for selecting A: First, agents may choose A as a matter of maximising expected utility. This objection is best countered by pointing to the fact that the we-reasoning explanation is meant to hold in other strategic interaction cases, too, where expected-utility considerations

will not deliver the Pareto optimal solution; for instance, it can explain the cooperative choice in a prisoners' dilemma (Bacharach 2006, Gold and Sugden 2007).

Second, is the choice not explicable by a simple preference transformation, that is, by the suggesting that agents have a preference for the best group outcome? The short answer to that is 'no' and it is perhaps best to point to the work of Hakli, Miller et al. (2010) who draw a useful distinction between what they call *pro-group I-mode reasoning* and *we-mode reasoning* ('we-reasoning' in my terminology).¹⁶ In pro-group I-mode, agents "select actions that, given their expectations of other agents' actions, best satisfy their preferences, which are group-regarding . . ." (Hakli, Miller et al. 2010, 299ff). Pro-group I-mode thinking cannot eliminate the Pareto-inefficient equilibrium in the hi-lo game. In general, "the we-mode tends to create more collective order than the pro-group mode: It can decrease the amount of equilibria but it cannot increase them" (Hakli, Miller et al. 2010, 306). Only the we-mode *guarantees* that the Pareto-optimal equilibrium be chosen in games like the hi-lo game (ibid.). According to Hakli, Miller et al., in the I-mode the individual agents can only select strategies, but in the we-mode agents adopt the point of the view of the group and can therefore select outcomes (ibid., 298).¹⁷ We-reasoning is not a type of preference transformation (where individual preferences are no longer self-regarding, but group-regarding), but it constitutes a kind of agency transformation.¹⁸ This idea of agency transformation is central also to Bacharach's work (2006) and echoed in Butler et al. (2011). For the purpose of this article I will assume that people do sometimes reason in this different mode and that work in experimental economics shows this as a matter of fact (Butler, Burbank et al. 2011, Butler 2012).

While the preceding paragraphs served to explain we-reasoning in the context of strategic interaction, let me now return to how I will be using the term in this article. The first obvious departure from that context is that I am interested in strategic *moral* interaction scenarios.¹⁹ The other difference is that I am not only interested in scenarios where agents make their decisions independently. I will come back to this point later.

I use the term we-reasoning to mean the type of moral reasoning where agents in considering their options for acting choose 'collectively available' options over individually available options. *We-reasoning about moral options—as I understand it—requires, first, we-framing the moral problem one is faced with and, second, choosing the appropriate individual action(s) for realizing a collectively available option.* Now, whether or not it makes sense for an agent to we-reason in a given situation will depend on their epistemic position, amongst other things.²⁰

I believe that this approach to understanding joint necessity problems is particularly helpful and in fact more so than alternative approaches. Further, it is helpful for thinking about moral obligations more generally. This latter issue, however, will not be discussed in this paper.

3. WE-REASONING IN COLLECTIVE RESCUE CASES

Returning to our first collective rescue case, *Hikers*, our basic intuition is that both hikers should approach the trapped man and try to lift the tree. But how do agents

make choices in such scenarios? Each hiker in our hypothetical scenario might reason in the following way:

- I have two options: I can contribute to lifting the tree or not. The first option only produces a morally desirable outcome if the other hiker cooperates. The second option produces the less desirable but more secure outcome of a continued hiking experience.
- Therefore, *if* the other hiker cooperates in lifting the tree, then I have an obligation to contribute to lifting the tree.²¹
- Until and unless the antecedent of this conditional is fulfilled, I have no actual obligation to do anything about the trapped man (this is assuming that nothing I could individually do would help him).

This would be an example of moral *I-mode reasoning* about options for moral action and moral obligations. The individual agent reasons about what is best given what others do. However, with regard to the trapped man, such reasoning will yield only conditional obligations. This generates at least two problems: that of deadlock (each of us will help only if the other helps) and the problem of mutual release (both do nothing therewith simultaneously denying the antecedent of each conditional obligation).²² There have been elaborate attempts in the literature to overcome these issues by developing more sophisticated accounts of I-mode obligations for joint necessity cases (Goodin 2012; Collins 2013). The solution offered by authors like Bob Goodin and Stephanie Collins has been to produce more complex conditionals.²³

I do not want to outright reject this picture of how we do (and possibly should) reason in joint necessity cases—and I lack the space here to directly address the individualist approaches described above. But I believe they give only part of the story and telling the other crucial and regularly overlooked part will make understanding these cases much easier.

It seems to me that when faced with joint necessity cases, moral agents regularly we-reason about their options. This is an empirical point and while I know of no empirical study focusing on moral decision making and we-reasoning, the results from studies conducted in behavioural economics and evolutionary biology on nonmoral decision making seem to support this view and I can see no principled reason why these results should not apply to moral deliberation (Bacharach 2006; Butler, Burbank et al. 2011; Butler 2012; Tomasello 2014).

To illustrate this point, let us look at how an agent faced with a strict joint necessity problem like *Hikers* would reason in we-mode.

- There are three options: We can lift the tree together, we can each go our way or one of us can try to lift it on their own. Clearly, the first option is the best from a moral point of view.
- I should play my part in making it the case that we lift the tree.

Instead of reasoning about her individual choices as best responses to the other person's choices (that is, as strategies,) the hiker reasoning in we-mode considers individually and collectively available options. My contention is this: if both agents have reason to engage in this kind of deliberation then they have a collective *pro tanto* obligation.

Note that my enquiry is limited to discussing such *pro tanto* obligations, because in order to make claims about collective *all-out* obligations, one would have to commit to a concrete normative theory. I am interested in the nature of plural obligations as such, rather than any specific substantive version thereof. However, I am confident that the view I put forward is compatible with several first-order normative theories.

Before moving on, let me further defend the plausibility of the proposal I just made. Consider this quick test for how we approach joint necessity cases: Our individual difference-making ability in a scenario such as *Hikers* depends on the other person's willingness and ability to contribute. But do we consider the other's willingness to help as part of the circumstances, that is, a given aspect of the situation, or do we treat their willingness as something that is not-yet-fixed, flexible, and that we could potentially change? I believe it is a mistake to primarily focus on our *individual* difference-making ability in explaining why we consider the option of helping the trapped man. An awareness of our *collective* difference-making ability explains better why we would consider the option to jointly assist in the first place. If an agent takes the collective difference-making ability as her starting point, she will not consider the willingness of the other to be merely part of the circumstances. For instance, most people would think it necessary in situations like *Hiker* to convince the second person to assist if they were unwilling to do so or hesitant, rather than take their unwillingness as a mere given. Most people would also think that we must not wait until the other has indicated their willingness to collaborate but that we each must be proactive by indicating our willingness. This seems to suggest that, regularly, the starting point of our reasoning process in such scenarios is what is collectively achievable. In other words: we-frame and we-reason.²⁴

Let me now illustrate how this process would look for a wide joint-necessity case like *Commuters*. In the case of the commuters pushing a 20+t train aside, there is no strict, but only wide joint necessity: there are more people present than necessary for achieving the desired outcome.²⁵ In other words, if everyone contributes the outcome is overdetermined.

Let us assume that it is known to the individual agent that this is the case. What they do not know is how many people are willing to help tilt the train. In this case, each individual agent when reasoning in pro-group I-mode is fully justified in thinking the following:

- I have two options: I can contribute to pushing the train aside or not.
- If sufficiently many others are willing to (and do) contribute then adding my contribution is not actually helping to produce the outcome. It is superfluous and makes no difference to the outcome.
- If an *insufficient number* of people contribute, then:

- If even when adding my contribution we still fall short of the required minimum number of contributions then my contribution makes no difference or, in other words, it is not the best response to the choices of others. Or
- If my contribution is the one contribution needed to help cross the threshold of minimally necessary contributions then it will make a difference, or, in other words, it is the best response to the others' choices.
- Only in the last case does my action secure the best outcome and is the best response to the others' choices.
- It is not clear whether or not I have an actual obligation to contribute to freeing the trapped man (this is assuming that nothing I could individually do would help him) and it is very difficult for me to find out (since I do not know the exact number of people willing to contribute and the number of those who are necessary). Further, it is unlikely that *I* will be the difference-maker.
- Also, since *my* contribution is not necessary for producing the desired outcome, it would never be *my* failure to contribute which undermines the joint cause. After all, my failure to contribute could always be compensated by any of the other undecided or unwilling agents.²⁶

The problem is, of course, that this kind of I-mode reasoning about her obligations is available to every single commuter on the platform. If everyone deliberates in this way, they may very well fail to rescue the trapped man for no other reason than because they are uncertain about their obligations or for arguing—correctly (!)—that they are not uniquely causing the joint effort to fail.

One might now contend that this is not how we would actually deliberate in a scenario like *Commuters*. I think that many (though probably not all) of us would reason quite differently. The flaws of the approach above are effectively avoided by we-reasoning. Individual commuters might deliberate along the following lines:

- There are three options: A sufficiently large subgroup of us push against the train together and succeeds in tilting it, or we all go our way, or some of us try to tilt the train even though our numbers are insufficient for making a significant difference. Clearly, the first option is the best from a moral point of view.
- I should take steps towards realising the first option (e.g., by encouraging others to push against the train with me).²⁷

I hope this shows how the counterintuitive implications of I-mode reasoning are exacerbated in wide joint-necessity cases. Likewise, I hope that this example lends plausibility to the idea that we could be required to participate in collective endeavours merely because it is the right thing to do, not because we are definitely making a difference to the outcome.

Let me now unpack my proposal—that agents have collective obligations precisely when they should employ we-reasoning—a little more. What does it mean to say that agents should engage in this kind of deliberation? I will answer this question using Christopher Woodard's notion of *group-based reasons* (2003; 2011). To say

that agents should we-reason is to say that they should give preference (or greater weight) to group-based reasons over individual-based reasons.

In the *Commuters* case, when the individual commuter notices that a man is trapped between the train and the platform, she can quickly conclude that help can only be provided collectively. But do others appear willing and ready to contribute? In many wide joint-necessity cases we have insufficient information on how many others are willing and ready to contribute. Do I have reason to contribute even if there is no overwhelming evidence of others contributing? Woodard defends the following claim:

[A] pattern of action by the group is capable of providing a reason to perform a constituent part of the pattern, so long as the group could perform this pattern of action—where that means that each member *could* play her part. (2003, 225)

A group-based reason to participate “would be provided by the consequences of the group action, of which the participation is only a part” (2003, 216). He explains that

[T]he idea is not that there is a reason to perform *A* because it will bring about *P*, or make *P* more likely. Instead it is that the goodness or rightness of *P* provides reasons to perform its parts, just in virtue of their being parts of it. (2011, 263)

Woodard rejects the idea that the others’ willingness as such is decisive for the existence of a group-based (or, as he later calls it, *pattern-based*) reason. Rather, such willingness will impact on the strength of one’s group-based reason (2003, 2011). According to Woodard:

[A]n individual can have a group-based reason to participate in a group action even when no other member of the group is willing to cooperate. This makes it possible for group-based reasons to conflict properly with individual-based reasons, and so increases their interest. (2003, 216)

The conflict between individual-based reasons and group-based reasons Woodard describes is precisely the conflict between I-reasoning and we-reasoning that is expressed in the different perspectives that the hikers and the commuters in our examples can take, as portrayed above. In other words, when we-reasoning about moral obligations, a deliberating agent will put emphasis on pattern-based or group-based reasons for acting in a particular way. When I-reasoning about moral obligations, a deliberating agent will emphasize individual-based reasons for acting in a particular way.

This means that I may have reason to we-frame and we-reason even if I do not know anything about the other persons’ actions or decisions. Think, for instance, of a modified *Hikers* case: as part of an experiment, two people find themselves in separate rooms without being able to communicate to each other, each facing a red button. They are each being told that there is another person in another room, who is

being given the same information: if both of them press the red button then a person who is trapped elsewhere under a log will be freed. If they both fail to press the button then the person will remain trapped and will most likely die. I hope that most people would think it is obligatory to press the red button, despite the fact that we do not know what the other one will do. Let me spell this out for the sake of clarity: each of the two agents has reason to consider the collectively available option as best (each pressing the button), they each have reason to include it in the set of options over which they deliberate, and they have reason to take corresponding steps towards realizing that option, namely pressing the button in front of them.

Often we will not need to have positive confirmation of what another person is doing and of whether they are contributing in order to justifiably we-reason about the problem at hand. We often (are required to) we-reason about our obligations *unless* we have reason to believe that the other(s) will *defect* (Lawford-Smith 2012; Aas 2015).²⁸ At other times, I believe, we may be required to we-reason *only if* we have reason to believe that others will *cooperate*.

There is no simple criterion for deciding whether to privilege we-reasoning over I-reasoning. This may seem unsatisfyingly vague, but I believe that any attempt to once and for all settle this matter in the abstract is bound to fail. It seems to me that much will depend on context and, furthermore, it will often be a matter of discretion.²⁹ However, where joint actions or jointly produced outcomes are part of our regular behavioural repertoire, agents should we-reason per default, e.g., when setting the table for dinner or helping to lift a pram into a bus.

Five things should be noted at this point. First, more needs to be said on the crucial question of when we use and should use each type of reasoning and when we are justified or required to abandon we-reasoning for I-reasoning. I will do so in the next section. Second, as I said earlier, the claims I have defended so far are not tied to any particular normative moral theory. In particular, when I speak of morally desirable or best options this need not be read in a consequentialist way. It could be that the best option is the one which entails respecting someone's autonomy.³⁰

Third, and relatedly, this enquiry is limited to exploring *pro tanto* moral obligations, rather than *all-out* moral obligations. We may say that the individual commuter has a *pro tanto* obligation to contribute to tilting the train even where she is uncertain whether sufficiently many others will do so or even where she knows that they will not. Whether or not it becomes an *all-out* obligation, depends on her knowledge of others' willingness to contribute, but also on competing normative claims and obligations. Further, it depends on which normative ethical theory is correct. Fourth, there might be a worry that the argument here leads to a proliferation of collective obligations and makes us positively responsible for an endless number of desirable outcomes that we could in principle collectively achieve. I will address this problem in the last section.

Fifth, to clarify: we-reasoning is reasoning at the individual level. This will be the only kind of reasoning possible in many scenarios where collectively available options are best, but the group of agents is large or dispersed, with very limited opportunity for potential helpers to communicate, discuss options, and choose joint strategies. However, there are also cases where in order to address a joint-necessity

problem agents need to directly communicate, develop a joint strategy, and distribute individual tasks, sometimes repeatedly and over a sustained period of time. This may often take the form of group-based reasoning (Sugden and Gold 2007; Hakli, Miller et al. 2010). However, in this paper, I am mainly interested in those cases where agents have to make their choices independently or with very little information on how the others choose or where only minimal group-based reasoning is required.

4. WHEN TO WE-REASON ABOUT OUR MORAL OBLIGATIONS

Presumably, the reader will grant the empirical point that in scenarios where the contributions of two or more agents are jointly necessary in order to produce a morally desirable outcome, we sometimes reason from the top down, so to speak, starting with the most desirable option even though that option is not available to us in the way that individually achievable options are available to us. Potentially, the reader will even grant the normative point, that we-reasoning might *sometimes* be a requirement of rationality or morality or both. But the obvious question is *when* should we engage in which kind of reasoning? While I said above that it is difficult to give a simple answer to this question, I will nonetheless try to say more about the circumstances under which we ought to we-reason. The answer to this question is crucial for my argument, because I claim that deliberating in I-mode often yields different results concerning our moral obligations than deliberating in we-mode.

At this point, a more fundamental problem needs to be discussed, namely whether or not the process of moral deliberation is or is not independent of what our actual moral obligations *are*. In other words, whether or not in deliberating about our moral obligations we are merely discovering (or failing to discover) what moral obligations we (objectively) have. I am talking about the difference between objective, subjective, and prospective views of moral obligation (Zimmerman 1996). According to the first, what we ought to do is ultimately independent of what we (can) know and even if we act on the best information available and are meticulously scrupulous in our decision making, we may still be violating our actual obligations. Moral deliberation serves to get us as close as possible to finding out what our—objective—obligations are. According to both the subjective and prospective view, our obligations depend on what we know or believe to be true (subjective view) or what we should know or believe to be true if we are conscientious and avail ourselves of the evidence (prospective view) (*ibid.*).

It seems to me that if we are objectivists about moral obligations, the idea of collective moral obligations as obligations that people hold together or jointly reduces to the question of what is—as a matter of fact—morally best or optimal (however moral optimality is understood, that is, according to the normative ethical theory one endorses). We-reasoning can then be seen as the most adequate deliberation method, since it requires each agent to consider collectively available options. On this view, we-reasoning can help us discover what is in fact morally best, where I-reasoning can make us fail to see that.

However, this is not the route I wish to take here. Rather, I propose to turn this argument on its head: Instead of assuming that our obligations (collective or

individual) are imperatives to produce morally optimal outcomes (or perform morally optimal actions for that matter), I suggest that what agents are morally obligated to do is to produce what they *reasonably believe* to be the best available outcome, given they have *conscientiously investigated* the options.

This picture disagrees with the objective view of moral obligations and I am aware that this move requires justification. Since I am not able to discuss the (de)merits of the objective view here, I will point to Michael Zimmerman's critique and respective rejection thereof.³¹ Zimmerman's main objection to that view is that it yields counterintuitive results in some scenarios with incomplete knowledge (2014, 30ff).³²

Further, if the *subjective* view of moral obligation were correct then agents' collective obligations would depend on what each individual agent believes to be morally best in a given set of circumstances, regardless of whether she has good evidence for that belief. Zimmerman rejects this view, mainly because it would render conscientious enquiries into the correctness of one's beliefs meaningless (2014, 26ff).

In contrast, on the *prospective* view, which he proposes and develops, what matters is what evidence there is and whether this evidence supports the belief regarding one's obligation (2014, 64). According to Zimmerman, this is "the best bet regarding the actual values at stake" (2014, 34).³³

My argument here aligns with a prospective view of moral obligations. That is, it makes sense only if we think moral obligations depend on what the individual deliberating agent can know, provided they are conscientious. If that is how we approach the issue of moral obligations then, I believe, some of the core propositions of existing accounts of collective obligations are more easily explained. On the subjective view, e.g., it is harder to justify the demand that agents should try to find out if the other is willing to cooperate, that is, to demand that agents collect evidence. On an objective view, it makes less sense to argue that one's beliefs concerning the other's willingness to cooperate impact on one's duties (as in Lawford-Smith [2013]). I am not suggesting that existing accounts of collective obligations have in fact endorsed the prospective view. I suspect that many of them implicitly rely on an objective view of moral obligation (Wringe 2005; Collins 2013; Pinkert 2014; Schwenkenbecher 2014) or else oscillate between different views (Lawford-Smith 2012; Lawford-Smith 2015). But I believe that the prospective view can best make sense of—and in fact simplify—such accounts or some of their core features, at least.

With Zimmerman, I believe that agents should avail themselves of evidence as part of discharging their moral duties. Taking his view one step further, I also hold that this means to recognize that others require evidence for their own moral deliberation, that one is often in a position to provide relevant evidence, and that one should in fact provide it, especially when faced with joint necessity cases.³⁴

This takes me back to our main issue and the one the reader has been waiting to hear about: When should agents we-reason about their obligations? When are we justified in abandoning we-reasoning for I-reasoning (or even be required to do so)? Remember that I am claiming that agents have obligations together (or jointly) with others precisely when they should we-reason about their options. Tying this back to what I said earlier: two (or more) agents have collective moral obligations precisely when they should act on group- or pattern-based reasons.

As I said before, I will leave aside the question of what substantial, ethical theory is correct. For the sake of argument, let us assume that people agree on what constitutes moral optimality.³⁵ If the prospective view is roughly correct, then agents should we-reason when the collective option is their best (moral) bet given the evidence. If both agents are equally in a position to view the same problem in the same way, and if to the conscientious agent pattern-based reasons outweigh individual-based reasons then both become the joint subject of a collective obligation, that is, they jointly hold that obligation.

In *Hikers*, the two potential rescuers have a collective obligation if each has conclusive reason to privilege the joint activity over individual activities. This implies that the joint activity is in fact available, which is easily established: each individual hiker can avail herself of evidence concerning the other hiker's ability and willingness to contribute. If they both have this evidence instantly, for example, where both signal their willingness to contribute instantly, then the collective obligation to assist the trapped man is instantaneous.

What if one of the hikers is not willing to contribute or does not indicate their willingness to the other? Her unwillingness may stem from a lack of evidence, e.g., evidence for the optimality of the collectively available option. In that case she may not have to we-reason and there is no collective obligation. Another possibility is that the evidence is there, but she fails to avail herself of it—she fails to take it in (perhaps in a negligent way) and as a result fails to act (perhaps she simply walks away from the scene). In this case we could possibly say that the hiker should have been aware of the optimality of the collective option and that there was a collective obligation on her and her counterpart even if *she* did not realize it. Her epistemic failure and lack of due care led to a moral failure.

Importantly, even in the first case where the oblivious hiker is not to blame for her ignorance of the optimality of the collective option, the other hiker can generate a collective obligation by providing her with the evidence and indicating her willingness to contribute. We often can (and regularly do) generate joint obligations with others in this way: “Hey, you, could you come over here and help me with this?” is all that it sometimes takes. If one of the two hikers refuses, despite evidence that a collective option is available and best, then she may be individually blameworthy for the joint failure to rescue the trapped man, since it is a strict joint necessity case (see, e.g., [Schwenkenbecher 2014]).

What if they cannot agree on how to go about carrying out the joint activity—a scenario discussed by Virginia Held (1970)? Held describes this as a case where the collective remains responsible for the failure to assist, but the individuals are absolved from responsibility. On my view, such absolution would depend on the evidence they had and the complexity of the task. However, these issues are difficult if not impossible to determine in purely the abstract.³⁶ Sometimes, a collective task is too complex to be accomplishable by a random congregation of agents. In such cases, there may be no collective obligation on the agents, because they lack joint ability. Or they may have an obligation to transform themselves into the kind of collective that can successfully address the task at hand (May 1992; Collins 2013). In other

cases, the task may be perfectly doable, but the chance of a coordination failure still high enough to excuse the agents for failing to perform it.

Summing up what has been said so far, we arrive at the following sufficient conditions for collective obligations:³⁷

Two or more moral agents jointly hold an obligation to perform an action or produce an outcome corresponding to a collectively available option under the assumption that each agent, if she is conscientious,

- i. has reason to believe that the collectively available option (e.g. joint rescue) is morally best;
- ii. has reason to include that option in her deliberation about her obligations (we-framing the problem);
- iii. has reason to deduce her individual course of action based on (i) and (ii) (we-reasoning about the problem) and the ability to do so; and the agent
- iv. has no overriding obligations, is not unduly burdened by the task, and is jointly capable with the other(s)³⁸ of discharging the task at hand.

Steps (i–iii) may take place consecutively or simultaneously. Further, (ii) and (iii) will usually (but not necessarily) involve some communication between the potential collaborators. This can be nonverbal communication. For instance, if I see the other person lay hands on that fallen tree in an attempt to remove it then I usually have a reason as per (ii) and (iii). Likewise, if I see lots of other people push against the train to tilt it, then I have reason as per (ii) and (iii).

We can now see that our two exemplary cases had a few implicit assumptions on which the initial plausibility of the claim that the passers-by and bystanders should jointly assist rests. Namely that rescuing the person is indeed best overall³⁹ and that this is obvious to each person.⁴⁰ The main question that had scholars worried was how to conceptually account for two issues: (a) that at time t_1 , when each agent is first confronted with the moral problem, she does not know if there are enough others willing to assist, in other words, whether there is joint ability, and (b) that the success of securing the best option does not depend on her alone, that is, that the potential duty of rescue cannot be neatly allocated to a specific agent.

My account has mitigated each problem: It proposes that two (or more) agents hold duties collectively when they each have reason to consider the collective option best, to include it in their set of options, and to actually take steps towards acting on that option. They will regularly have reason to do so when the collective activity is a contextually given default position or where they have some positive indication of willingness to cooperate from the other agent(s). The latter case includes those cases where one person takes the lead and distributes tasks.

If they fail to we-frame the situation or to we-reason about their individual actions, they make a mistake in moral deliberation (which may or may not be excusable). Flawed reasoning can sometimes still yield the right result. But if agents fail to collaborate as a result of flawed reasoning they violate a moral obligation. This is no different for individual duties, except for one aspect: if we accept that agents can

hold moral obligations jointly with others it looks like *I* may be on the hook for *your* flawed reasoning. Let me briefly comment on this issue. If we accept that obligations are held jointly in some joint-necessity cases then it seems that automatically each agent who is part of the obligation has some reason for ensuring the obligation is discharged and for correcting others' flaws in moral reasoning. This to me seems to be a good thing. If *I* recognize that *we* ought to jointly save the trapped man by lifting the tree and *you* fail to see that (without good reason) then I should make an effort to convince you.

Further, let us suppose that I fail to get you to acknowledge what you should acknowledge, namely that we jointly ought to help the trapped man. We might say that the resulting failure to act on our duty to help him is a joint failure. That is, the wrongdoing involved in failing to rescue the man is joint wrongdoing. But we need not commit to making a similar claim about blameworthiness. It seems perfectly intelligible to say that we jointly failed the trapped man and *you* are to blame. Admittedly, this question needs further scrutiny. Unfortunately, I do not have the space to discuss this here, but will instead point to my earlier discussion of this issue (2014).

5. PROBLEMS AND OPEN QUESTIONS

Let me finally address five further qualifications of and potential objections to the above-sketched account of collective obligations.

First, it might be objected that sometimes we (need to) make moral decisions very quickly and cannot go through the steps indicated above. Often, so the objection goes, we have no clear idea about exactly how many people are necessary in order to complete some collective task, how many of those around are willing to help, and whether the task is doable at all, but we nonetheless get to work and make a start. I agree that this is often how we approach collective tasks, but this does not undermine my argument. The above steps will often take place simultaneously or not be explicitly followed at all. They constitute, of course, an idealized account of moral reasoning.

Also note that if we make decisions very fast we often get things wrong. Picking the collective option (and we-reasoning about the problem) may sometimes be our first impulse, but it may not be the best bet. Depending on the circumstances we may or may not be excused for getting it wrong. Further, our social and cultural conditioning, of course, impacts on our reaction to joint-necessity cases and other collective moral action problems. For some, picking a collective option (and acting on it) may be an almost instinctive response, whereas for others in the same situation it is not. We have different aptitudes for recognizing collective moral action problems as problems for 'our group'. To what extent these different aptitudes and the impact of social factors should be reflected in our ethical theories is a topic for another paper.

Second, one might want to insist that regardless of whether it turns out that we are actually able to help the trapped person, each of us is under an obligation to approach the person and to make an effort to lift the tree even if the other potential helper is doing nothing. According to this objection, we should first and foremost try to help, even if it is pointless. There seems to be something morally wrong about not

even trying even if the other person clearly does not want to collaborate. And if we grant this point then does this not suggest that first and foremost we have an—individual—obligation to *try* to assist?

I believe that the idea that we should try to assist regardless can be explained in a way that does not undermine my argument. In my view, the appeal of this idea rests on the intuition that we should try to find out if we have individual ability to lift the tree. It would be odd indeed, given what is at stake, to not even investigate that possibility. We ought to explore our options, where these are not obvious at first glance, especially in cases of great urgency with a high potential for significant moral loss. In other words, rather than suggesting that really we only have individual obligations to *try* to assist in a joint necessity case like *hikers*, the intuition that we ought to do *something* even if the other does not cooperate points to an epistemic duty to investigate, which we may have even where we are objectively powerless to assist.

Third, it is important to note that there will be more complex and less urgent collective tasks, which require more detailed elaboration of individual action choices than individual-based we-reasoning can achieve. These cases are not covered in the account presented here. Where group membership is relatively stable and communication possible, potential collaborators may (need to) discuss different options together in a way that is best described as team-reasoning (Sugden and Gold 2007; Hakli, Miller et al. 2010). This process happens at the level of the group and requires certain epistemic conditions to be in place.⁴¹

Fourth, one might nonetheless wonder what the advantage of my view is over individualist views, when, in fact, it might appear more complicated than those. However, it is important to distinguish between what may be called ‘basic’ individualist views and more sophisticated accounts such as Stephanie Collins’s for instance (2013). The view proposed here is in fact simpler than Collins’s.⁴² Further, it is an advantage of a collectivist approach that it can explain and justify the emergence of individual contributory obligations in joint necessity cases in a way that individualist accounts cannot (Wringe 2016). Further, it can deliver unconditional obligations and it avoids many of the counterintuitive implications of the individualist view, as discussed in section three.

Fifth, there might be a worry that the argument here opens the floodgates to a multitude of collective obligations for pretty much anything that we could in principle achieve collectively—which is a lot. In other words: does this approach not lead to a morality that is overly demanding? But note that this worry is misguided, since I only discussed *pro tanto* collective obligations. Whether or not any specific set of individual agents have such joint obligations, does not merely depend on their ability to effect positive change in the world, but on what their evidence gives them reason to act on. By contrast, one might believe that the view proposed here makes the notion of collective obligations impotent. This objection, however, fails to acknowledge that most of us are already acting on collective obligations, that is, for group-based moral reasons. Many of us who reduce our carbon footprint, who reduce their meat consumption, or who simply vote in a general election think of our actions in terms of contributing to or playing our part in a collective obligation. After all, when we

justify acting in ways that make no discernible difference to a desired outcome we tend to invoke pattern-based reasons for our choices.⁴³

NOTES

1. Agents will count as organised or nonorganised always with regard to a specific goal. That is, a group may be organised with a view to producing x , but may be unorganised with a view to producing y . My colleagues and I are organised as a university department for the purposes of teaching and administration, but we are not organised for the purpose of holding a political rally (as a political activist group would be) or building a house (as a building firm would be). Even though we could perhaps organise to do those things, they are not activities embedded in the current structure of the organisation. Hence, with regard to that (alien) activity, we count as nonorganised.
2. This means that certain issues are outside the scope of this enquiry. For instance: when should we form an organised group in order to resolve a morally urgent problem? However, that the issue of forming an organised group raises the exact question discussed here: when do individuals who are not (yet) organised in a group have an obligation to do so (and how do they know)?
3. I am adopting the term *joint necessity* from Holly Lawford-Smith (2012), adding the distinction between *strict* and *wide*, as well as that between *analytic* and *circumstantial* joint necessity.
4. Variations of this scenario abound in the literature on collective agency and responsibility (Held 1970; Cullity 2004; Lawford-Smith 2012; Collins 2013; Schwenkenbecher 2013; Schwenkenbecher 2014; Wringer 2016), to name only a few. In all of these scenarios, a person's life is threatened and it takes at least two people to avert the threat.
5. In both cases a lot is at stake, the need for action is urgent, and it is within our capacity to help. Even though the argument that if we can produce something morally valuable at little cost to ourselves then we are morally obligated to do so has been most prominently made by utilitarians (Singer 1972), all major moral theories would agree with the gist of this conclusion.
6. More precisely, they disagree on how we should conceptualize the idea that two agents together ought to produce a morally optimal outcome. 'Revisionists' claim that collections of two or agents can hold obligations as a group or together even if they do not form a group agent in the strict sense. 'Conservatives' argue that individuals (only) have individual obligations to promote morally optimal outcomes in such cases.
7. This happened in Perth, Australia, on 5 August 2014. See <http://www.abc.net.au/news/2014-08-06/man-freed-after-leg-trapped-in-gap-on-perth-train-station/5652486>. A similar event occurred in Japan the year before: <https://www.theguardian.com/world/2013/jul/22/tokyo-train-passengers-rescue-woman-trapped>. See also: <https://www.youtube.com/watch?v=39i89NJCRCQ>. Both accessed on 22 Feb 2017.
8. For a short explanation of the origins of the concept of team reasoning see Robert Sugden (2015, 152).
9. I am using this term differently from how, e.g., Butler et al. use it (2011, 2012).
10. Let me add a quick comment on using the decision-theoretic tools of game theory in ethics. Our moral decision making can concern problems that are structurally similar to the scenarios studied in game theory, including coordination problems or prisoners' dilemmas. Therefore, I believe it is fruitful to adopt the concept for our purposes with one (important) caveat: traditionally, the study of strategic interaction focuses on scenarios where individual agents make their choice(s) independently of one another. In real life, and in the moral decisions it imposes on us, this is often not the case.
11. "A Nash equilibrium is a set of strategies such that no player can improve her payoff by changing her strategy, given the strategies of the other players." (Reiss 2013, 58). In other words, if player 1 chooses 'A' (or 'B') then player 2 cannot improve her payoff by choosing anything other than 'A' (or 'B'). She is best off choosing the same strategy given the other's strategy.
12. (A/A) is optimal in that there is no solution that leaves one of the players better off while leaving the other one at least not worse off. (B/B) is not optimal, because there is a solution that would make both players better off.
13. This is usually an implicit, not an explicit assumption. In prisoners' dilemma narratives, e.g., both suspects are being told the same story, so the payoffs are at least shared knowledge between. Further, it is usually assumed that each prisoner knows that the other knows these payoffs and vice versa.

14. The concept of we-reasoning is meant to improve classic game theory in at least three ways: (1) It redefines what a *rational* choice is for individual players, allowing the cooperative choice in one-shot prisoners' dilemmas to be rational, (2) it allows us to better *predict* players' choices, (3) it *explains* why players make those choices.
15. These are merely the two most obvious alternatives. There are other possible explanations, e.g., that players might be pursuing a *maximax* strategy.
16. Confusingly, Hakli, Miller et al. (2010) use the term 'we-reasoning' for both pro-group I-mode reasoning and we-mode reasoning.
17. Hakli, Miller et al. say that "a group agent can in a sense select outcomes" (2010, 298). I think this is misleading, as we-mode reasoning is employed by individuals. Framing the situation as one for the group and choosing individual actions accordingly will not require a group agent. Nor does acting together necessarily transform a group of individuals into a group agent (Pettit and Schweikard 2006).
18. According to some authors, we-mode reasoning involves forming a collective intention (Hakli et al. 2010), that is, an intention with regard to the action (to be) performed by the group as different to the intention concerned with the action (to be) performed by the individual. Gold and Sugden (2007) concur that the modes of reasoning behind the 'we-intentions' people form when cooperating differ from those behind 'I-intentions'. Collective intentions, in their view "are the product of a distinctive mode of practical reasoning, team reasoning, in which agency is attributed to groups" (2007, 137). However, we-mode reasoning does not necessarily involve the forming of collective intentions, at least not if such intentions are thought to require interdependent beliefs about others' intentions. In cases like the 'hi-lo' game and the PD it is neither necessary nor likely that agents have no interdependent beliefs about others' intentions.
19. It is, of course, perfectly compatible with game theoretic analysis that the players' preferences are of moral nature. In that case, the highest payoffs correspond to the morally best outcome.
20. Numerous factors may impact on how an agent frames a given scenario, including their level of identification with the group, their previous experiences with spontaneous collaboration, or contextual cues as to the social expectations of the given situation.
21. Assuming that I have an obligation to produce morally desirable outcomes.
22. Another problem is this: If they fail to free the trapped man and continue on their hike, the wrongness of failing to save him would not seem to go over and above the wrongness of failing to contribute to saving him. If there is no obligation to save him (as there is no appropriate subject of such an obligation) then no one can be blamed for having failed to save him. This seems counterintuitive (see Schwenkenbecher 2014). Further, and perhaps more controversially, the I-mode approach seems to challenge our moral phenomenology (see Wringer 2016).
23. While I am not able to argue for this in any detail, I should point out that I am not convinced that Bob Goodin solves the problem. He rightly shows that simple conditional commitments are not enough even if both are committed to the maximally ethical action. "I will lift the tree if you will" does not commit them to acting unilaterally and satisfying the conditional (Goodin 2012). Goodin's solution is to argue that individuals must commit to the following conjunction of two conditionals: *I will if you will* and *I will if (you will if I will)*. This solution does not escape the problem of failing to match our intuitions. If you and your friend fail to lift the tree and free the trapped man you have each failed in a duty to contribute to that action, not in a duty to perform that action. Further, how do I as a deliberating agent know that you will contribute if I will contribute? We may not be able to find out (e.g., if we cannot directly communicate) in which case, on Goodin's account, we either would not have an obligation to contribute or not know whether we do. I believe that my alternative *Hikers* scenario shows, contra Goodin, that we may have an obligation regardless.
24. It may be tempting to ascribe collective obligations and obligations to we-reason whenever a collective outcome is optimal, but I approach the problem via the perspective of the agents who find themselves in such situations. The view defended, therefore, commits me to a nonobjective view of moral obligations (see Zimmerman 2014).
25. I thank Matthew Kopec for pointing out the importance of this case for my argument.
26. This marks the difference between wide joint necessity cases, like *Commuters*, and strict joint-necessity cases like *Hikers*. For the former, each agent is correct in arguing that *her* refusal does not uniquely undermine the collective endeavour.

27. This reasoning reflects the case of the commuters where there is no certainty as to exactly how many people are required to help and decisions have to be made fast. This will often be the case in rescue scenarios, but not in other collective assistance contexts. Where there is the opportunity to discuss decisions yet other, more sophisticated, forms of collective reasoning may occur instead.
28. In contrast, Derek Parfit seems to have suggested that we only have reason to act as a member of a potential group if we know that enough others are contributing (Parfit 1984).
29. This means that sometimes we may not be able to establish whether a set of agents really did have a collective moral obligation to produce some outcome or action.
30. For a discussion of this issue, see Zimmerman (1996).
31. In Zimmerman's own words: "... if one should ever be so fortunate as to know which of one's options actually is best, then the verdict issued by the Prospective View is the same as that issued by the Objective View... It is only when ignorance regarding the relevant empirical or evaluative facts enters the scene that a gap may emerge between what is actually best and what is prospectively best" (2014, 92).
32. It should be noted that Zimmerman rejects the idea of collective obligations (1996), but there is no reason why those who do not share his scepticism cannot endorse his analysis of moral obligations.
33. Zimmerman uses a case by Frank Jackson, which has a doctor provide a partial cure to a patient even though a complete cure is available and the doctor is aware of that. In this case, there is not enough evidence for the doctor to distinguish between the drug causing a complete cure and one causing certain death (Jackson 1991). However, there is clear evidence for which drug provides the partial cure, and the doctor—rightfully in Zimmerman's view—gives the patient the latter despite knowing that it is not the best option (Zimmerman 2014, 30). Zimmerman points out that on the objective view she has acted wrongly.
34. This explains the obligation to signal one's willingness to cooperate to other agents.
35. One might object that since people in fact disagree on substantive moral issues my account of collective obligations never gets off the ground, since it requires people to agree on what matters (most) morally. But I disagree: we can perfectly well speak about individual moral obligations in an abstract manner without adopting a specific substantive view, so this move should be available for discussing collective obligations, too.
36. One recent suggestion to determine this in the abstract is Pinkert's (2014) proposal that scenarios triggering collective obligations are those that have one salient solution, where it is clear what needs to be done collectively and what each individual needs to do (immediate joint ability) or where this can easily be established (mediate joint ability). His proposal shares features with mine, but is, I believe, (implicitly) committed to an objective view of moral obligations.
37. These conditions are jointly sufficient for collective obligations. Plural obligations can also arise without joint necessity, that is, in cases where one person alone can produce a desirable outcome, but as a matter of fairness the burden of producing it should be distributed amongst several people. These types of cases, however, are not discussed here.
38. What exactly it means for two or more agents to be jointly capable of doing something (performing an action or producing an outcome) is another issue deserving more detailed discussion, which I provided elsewhere (Schwenkenbecher forthcoming).
39. For the cases discussed, this assumption was made more plausible by the fact that all major moral theories and common sense morality would converge on the judgment that rescuing someone's life at little cost to oneself is morally obligatory.
40. Other features that support the plausibility of an instant duty to collectively assist are epistemic simplicity (it is obvious what the solution to the problem is and what the contributory roles involve, it is clear that the problem needs a quick solution, that others are willing to help and agents can communicate effortlessly) and moral simplicity (it is a one-off effort, posing a negligible burden on helpers, and there are no competing duties).
41. Sometimes it is assumed that team reasoning requires common knowledge (Gold and Sugden 2007). However, I believe that some lower-level group knowledge may suffice.
42. On Collins's account, the passers-by in our collective rescue cases would have duties to collectivise, that is, duties to work towards forming a group agent, which then acquires a duty to perform the rescue action. The duty of the group agent entails contributory duties for group members. Her five criteria for collectivization duties (CCD) are as follows:

(1) φ is morally pressing, and (2) at t_1 , either: no (collective or individual) agent/s have duties, either to φ ? or to take responsive actions with a view to there being the morally desirable outcome that φ produces; or too many agent/s with such duties default, and (3) if, at t_1 , A, \dots, N each took responsive steps towards there being a collective-that-can- φ then, at t_2 , that collective would incur a duty to φ , and (4) at t_1 , A, \dots, N are each able to take responsive steps towards there being a collective-that-can- φ at a reasonable expected personal and moral cost, and (5) other individuals will not successfully take responsive steps towards there being a collective that will incur a stronger duty to φ .

If these are met then

(6) at t_1 , A, \dots, N each has a duty to take responsive steps towards there being a collective-that-can- φ , and (7) at t_2 , once a $\{A, \dots, N\}$ collective-that-can- φ is formed, that collective has a duty to φ , which entails (8) at t_3 , once the collective has distributed φ -related roles, each member with a φ -related role has a duty to perform that role. (Collins 2013, 244).

See also my discussion of Collins's account in Schwenkenbecher (2018).

43. The author gratefully acknowledges financial support from the School of Arts at Murdoch University (2016, 2017), and from an Oxford Martin Visiting Fellowship by the Oxford Martin School (2017–2018) to conduct research for this article. Further thanks go to Matthew Kopec, Olivier Roy, David Butler, and Christopher Woodard for very helpful comments. This article has also benefitted from discussions with colleagues at conferences (Australasian Association of Philosophy's Annual Conference, Manchester Political Theory Workshops) and research seminars (Charles Sturt University, Oxford University, The University of Western Australia), as well as from a writing retreat run by the Centre for Responsible Citizenship and Sustainability at Murdoch University (with funding from the School of Business and Governance).

REFERENCES

- Aas, S. 2015. "Distributing Collective Obligation," *Journal of Ethics and Social Philosophy* 9(3).
- Bacharach, M. 1999. "Interactive Team Reasoning: A Contribution to the Theory of Co-operation," *Research in Economics* 53(2): 117–47.
- . 2006. *Beyond Individual Choice: Teams and Frames in Game Theory*, Princeton: Princeton University Press.
- Björnsson, G. 2014. "Essentially Shared Obligations," *Midwest Studies in Philosophy* 38(1): 103–20.
- . forthcoming. "On Individual and Shared Obligations," *Philosophy and Climate Change*.
- Butler, D. 2012. "A Choice for 'Me' or for 'Us'? Using We-Reasoning to Predict Cooperation and Coordination in Games," *Theory and Decision* 73(1): 53–76.
- Butler, D.J., V. K. Burbank, and J. S. Chisholm 2011. "The Frames behind the Games: Player's Perceptions of Prisoners Dilemma, Chicken, Dictator, and Ultimatum Games," *The Journal of Socio-Economics* 40(2): 103–14.
- Collins, S. 2013. "Collectives' Duties and Collectivisation Duties," *Australasian Journal of Philosophy* 91(2): 231–48.
- Gold, N. and R. Sugden 2007. "Collective Intentions and Team Agency," *Journal of Philosophy* 104(3): 109–37.
- Goodin, R.E. 2012. "Excused by the Unwillingness of Others?" *Analysis* 72(1): 18–24.
- Hakli, R., K. Miller, and R. Tuomela 2010. "Two Kinds of We-Reasoning," *Economics and Philosophy* 26(3): 291–320.
- Held, V. 1970. "Can a Random Collection of Individuals Be Morally Responsible?" *Journal of Philosophy* 67(14): 471–81.
- Isaacs, T.L. 2011. *Moral Responsibility in Collective Contexts*, Oxford: Oxford University Press.

- Jackson, F. 1991. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," *Ethics* 101(3): 461–82.
- Lawford-Smith, H. 2012. "The Feasibility of Collectives' Actions," *Australasian Journal of Philosophy* 90(3): 453–67.
- . 2015. "What 'We'?" *Journal of Social Ontology* 1(2): 225–49.
- List, C. and P. Pettit 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*, Oxford and New York: Oxford University Press.
- Ludwig, K. 2016. *From Individual to Plural Agency: Collective Action I*, Oxford: Oxford University Press.
- May, L. 1992. *Sharing Responsibility*, Chicago: University of Chicago Press.
- Parfit, D. 1984. "Five Mistakes in Moral Mathematics," in *Reasons and Persons*, Oxford: Clarendon Press, 1: 55–83.
- Pettit, P. and D. Schweikard 2006. "Joint Actions and Group Agents," *Philosophy of the Social Sciences* 36(1): 18–39.
- Pinkert, F. 2014. "What We Together Can (Be Required to) Do," *Midwest Studies In Philosophy* 38(1): 187–202.
- Schwenkenbecher, A. 2013. "Joint Duties and Global Moral Obligations," in *Ratio* 26: 310–328.
- . 2014. "Joint Moral Duties," *Midwest Studies In Philosophy* 38(1): 58–74.
- . 2018. "Making Sense of Collective Moral Obligations: A Comparison of Current Approaches," in *Collectivity: Ontology, Ethics and Social Justice*, ed. T. Isaacs, K. Hess and V. Ignjeski, Lanham, MD: Rowman and Littlefield, 109–32.
- . *Collective Moral Obligations*, unpublished manuscript.
- Singer, P. 1972. "Famine, Affluence, and Morality," *Philosophy and Public Affairs* 1(3): 229–43.
- Sugden, R. 2015. "Team Reasoning and Intentional Cooperation for Mutual Benefit," *Journal of Social Ontology* 1(1): 143–66.
- Sugden, R. and N. Gold 2007. "Theories of Team Agency," in *Rationality and Commitment*, ed. F. Peter and H.B. Schmid, Oxford: Oxford University Press.
- Tollefsen, D. 2015. *Groups as Agents*, Cambridge, UK: Polity Press.
- Tomasello, M. 2014. *A Natural History of Human Thinking*, Cambridge, MA and London, England: Harvard University Press.
- Tuomela, R. 2013. *Social Ontology: Collective Intentionality and Group Agents*, New York: Oxford University Press.
- Woodard, C. 2003. "Group-Based Reasons for Action," *Ethical Theory and Moral Practice* 6(2): 215–29.
- . 2011. "Rationality and the Unit of Action," *Review of Philosophy and Psychology* 2(2): 261–77.
- Wringe, B. 2005. "Needs, Rights, and Collective Obligations," *Royal Institute of Philosophy Supplement* 80(57): 187–207.
- . 2010. "Global Obligations and the Agency Objection," *Ratio* 23(2): 217–31.
- Zimmerman, M.J. 1996. *The Concept of Moral Obligation*, Cambridge: Cambridge University Press.
- . 2014. *Ignorance and Moral Obligation*, Oxford: Oxford University Press.