


ARTICLE

Collective Reasons and Agent-Relativity

Alexander Dietz 

Cardiff University, Cardiff, Wales

Corresponding author. E-mail: DietzA4@cardiff.ac.uk

Abstract

Could it be true that even though we as a group ought to do something, you as an individual ought not to do your part? And under what conditions, in particular, could this happen? In this article, I discuss how a certain kind of case, introduced by David Copp, illustrates the possibility that you ought not to do your part even when you would be playing a crucial causal role in the group action. This is because you may have special agent-relative reasons against participating that are not shared by the group as a whole. I defend the claim that these are indeed cases in which you ought not to do your part in what the group ought to do. I then argue that we can expect these cases to produce a troubling kind of rational conflict.

1. Introduction

Many of us may be inclined to think that ethical considerations can bear not only on what each of us does as an individual, but also on what we do together. We may think that groups as well as individuals can have reasons to act in certain ways; that there are not only things that each of us ought to do, but also things that we as a group ought to do.¹ For example, you might think that the members of your band collectively ought to put on the best show they can, or that humanity as a group ought to take action to address climate change.

If there can be things that groups ought to do, it is natural to be interested in what this might mean for what individuals ought to do. In particular, we might wonder, could it be possible that, even though we together ought to do something, you as an individual ought not to do your part? And under what conditions, in particular, could that happen?

Now, one important kind of case in which you might be especially inclined to think that you ought not to do your part are those in which it seems that your participation would not have the right kind of causal influence on the collective activity. However, as I will point out, another kind of case, introduced by David Copp, illustrates the possibility that you ought not to do your part even when you would be playing a crucial causal role. This is because you may have special *agent-relative* reasons against participating that are not shared by the group as a whole.

In this article, I first defend the claim that these are indeed cases where we should accept that you ought not to do your part in what the group ought to do. I then

¹See Jackson 1987, Dietz 2016, Wringer 2016, Schwenkenbecher 2018, and Collins 2019.

argue that we can expect these ethical conflicts in turn to produce a troubling kind of rational conflict. I conclude by comparing the cases under discussion to other kinds of ethical conflicts.

2. When you ought not to do your part

Suppose, then, that you are part of the group that ought to perform some activity. What sorts of factors could make it the case that you ought not to do your part?

One important kind of reason is this: your participation might fail to have the right kind of causal influence on the group activity. If so, we might think, then even if there is a presumption that you ought to participate, this presumption is defeated. Your reasons to participate are undercut or canceled, and so can be easily outweighed by your reasons to perform alternative actions.

There are several ways in which your participation might fail to have the right kind of causal influence. First, there are cases where the combined actions of many people result in some morally significant outcome, and in these cases, it often seems that no single person's actions could make a morally significant difference to the group activity. For example, you might think that while humanity ought to reduce its carbon emissions, decreasing your own carbon emissions is surely not going to make a morally significant difference.²

There are also other kinds of case where your participation does not seem to have the right kind of causal influence. For example, your participation might not be sufficient for the group activity to take place because not enough other people would join you; or it might not be necessary, because there are already enough other people who are going to contribute.³

However, I suggest that a kind of case first introduced by Copp illustrates the possibility that a conflict between individual and collective reasons for action might arise even when we are not worried about whether your participation would play the right kind of causal role.⁴ This might happen because you might have independent reasons not to participate in the group activity, reasons which are not shared by the group as a whole.

I will rely here on a simplified variant of Copp's case:

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.⁵

I suggest that we have at least *prima facie* reason to accept the following judgments about this case: I ought to rescue my child, but we as a group ought to rescue the strangers.

²There has been significant discussion of these cases. For an overview, see Nefsky 2019.

³On individual reasons to participate in group activities when other members of the group will fail to do their parts, see Woodard 2008 and Dietz 2016: 969–73.

⁴See Copp 2007: 376–77. For discussion, see Ludwig 2007, Miller 2007, and Copp 2012. A related case is also discussed briefly in Björnsson 2020.

⁵I offered this case in Dietz 2016: 973–74. Copp's original example involves a prime minister who, Copp suggests, has a moral obligation to meet the demands of an outlaw group which is holding her child hostage, though the government as a whole has a moral obligation to reject the demands. Several of the objections to Copp that are raised in the exchange between Copp, Ludwig, and Miller revolve around the prime minister's institutional role, which is not a feature of my example.

Why should we think that I ought to rescue my child? On the one hand, we might think, I do have some reason to do whatever would result in lives being saved, and, other things equal, I have more reason to do what would result in more lives being saved. In addition, if it is true that the group ought to rescue the strangers, then I may have at least some reason to do my part in that activity, in virtue of my membership in the group.⁶ On the other hand, we might think, I also have reason to give extra weight to my child's life, in virtue of my relationship to my child. And this reason, we might think, is quite strong: strong enough to outweigh both the larger number of strangers' lives at stake, and any reasons I might have deriving from my membership in the group. So, all things considered, I ought to rescue my child.

It will be helpful to think about how these reasons can be classified in terms of the common distinction between *agent-relative* and *agent-neutral* reasons for action. It is a matter of debate how exactly this distinction should be drawn.⁷ But roughly, a reason is agent-relative when it is a fact that has some special relation to particular agents, whereas a reason is agent-neutral if it does not. For example, in this case, my reasons to do whatever would result in lives being saved, regardless of their relationship to me, are agent-neutral. My reason to do what would result in the survival of my own child, by contrast, is agent-relative. (My reason to do my part in what my group ought to do may also count as agent-relative, insofar as it involves how I myself would count as participating in a larger pattern of action.)

Next, why should we think that we as a group, by contrast, ought to rescue the strangers? Well, assuming that the group has the kind of agency enabling it to have reasons for action at all, then it seems that the group would also have agent-neutral reasons to do whatever would result in lives being saved, and other things equal, it would have stronger agent-neutral reasons to do what would result in more lives being saved. So the group's agent-neutral reasons favor the rescue mission over abandoning it so that I can rescue my child. But whereas the fact that one of the lives that could be saved is that of my child gives me an agent-relative reason, this fact does not, we can assume, apply to the group as a whole. So, at least initially, it seems that the group's only reasons are its agent-neutral reasons to do what would result in lives being saved. So, all things considered, the group ought to rescue the strangers.

Thus, it seems that we as a group ought to rescue the strangers, but I ought not to do my part, but rather ought to save my child. Again, I have only tried to show why we have *prima facie* reason to accept these conclusions; we will next consider possible strategies for resisting this argument.

If these are the right conclusions to reach about this case, then why would that be significant? Note that in this case, my participation would clearly play a significant causal role: my participation would be both necessary and sufficient (given your help) for the success of the rescue mission. Thus, this case illustrates the possibility that collective and individual reasons might come apart even when we bracket issues about the individual's causal role.

Before we consider this case in more detail, there are two issues that I want to point out. First, while Copp's discussion of his similar example was concerned with what the

⁶However, see Dietz 2016: 974–75 on concerns about “double-counting” reasons for action that individuals might have in virtue of group membership with closely related reasons for action that they might have independent of their group membership.

⁷For discussion of how to draw the distinction, see Ridge 2017.

group and its individual members each have a *moral obligation* or *duty* to do, my focus is on what they each *ought all things considered* or *have most reason* to do.

Second, much of the recent literature on the ethics of collective action has focused on the contrast between groups that are organized in ways that make it more plausible to consider them agents in their own right, and groups that lack these sorts of organization.⁸ The focus of this article will not be on this issue, but will rather be on the general issue of how the presence of agent-relative reasons may create tension between what a group ought to do and what its individual members ought to do. However, there are several points at which the structure of the group might seem to bear on the argument, so we will be returning to this issue below.

3. Resisting the conflict

I have just argued that we have at least *prima facie* reason to think that, in Rescue Mission, we as a group ought to rescue the strangers, but I ought not to do my part, since I ought instead to rescue my child. I will now consider, and respond to, two possible lines of resistance to these conclusions. First, opponents could focus on resisting my claims about our individual reasons. In particular, they could grant that the group ought to save the strangers, but argue that I in fact ought to do my part. Alternatively, opponents could focus on resisting my claims about our collective reasons. In particular, they could grant that I should rescue my own child, but deny that the group ought to rescue the strangers.

Let's start by seeing whether we can try to defuse the threat coming from our individual reasons. Was I wrong to suggest that we sometimes have overriding agent-relative reasons to defect from what the group ought to do?

Now, one way to resist this suggestion would simply be to deny that individuals ever have agent-relative reasons for action. But these kinds of reasons are standardly thought to be part of commonsense morality, and many people find them highly intuitive. So this option would come at a very high cost. (Those who are independently motivated to argue that we have only agent-neutral reasons for action, however, should take note: the fact that denying agent-relative reasons would get us out of these conflicts could at least somewhat strengthen their case for doing so.)

Second, we could claim that while individuals do have agent-relative reasons for action, they are always overridden in the relevant group contexts. This option is also unattractive. It is highly intuitive to think that I have much stronger reasons to save my own child from danger than to save two strangers, or perhaps even a large number of strangers. If we share this intuition, then we are likely to resist the idea that being part of a group which ought to save the strangers is enough to override my personal reasons.

Next, let's see if we can defuse the problem by focusing on the collective side of the conflict. As before, there are more and less radical ways we could try to do this.

One of the more radical options would be to abandon the idea of collective reasons altogether. But we would then be losing out on what a number of philosophers find to be a highly plausible and attractive idea, and a promising way to make sense of the judgments we want to make about a variety of examples.

Another strategy for resisting the idea that the group ought to save the two strangers which many will find radical would be to resist the idea that, quite generally, agents

⁸For introductions to this literature, see Schwenkenbecher 2018 and Collins 2019.

ought to save more lives rather than fewer, other things equal.⁹ There is much to say about this issue, but because my focus here is on questions specifically about the relationship between individual and collective reasons for action, I will bracket this issue for the purpose of this discussion.

A less radical option would be to claim that what we together ought to do can itself be affected by the kinds of considerations that give rise to our individual agent-relative reasons. For example, in the rescue mission case, it might be claimed that the group does have a special reason to save my child, or at least to act in a way that allows me to save my child. More generally, we might claim that even if I have a reason to bring about some outcome in virtue of some fact that does not apply to the group to which I belong, this always provides the group with a reason to bring about, or at least to allow, this outcome. We can think of this as the idea that just as the group's reasons can transmit down to me, giving me group-based reasons to do my part, my personal reasons can also transmit up to the group.

However, even if we did think that individual reasons transmit up to the group, it's not clear that it would solve the problem. Suppose that rather than saving strangers, our rescue mission would save the lives of your two children, but that my doing my part would once again make it impossible for me to rescue my child. In that case, the group would get reasons for action corresponding to your reasons to rescue your two children, which would presumably be stronger than whatever reasons it got corresponding to my reason to rescue my one child. So again, it seems that we together ought to undertake the rescue mission. But it is still plausible that I ought to defect, and rescue my own child.

Next, rather than claiming that a group has extra reason to *favor* actions in which its members have an interest, we might instead claim that groups should *disregard* actions that would be incompatible with its members' agent-relative reasons. More precisely, we might claim that if an individual would have decisive reason not to do her part in a group action even when she would be playing the right causal role, then this prevents that action from being one that the group ought to perform. In that case, in Rescue Mission, we would say that because A ought to save her child, it is not the case that the group ought to save the strangers.¹⁰

Here is why I am skeptical of this proposal. What would we have to say about the group's reasons for action in order for the proposal to be true? The most natural answer seems to be that when a particular group action would be incompatible with its members' agent-relative reasons, then the group's reasons to perform that action are thereby either canceled or excluded from consideration. In that case, we would have to say that

⁹See Anscombe 1967 and Taurek 1977. For a recent overview and discussion of this issue, see Hirose 2014.

¹⁰A related view that we might be attracted to is that if an individual is not *morally permitted* to do her part in a group action, that prevents the group from having a moral obligation to perform the action (see Collins 2019: ch. 7). In our case, it is plausible that you are not morally permitted to participate in the rescue mission, because you have a moral obligation to rescue your child, so this view would imply that the group would therefore lose any moral obligation it might have to rescue the strangers. However, there are structurally similar cases where this view would not apply: cases where you may have strong agent-relative reasons to defect from the group action but where you may still be morally permitted to participate. For example, it might be that in participating, you would be sacrificing your own life, rather than your child's life. Plausibly, sacrificing your own life would be something that you have strong agent-relative reasons not to do but are nevertheless morally permitted to do. So in this case, the view under discussion would not tell us anything about whether the group ought to save the strangers.

the group has *no reason*, or no reason that it should take into account, to save the strangers. But I find this implausible: surely there is at least something to be said for an action that would save two people's lives.

Here is an additional argument to support this intuition. Suppose that rather than rescuing any of the people in danger, the group could instead spend the day playing checkers. The following claim seems hard to deny: the group ought to save the strangers *rather than* play checkers. So surely the group must have *more reason* to save the strangers than to play checkers. So surely we should conclude, again, that the group has at least *some* reason to save the strangers. There are two ways in which we might avoid this conclusion. First, instead of claiming that the group has some positive reason to save the strangers, we might claim that the group has reason not to play checkers, and that this is why it is true that the group has more reason to do the former. However, I find this implausible: surely there is nothing wrong in itself with playing checkers; the point is just that the group would be failing to do something (save lives) that it has a compelling reason to do. Second, we might adopt the thesis that reasons can be essentially contrastive, and therefore that the claim that an agent (whether a group or an individual) has reason to do X rather than Y need not imply that she has any reason simpliciter either to do X or not to do Y.¹¹ I grant that this is a possible way out, but note that this route requires us to adopt a controversial view about the structure of reasons for action.

Now, again, much of the recent literature on the ethics of collective action has focused on the contrast between groups that are organized in ways that make it more plausible to consider them agents in their own right, and groups that are not. And, again, this article is not defending any particular view on this issue. But I want to note here why I do not think that this issue materially affects the case for thinking that there can be normative conflicts of the sort that I have been defending.

On the one hand, we might think that if the group does not have the right sort of structure, then it is not appropriate to talk of it as having normative reasons for action at all. In this event, in order for the Rescue Mission case to generate normative conflict, this just means we need to stipulate that the group *does* have sufficient structure: that, for example, the group has established roles and decision-making procedures.

Alternatively, we might think while that relatively unstructured groups can possess normative reasons for action, what the group has reasons to do is derived from its members' reasons, in such a way that the group could have a reason to do something only if its members would all have sufficient reason to participate. For example, Gunnar Björnsson has defended a related view about moral obligation, as opposed to normative reasons.¹² Björnsson considers a case in which each of three parents has to decide whether to save her own child or participate in a risky attempt to save a larger number of children (along with her own child). In this case, Björnsson suggests, supposing that each parent is not morally obligated to risk her own child's life by participating, it is intuitive that the group as a whole cannot have a moral obligation to save the larger number of children.

This type of view, however, faces the same problems that we saw earlier. If we want to deny the claim that the group in Rescue Mission ought to save the strangers, then what would we have to say about the group's reasons for action? As we have seen, we have two options. We could say that the group itself has an agent-relative reason to favor options in which members have a stake, but this would not solve the problem in the

¹¹See Snedegar 2017.

¹²Björnsson 2020. Björnsson also provides a helpful summary of other views along these lines.

“my child vs. your children” version of the case. Or we could say that the group lacks the agent-neutral reason to save the strangers, but this claim, I have argued, is implausible. The view in question, that the group has a reason to do something only if its members would all have sufficient reason to participate, simply takes this latter horn of the dilemma. (Note that it matters that we are talking here in terms of the group’s reasons for action rather than in terms of its moral obligations. I do not think that it would be implausible to claim that the group lacks a moral obligation to rescue the strangers, as Björnsson’s view implies. But it does seem implausible to me to say that the group does not have any reason to do this.) And finally, even if we do accept a view like this about unstructured groups, then, again, we can simply stipulate that the group in Rescue Mission *does* have the right sort of structure to count as a full-fledged collective agent.

I have defended the idea that, in cases like Rescue Mission, it can be true both that we as a group ought to perform some action, and also that I ought not to do my part. To conclude this section, I want to note that accepting that there can be normative conflicts of this kind does not mean denying that there is any important relationship between collective and individual reasons or related normative notions. For example, Bill Wringer has argued that collective obligations entail individual obligations to do one’s part so long as others would be reasonably likely to do theirs.¹³ However, in response to the type of case highlighted by Copp, Wringer specifies that these may only be *pro tanto* individual obligations, and so potentially overridable by other considerations. Similarly, I have granted that if the group ought to rescue the strangers, then this might mean that I must have *some* reason to do my part, but have suggested that this reason is plausibly outweighed by my agent-relative reason to save my own child.

4. Rationality and paralysis

I have argued that even if a group to which you belong ought to do something, and even if your participation would play a significant causal role, it might nevertheless be true that you ought not to do your part. I will now argue that such cases, if they exist, would put us into a troubling sort of rational conflict.

Suppose that, in such a case, both the group as a whole and the individual members of the group know both the facts of the situation and what they ought to do. If so, I argue, then it will be impossible for both the group and the individuals to be fully rational.

To start, it is plausible that rational agents will be enkratic: they will intend to do what they believe they ought, all things considered, to do. And if we accept enkrasia as a standard for individuals, and we also accept that there can be things that groups of people ought to do, then it seems that we should likewise accept enkrasia as a standard for groups.¹⁴

What will this mean in cases like Rescue Mission? Suppose that, as I suggested, I ought to rescue my child, but we as a group ought to rescue the strangers. And suppose that I individually and we collectively agree with these judgments. In that case, if I am

¹³Wringer 2016: 487–88.

¹⁴Again, note that Copp’s claim about his original case is about the moral obligations of the relevant parties, rather than about what they ought to do all things considered. This is significant here because if we think that agents might sometimes have most reason to violate their moral obligations, we may think that they could rationally do so, whereas we might think that it can never be rational for agents to fail to intend to do what they believe they ought to do all things considered.

enkratic as an individual, then I will intend to rescue my child. And if we are enkratic as a group, then we will intend to rescue the strangers.

However, these intentions seem to be in tension. This, I suggest, is because if we collectively intend to act in some way, and doing my part is necessary for the performance of the group activity, then it must be true that I intend to do my part.¹⁵

If collective intentions do imply corresponding individual intentions, and we as a group intend to rescue the strangers, then I will intend to do my part. But again, if I am enkratic as an individual, and know that I ought to rescue my child, I will also intend to do that. So if we are enkratic both individually and collectively, then I will find myself with two conflicting intentions. But plausibly, I cannot rationally retain two intentions if I know that I cannot carry out both. Thus, if the group is fully rational, then I must not be: for the group to be enkratic, I must either fail to intend to do what I believe I ought to do, or retain two transparently conflicting intentions. And if I am rational, then the group must not be rational: I must abandon the intention to do my part in rescuing the strangers, in which case the group will no longer meet the standard of enkrasia.

To be clear, I have not been arguing that in this situation it would be impossible for me as an individual to be fully rational. (Nor have I argued that it is impossible for us as a group to be fully rational.) Instead, what I have argued is that it is impossible for both the individual and the group to be fully rational at the same time, as long as we recognize what we ought to do.

I have argued that the kind of rational conflict that I have described will arise in cases like Rescue Mission. Again, this case illustrates how it might be possible for collective and individual reasons to come apart even when the individual's participation would be both necessary and sufficient for the performance of the group activity. However, this conflict will not arise in cases where we know that my participation would not be necessary, and/or would not be sufficient. If we knew my participation would not be necessary, then since doing what I individually ought to do will not prevent us from doing what we collectively ought to do, we could intend to do what we collectively ought to do without my intending to do my part. If we knew that my participation would not be sufficient, then we must believe that we will not do what we collectively ought to do. But in that case, if, as is widely accepted, one cannot intend to do something that one believes one will not do, we will not be able to form an intention to do what we collectively ought to do.

Now, we saw earlier that if the group does not possess the right kind of structure, then some might be skeptical that it can possess reasons for action at all, or might claim that if it does, then it could only have reasons for activities in which its members would all have sufficient reason to participate. In response, I suggested that we can simply stipulate that the group in Rescue Mission does have the sort of structure to qualify as a full-fledged collective agent. However, might this stipulation undermine my claim that cases like Rescue Mission will produce rational conflict?

The most tempting reason for skepticism here, it seems to me, would be to think that, in structured groups, the connection between collective intentions and individual intentions is not as strong as my argument requires. After all, we might think, when it comes to paradigmatic structured groups, it seems natural to talk about the group intending to do things in which some of its members do not intend to participate.

¹⁵For discussions defending a strong link between collective and individual intention, see Tuomela and Miller 1988, Searle 1990, Bratman 1999, and Kutz 2000.

For example, we might talk about a country planning to launch a war even if many of its citizens are not willing to fight or to contribute to the war effort.

However, recall that the claim that my argument relied on was that we can collectively intend to do something only if I intend to do my part *when doing my part is necessary for the performance of the group activity*. When we think about structured groups, and particularly very large structured groups such as countries, it is natural to think of examples where an individual's participation is not necessary, and that, I suggest, may be why it seems natural to think that collective intentions need not imply individual intentions in these contexts. But even in structured groups, it is clearly possible for an individual's participation to be necessary (and sufficient) for the performance of the group activity. After all, structured groups need not be large, and even when they are, particular individuals may still play a crucial role. And when we do suppose that the individual's participation is necessary for the group activity, then it still seems to me that the group will count as intending to perform the activity only if the individual intends to do her part.

Again, I have argued that in cases like Rescue Mission, it will be impossible for both the group and its individuals to be rational. This conclusion is theoretically interesting. But what I find especially compelling about this situation is how it would feel from the inside. While I have argued that either the group or the individuals in this situation must fail to be *fully* rational, it does seem possible that they might both be *moderately* rational. In other words, it seems possible that, both as a group and as individuals, we might mean well; we might try to do the right thing. To make this more precise, suppose that both we as a group and I as an individual are moderately enkratic, in that we are disposed to form an intention to do what we believe we ought to do when we are *consciously thinking* about what we ought to do. If so, then when I think about what I ought to do, I will tend to form an intention to rescue my child, but when we together consider what we ought to do, we will tend to form an intention to rescue the strangers, and I will form an intention to do my part.

As a result, it seems to me, this 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. And I will next argue that this conflict, unlike other more familiar kinds of ethical conflict, is one which ethical reflection is powerless to resolve.

Now, it might be suggested that this paralysis is unremarkable. After all, it might be argued, given that I believe I ought to rescue my child, and cannot do this if I do my part in rescuing the strangers, what rationality requires is for me to intend simply to rescue my child. Insofar as I find myself intending to help rescue the strangers, then, I am being irrational. And we should not find it remarkable that an agent could be paralyzed in virtue of alternating between a rationally required intention and a rationally prohibited intention.

However, this is not an ordinary conflict between rationality and irrationality. Instead, what is significant about this situation is that the rationally prohibited intention, for the agent to help rescue the strangers, is itself a product of rationality. That is, I will form this intention only because I am part of a group that has a rational disposition to intend to do what it ought to do. So the conflict in this case is not simply between my own rationality and my own irrationality, but rather between the rationality that I possess as an individual and the rationality of a group of which I am a member.

5. Unit-of-agency conflicts and moral dilemmas

I have argued that there could be situations where, although we as a group ought to act in some way, and although my participating would play a significant causal role, I nevertheless ought not to do my part. And such cases, I have argued, can put us in a paralyzing situation. I will conclude by considering how these cases are related to a more familiar sort of ethical conflict: moral dilemmas.

Broadly speaking, we can divide moral dilemmas into two categories. The first and most commonly discussed category is that of *single-agent dilemmas*. These are cases in which a single agent faces a conflict between ethical requirements of some kind: for example, between the requirements of morality and the requirements of self-interest, or between two incompatible moral obligations. The strongest sort of single-agent dilemma would be a case in which (1) an agent has to choose between two incompatible options A and B, and (2) it is true both that the agent ought *all things considered* to do A, and that the agent ought *all things considered* to do B. It is controversial, however, whether cases of this sort are possible.¹⁶

A second sort of case is what we might call *multi-agent dilemmas*: cases in which one agent ought to perform one action, and another agent ought to perform another action, but where it is not possible for both actions to be performed.¹⁷ For example, we might think that Antigone ought to arrange for her brother's burial, while Creon ought to prevent it.¹⁸

It is worth emphasizing that the type of conflict that we have been focusing on is not a type of single-agent dilemma. What we have been focusing on is not the idea that a particular agent might be subject to two incompatible ethical requirements, but rather the idea that there might be a conflict between the ethical requirements that apply to different levels or units of agency. For example, the issue is not simply that the agent seems to have both an obligation to her child and an incompatible obligation to do her part in the rescue mission. Rather, the issue is to understand the nature of conflicts between what the *group as a whole* ought to do and what individual members of the group ought to do.¹⁹ In particular, we have been interested in how it might turn out that *we* ought to act in some way, but *I* ought not to do my part – where this claim about what I ought to do has already taken into account any group-based normative reasons that I might have.

The cases that we have been focusing on might be classed as multi-agent dilemmas. Typically, those who discuss these cases have in mind examples where both agents are individual people, as in the Antigone/Creon example. But our cases could be thought of as multi-agent dilemmas where one of the agents is an individual, and the other agent is a group. However, if so, the type of conflict that we have been focusing on would clearly represent a peculiar sort of multi-agent dilemma: one in which one of the agents is part of the other.

¹⁶See Thomson 1990, Brink 1994, and Goble 2009.

¹⁷See Marcus 1980: 122, McConnell 1988, and Taylor 2013.

¹⁸This example is taken from Marcus 1980: 122. Marcus describes the case in terms of what agents are *obliged* to do rather than *ought* to do.

¹⁹Similarly, the conflict is not a conflict between any particular agent's agent-neutral and agent-relative reasons. The reason why it is significant that my reasons to prefer that my child be saved are agent-relative is that this means that the group as a whole (a different agent, or unit of agency) may not have the same reasons. This is what makes it possible that we as a group ought to do what brings about some outcome but I ought not to do my part, even if I would be playing a crucial causal role in bringing about that outcome.

Although our cases are not single-agent dilemmas, there is one important feature that they have in common: single-agent dilemmas of the “all things considered” kind could also be paralyzing.²⁰ And we can explain why this would happen on grounds similar to those offered above. If I am enkratic, and I recognize that I ought all things considered to do A and ought all things considered to do B, then I will intend to do A, and will intend to do B. But since I also know these intentions to be in conflict, it seems that I will be unable to decide what to do until I can somehow settle on one coherent intention.

In contrast, it does not seem that the typical examples of multi-agent dilemmas, where the agents are both individuals, would be paralyzing. After all, even if both Antigone and Creon, for example, are enkratic, and so Antigone intends to arrange the burial while Creon intends to prevent it, there is no reason for thinking that either agent will themselves end up with conflicting intentions. Even if Antigone intends one thing, that does not imply that Creon intends to act in accordance with Antigone’s intentions, or vice versa. But in our cases, where one agent is a group and the other agent is an individual member of that group, and where that individual would be playing a crucial causal role in the group action, I have suggested that things are different: the group’s intending to perform the group action does imply that the agent will intend to do her part.

However, I will now argue that it is significant that our cases are not single-agent dilemmas, because one of the most promising strategies we might use to resolve our paralysis in single-agent dilemmas cannot be applied to our cases.

As we have seen, single-agent dilemmas are cases in which some agent faces a conflict between ethical requirements. We have seen that the strongest sort of single-agent dilemma, in which an agent faces incompatible options and ought to perform each of them all things considered, would be paralyzing. But there are also other weaker forms of single-agent dilemma, such as cases in which an agent has to choose between two incompatible moral obligations, or between doing what she morally ought to do and doing what she ought to do from a self-interested point of view. While it may be difficult to decide what to do in these cases, however, that does not mean that ethical reflection cannot provide us with a way to choose. After all, we might think, we can understand the conflicting ethical requirements as each contributing to a common standard, and we can then refer to this common standard in deciding what to do.

For example, suppose that I am confronted with what seem to be two incompatible moral obligations. In this case, we might think, these are really only *pro tanto* obligations, which contribute to what I am morally obligated to do all things considered; my ultimate moral obligation is to fulfill the more stringent obligation, or, if they are equally stringent, just to fulfill one or the other. Or suppose that I have to decide between doing what I morally ought to do and acting in accordance with self-interest. In this case, we might think, while I might have moral reasons for favoring one option and self-interested reasons favoring another option, these are both just two species of a broader genus, reasons for action simpliciter. I can decide what to do by considering all of my reasons for action, which determine what I ought to do all things considered.

This analogy might lead us to think that we can rationally adjudicate the conflict between individual and collective reasons in the same way. Individual and collective reasons, we might think, are just another way of carving up reasons for action simpliciter.

²⁰This seems to be suggested in Thomson 1990: 83 and Thomson 2008: 174–75. For a reply, see Horty 2003: 588–89.

We can resolve our struggle to decide what to do by considering what we ought to do all things considered, taking both individual and collective reasons into account.

However, the analogy, I think, is misleading, and the problem is much more intractable than it suggests. For there does not seem to be any agent who has both individual and collective reasons to balance against one another; again, as in more familiar multi-agent dilemmas, the conflict is a conflict between what two different units of agency ought to do. By definition, all of the reasons that I have – including any reasons I have to do my part in what the group ought to do – are individual reasons, and all of the reasons that the group has are collective reasons. I already know what I ought to do all things considered, and the group already knows what it ought to do all things considered. For the suggestion to work, it seems, we would have to be able to specify some third unit of agency that encompasses both the individual and the collective, and claim that individual and collective reasons are both ultimately possessed by this neutral agent. But it seems difficult to make sense of what this third unit of agency would be. So it seems difficult to see how this type of strategy could help us to resolve these conflicts.

References

- Anscombe, G. E. M.** 1967. Who Is Wronged? *Oxford Review* 5, 16–17.
- Björnsson, Gunnar.** 2020. Collective Responsibility and Collective Obligations Without Collective Moral Agents In Saba Bazargan-Forward and Deborah Tollefsen (eds), *The Routledge Handbook of Collective Responsibility* (New York: Routledge, 2020), 127–41.
- Bratman, Michael.** 1999. I Intend that We J. In *Faces of Intention* (Cambridge: Cambridge University Press, 1999) 142–61.
- Brink, David O.** 1994. Moral Conflict and Its Structure. *Philosophical Review* 103.2, 215–47.
- Collins, Stephanie.** 2019. *Group Duties: Their Existence and Their Implications for Individuals* (Oxford: Oxford University Press).
- Copp, David.** 2007. The Collective Moral Autonomy Thesis. *Journal of Social Philosophy* 38.3, 369–88.
- Copp, David.** 2012. The Collective Moral Autonomy Thesis: Reply to Ludwig and Miller. *Journal of Social Philosophy* 43.1, 78–95.
- Dietz, Alexander.** 2016. What We Together Ought to Do. *Ethics* 126.4, 955–82.
- Goble, Lou.** 2009. Normative Conflicts and the Logic of “Ought”. *Noûs* 43.3, 450–89.
- Hirose, Iwao.** 2014. *Moral Aggregation* (Oxford: Oxford University Press).
- Horty, John F.** 2003. Reasoning with Moral Conflicts. *Noûs* 37.4, 557–605.
- Jackson, Frank** (1987) Group Morality. In J. J. C. Smart, Philip Pettit, Richard Sylvan, and Jean Norman (eds), *Metaphysics and Morality: Essays in Honour of J. J. C. Smart* (Oxford: Blackwell) 91–110.
- Kutz, Christopher.** 2000. Acting Together. *Philosophy and Phenomenological Research* 61, 1–31.
- Ludwig, Kirk.** 2007. The Argument from Normative Autonomy for Collective Agents. *Journal of Social Philosophy* 38.3, 410–27.
- Marcus, Ruth Barcan.** 1980. Moral Dilemmas and Consistency. *Journal of Philosophy* 77.3, 121–36.
- McConnell, Terrance.** 1988. Interpersonal Moral Conflicts. *American Philosophical Quarterly* 25.1, 25–35.
- Miller, Seumas.** 2007. Against the Collective Moral Autonomy Thesis. *Journal of Social Philosophy* 38.3, 389–409.
- Nefsky, Julia.** 2019. Collective Harm and the Inefficacy Problem. *Philosophy Compass* 14.4, <https://doi.org/10.1111/phc3.12587>.
- Ridge, Michael.** 2017. Reasons for Action: Agent-Neutral vs. Agent-Relative. In *Stanford Encyclopedia of Philosophy*, <https://plato.stanford.edu/archives/fall2017/entries/reasons-agent/>.
- Schwenkenbecher, Anne.** 2018. Making Sense of Collective Moral Obligations: A Comparison of Existing Approaches. In Kandy Hess, Violetta Ignieski, and Tracy Isaacs (eds), *Collectivity: Ontology, Ethics, and Social Justice* (London: Rowman and Littlefield, 2018), 109–32.
- Searle, John.** 1990. Collective Intentions and Actions. In Philip R. Cohen, Jerry Morgan, and Martha E. Pollack (eds), *Intentions in Communication* (Cambridge, MA: MIT Press, 1990), 401–15.

- Snedegar, Justin.** 2017. *Contrastive Reasons* (Oxford University Press).
- Taurek, John M.** 1977. Should the Numbers Count? *Philosophy and Public Affairs* 6.4, 293–316.
- Taylor, Erin.** 2013. Irreconcilable Differences. *American Philosophical Quarterly* 50.2, 181–92.
- Thomson, Judith Jarvis.** 1990. *The Realm of Rights* (Cambridge, MA: Harvard University Press).
- Thomson, Judith Jarvis.** 2008. *Normativity* (Chicago: Open Court).
- Tuomela, Raimo and Miller, Kaarlo.** 1988. We-Intentions. *Philosophical Studies* 53, 367–89.
- Woodard, Christopher.** 2008. *Reasons, Patterns, and Cooperation* (New York: Routledge).
- Wringe, Bill.** 2016. Collective Obligations: Their Existence, Their Explanatory Power, and Their Supervenience on the Obligations of Individuals. *European Journal of Philosophy* 24.2, 72–497.