

What We Together Ought to Do*

Alexander Dietz

I argue that we have not only individual reasons for action but also collective reasons for action: reasons which apply to us as a group. I next argue that if we together have a reason to act, then I may have a reason to do my part, but only when others will do theirs. Finally, I argue that collective reasons to do good can never make a difference to what individuals ought to do, but that other kinds of reasons can.

I. INTRODUCTION

When we think about what we should do, we are often thinking about what each of us should do. But we can also think about whether there is anything that we together ought to do. For example, consider the effect that humans are now having on the global climate. It may be that, given what everyone else is doing, no one of us is making things any worse. But our combined activities are clearly doing harm. This fact might give us, as a group, a reason to change our behavior.

In this article, I will ask two questions:

- (1) What, if anything, do we together have reasons to do?

and

- (2) If we together have reasons to act in certain ways, how, if at all, does this affect what each of us ought to do?

* I would like to thank a number of people for their feedback and discussion, including Stephen Bero, Renee Jorgensen Bolinger, Erik Encarnacion, Joe Horton, Nathan Robert Howard, Nicholas Laskowski, Julia Nefsky, Alexander Sarch, Aness Webster, Ralph Wedgwood, Christopher Woodard, the editors and referees at *Ethics*, and audiences at the University of Southern California Speculative Society, the Rocky Mountain Ethics Congress, Collective Intentionality IX, and the University of Rennes. I'm especially grateful to Abelard Podgorski, Jacob Ross, Mark Schroeder, and Jonathan Quong.

Ethics 126 (July 2016): 955–982

© 2016 by The University of Chicago. All rights reserved. 0014-1704/2016/12604-0004\$10.00

I will not offer complete answers to these questions here but will try to make progress, building on work by Derek Parfit, Frank Jackson, Donald Regan, Christopher Woodard, and others.¹

In response to the first question, I will argue that groups do in fact have certain reasons for action. More specifically, I will argue that, to a limited but still significant extent, groups have reasons of the same kinds as those which individuals have. These include reasons to make outcomes better, not to harm people in certain ways, and to benefit themselves.

In response to the second question, I will argue that if we together have a reason to act in some way, then I have a reason to do my part, but only so long as others will do theirs. I will therefore argue that we should reject stronger claims advanced by Woodard and Jackson. In addition, I will argue, my reasons to do my part should not add up with certain similar reasons that I might already have.

I will then discuss what these answers imply. One natural aspiration we might have for a theory of collective reasons is to help support the conclusion that each of us should contribute to good causes such as preventing climate change, even when our own contribution would not make a difference. But I will argue that in fact, our general collective reasons to do good might never affect what each of us should do. Still, I will argue, claims about what we together ought to do are worth making even where they do not have implications for individuals. And in some cases, I will argue, such claims can have important implications for what each of us ought to do.²

II. PRELIMINARIES

Groups and Joint Action

We can start with some terminology. I will be using “group” and “we together” to refer to sets of persons, considered as units. In my general prin-

1. Compared with discussions of collective responsibility, and of the nature of collective action, very little has been written on collective duties and reasons for action. Important exceptions include Derek Parfit’s discussion in *Reasons and Persons* (New York: Oxford University Press, 1984), sec. 26; his later “What We Together Do” (unpublished manuscript); and Frank Jackson, “Group Morality,” in *Metaphysics and Morality*, ed. Philip Pettit, Richard Sylvan, and Jean Norman (Oxford: Blackwell, 1987), 91–110. Donald Regan’s *Utilitarianism and Co-operation* (New York: Oxford University Press, 1980) and Christopher Woodard’s *Reasons, Patterns, and Cooperation* (New York: Routledge, 2008) also discuss closely related issues. However, in recent years, more work directly addressing collective duties and reasons for action has started to appear, including Bill Wringe, “Collective Obligations: Their Existence, Their Explanatory Power, and Their Supervenience on the Obligations of Individuals,” *European Journal of Philosophy* (2014), <http://onlinelibrary.wiley.com/doi/10.1111/ejop.12076/abstract>; and a special issue of *Midwest Studies in Philosophy* on “forward-looking collective responsibility” (September 2014).

2. I am using “ought” and “should” in the “all-things-considered” sense, not in a more narrowly moral sense. I will be assuming that we ought to do whatever we have the strongest reasons to do but will not be assuming that what we ought to do is the same as what we have an obligation to do, moral or otherwise.

ciples, I will use them to refer to any set of persons capable of acting together in a certain way, specified below. I will not be referring only to organized groups or to related entities such as clubs, corporations, or sports teams and will not be discussing issues specific to these kinds of groups.

Next, some assumptions. I will assume that it is possible for people to intentionally act together. I will not endorse any particular account of the nature of this kind of action.³ But I will assume that when we make an explicit, sincere agreement for us to carry out some plan of action, each trusting the others to do their parts, and when each in fact does his part, this is enough for it to be true that we have intentionally acted together. I will use the terms “group action,” “joint action,” and “collective action” to refer to action of this explicit-agreement kind, except when otherwise noted.

I intend my normative claims to apply to those sets of persons who can act together in this explicit way and will assume that the people in my examples can act together in this way. But my claims will not be limited to sets of persons who do ever act together in this way. That is, there might be things that certain groups ought to do, even if they have never coordinated in the past and will never coordinate in the future.

Collective Reasons

When we say that some group of people should act in some way, we might merely be saying that all or most of the members of this group should act in this way. But when I talk about what we “together ought” or “together have a reason” to do, or what a group of people ought or has a reason to do, or our “collective reasons,” I will not merely be saying that all or most of the members of this group have reasons to do their parts in the relevant course of action. As I will argue later, there can even be cases in which, though you and I together have decisive reasons to act in some way, I have no reason to do my part. And when we each do have reasons to do our parts, this may be partly in virtue of the reasons we have as a group, rather than the other way around. Rather, I will be talking about a kind of reason whose bearer is the group itself.

In addition, I will focus on what collective reasons we have for acting given the facts of our situation, rather than what reasons we have given

3. Leading accounts of collective action include those offered by Margaret Gilbert and Michael Bratman. See Gilbert, “Walking Together: A Paradigmatic Social Phenomenon,” *Midwest Studies in Philosophy* 15 (1990): 1–14; and Bratman, *Faces of Intention: Selected Essays on Intention and Agency* (New York: Cambridge University Press, 1999). For more recent treatments, see Gilbert, *Joint Commitment: How We Make the Social World* (New York: Oxford University Press, 2014); and Bratman, *Shared Agency: A Planning Theory of Acting Together* (New York: Oxford University Press, 2014). See also Christian List and Philip Pettit, *Group Agency: The Possibility, Design, and Status of Corporate Agents* (New York: Oxford University Press, 2011).

our evidence or beliefs. So I will mostly ignore questions about what to do when we don't know all the relevant facts.

Collective Reasons and Collective Responsibility

It might be assumed that this article is intended to add to the extensive literature that has been written on collective responsibility, which, as Tracy Isaacs puts it, is typically understood in terms of "retrospective assessments of the actions of collective agents according to which they are blameworthy or praiseworthy."⁴ But the notions of what we ought to do, and of what we have reasons to do, are distinct from this notion of responsibility.⁵ For example, I might have strong reasons not to act in some way even if, because the temptation is hard to resist, I could not be blamed for doing so.

These issues may not be entirely independent. For example, some of our strongest reasons for action may be tied to our moral obligations. And it may be that we can be morally obligated to act in some way only if we would be blameworthy if we failed to do so without an excuse. If so, then significant questions about collective reasons for action will depend on what, if anything, groups can be blamed for. Nevertheless, I will here focus more directly on what groups have reasons to do.

III. WHAT SHOULD WE DO?

Are Collective Reasons Like Individual Reasons?

We can now address our first question: What, if anything, do we together have reasons to do?

We can start to answer this question, I propose, by considering the reasons that individuals have for acting, and asking whether groups might have similar reasons. For example, if individuals always have a reason to keep their promises, then groups might always have a reason to keep *their* promises. There is a range of possibilities here. At one extreme, there might be full parity. The very same reasons might apply to groups and to individuals. At the other extreme, the kinds of facts that give groups reasons for acting might be entirely unlike those that give reasons to individuals.

4. Tracy Isaacs, "Collective Responsibility and Collective Obligation," *Midwest Studies in Philosophy* 38 (2014): 40–57, 40. Some key articles on collective responsibility are collected in Larry May and Stacey Hoffman, eds., *Collective Responsibility: Five Decades of Debate in Theoretical and Applied Ethics* (Savage, MD: Rowman & Littlefield, 1991). See also Peter A. French, *Collective and Corporate Responsibility* (New York: Columbia University Press, 1984); Larry May, *The Morality of Groups* (Notre Dame, IN: University of Notre Dame Press, 1987); and Christopher Kutz, *Complicity: Ethics and Law for a Collective Age* (New York: Cambridge University Press, 2000).

5. Both Isaacs and Bill Wringe make a similar claim about the need to distinguish collective responsibility and collective obligation. See Wringe, "Collective Obligations," sec. 2; and Isaacs, "Collective Responsibility and Collective Obligation," 40.

To defend a full-parity thesis, it might be argued that our ethical theory should apply to all agents in the same way. Insofar as a group is an agent, it must have the same basic reasons for action as any other agent. While what groups and individuals have reasons to do might often differ in particular cases, this will merely be due to differences in the circumstances they happen to face.⁶

However, there also seem to be fundamental differences between the nature of an individual person and the nature of a group, and not merely in their circumstances. And at least two such differences seem to have normative importance.

First, individuals are capable of certain mental states, including beliefs, desires, guilt, pleasure, pain, and continuous conscious experience, mental states which many people take to be closely linked to our reasons for action. However, it might be pointed out, it is controversial whether or in what ways groups have mental states of these kinds. And it is certainly doubtful whether all groups who can act together have such mental states.

Second, individuals and groups do not seem to be capable of the same kinds of relationships. For example, individuals might have reasons for action related to love, friendship, and family relationships that could not apply to groups. In addition, since groups, unlike individuals, have members, certain kinds of groups, such as states or corporations, might have obligations of justice toward their members that could not apply to individuals.

Unless it can be shown that these differences are either illusory or unimportant, I conclude that we should be skeptical of the claim that the very same reasons apply to groups as apply to individuals. Therefore, to fully understand what we together ought to do, we would have to figure out which individual reasons do not apply to groups, and which reasons groups have that are unique to them.

Nevertheless, in the following sections, I will offer arguments in support of the idea that there is significant parity between individual and group reasons. I will focus on three important kinds of reasons: reasons to make outcomes better, to avoid harming people in certain ways, and to benefit ourselves. I will assume, as many people believe, that we have individual reasons to act in these ways. However, I will argue that there is also significant evidence that we have collective reasons to do so.

I will support these claims, not by general considerations about the sense in which ethical theory should apply to all agents universally, but rather by focusing on cases. In order to capture our intuitions about these cases, I will argue, it will not be enough to appeal to individual reasons.

6. I owe this suggestion to Mark Schroeder.

However, we can capture these intuitions in a simple and satisfying way by appealing to collective reasons of the relevant kinds. Therefore, these cases represent significant evidence in favor of the claim that we have these collective reasons.

Beneficence

Start with our reasons to make outcomes better. Most people believe that each of us has some reason to do what will make things better. I claim that, in addition, we together have such reasons. I will support this claim with the following case described by Allan Gibbard.⁷ Suppose that you and I each have two options, with outcomes as follows:

		You	
		do A	do B
I	do A	Second-best	Bad
	do B	Bad	Best

Suppose further that the second-best outcome would be much worse than the best outcome, and that the bad outcomes would be extremely bad. We can call this *Gibbard’s Case*.⁸

As Derek Parfit writes, it is “in some sense obvious” that, other things equal, we should both do B and bring about the best outcome.⁹ Likewise, it seems to me in some sense obvious that we should not both do A, and bring about the second-best outcome. But in what sense are these claims correct?

On the account I will defend, these claims are correct in the sense that we together ought to carry out the course of action in which we each do B. We have decisive reasons as a group to carry out this course of action, in virtue of the fact that we would together be making things better.

7. Allan Gibbard introduced a case with this structure in “Rule-Utilitarianism: Merely an Illusory Alternative?”, *Australasian Journal of Philosophy* 43 (1965): 211–20, at 214–15. Donald Regan extensively discusses a version of this case in *Utilitarianism and Co-operation*.

8. For a more specific version of this case, imagine that you and I are industrialists who have been asked to release our supplies of a certain chemical into the atmosphere (option B), as doing so will reverse the effects of global warming. If we both refuse (option A), global warming will proceed unchecked, causing millions of deaths. However, this chemical has certain dangerous properties which are neutralized only when there is enough of it dispersed throughout the air. The chemical will be harmless if we both release our supplies. But if only one of us does, this will cause changes to the environment far more catastrophic than the effects of global warming.

9. Parfit, “What We Together Do,” 11.

And we together have decisive reasons not to do A, in virtue of the fact that we would together be making things worse than they could be.¹⁰

This account, I will argue, has significant advantages over two alternative ways of accounting for Gibbard's Case. On the first such alternative, we might try to show that we each ought to do B, and not A. Second, we might make claims not about how we ought to do B and not A in the usual sense, but rather about other things, such as about what merely ought to happen, or about our motives. I will start by considering these alternatives and showing how they fall short. I will then show that the appeal to collective reasons can avoid these problems and can thus provide a more satisfying account of the sense in which we should both do B, and not A.

Let's start by considering noncollectivist accounts of the first type. On these accounts, our sense that "we should both do B, and not A" should be understood not as a claim about what we as a group should do, but rather merely as equivalent to the conjunction of the purely individualistic claims that I ought to do B, and not A, and that you ought to do B, and not A.

How could we explain why each of us should do B, and not A? Note that it would not be enough to point to the effects of what each of us would be doing. For suppose that we both do A. And suppose further that each of us is stubborn and would do A regardless of what the other does. In that event, we would each be making things as good as possible, given what the other is doing. So appealing to our individual effects would not explain why we should not do A. However, we might find some alternative explanation. For example, we might claim, with some rule consequentialists, that each of us ought to do what, if everyone did it, would have the best consequences.

However, even if we can find a plausible explanation of some such kind, all accounts according to which each of us ought to do B, and not A, are open to the following objection. Suppose that you will do A. In that event, my doing B would make the outcome extremely bad. So it seems clear that I should not do B. But the accounts we are considering claim that "we should do B, and not A" in the sense that each of us should do B, and not A. These accounts therefore claim that I should do B, even though this would lead to disaster. That is unacceptable.

Rather than making the general claim that each of us ought to do B and not A, we might instead merely claim that I ought to do B if you do B, but ought to do A if you do A. But again, this would not help to explain in what sense we should not both do A.

10. In using Gibbard's Case to support the idea of collective reasons to make outcomes better, I am following a suggestion that Parfit makes, but ultimately rejects, in "What We Together Do." I discuss two of his reasons for rejecting this idea below.

Let's turn, then, to accounts of the second kind, which make claims not about the rightness or wrongness of our doing A or B, but rather about other aspects of the case.

First, in some versions of this case, it might be true that we would do B just in case one of us suggests that we do so. As Parfit points out, in these versions of the case, each of us would then have acted in the wrong way by failing to make such a suggestion, in virtue of its consequences.¹¹

However, this account will not succeed in other versions of the case. If we are each determined to do A regardless of what the other says, then, given the other, none of our prior choices would make things worse.

Next, it might be claimed that if we are each determined to do A even if the other were willing to do B, then this reveals something bad about our motives or character.¹² And this, it might be suggested, can explain our feeling that if we both do A, something has gone wrong.

However, this seems inadequate. It seems obvious, not just that we should not be the sort of people who would do A, but also simply that we should not do A.¹³

Another alternative is to claim that when we say that "we ought to do B," we are using "ought" in what is called the evaluative sense, rather than the deliberative sense. Whereas the deliberative "ought" is closely tied to particular agents, we use "ought" in the evaluative sense to make more impersonal claims about what would ideally be the case, such as "there ought to be less misery."¹⁴ On this suggestion, in claiming that "we ought to do B," what we mean is that "it ought to be the case that we do B." This is roughly equivalent to claiming that it would be a good thing if our both doing B is what happens.

However, this too seems inadequate. It might seem sufficient to claim that it would be a good thing if our both doing B is what happens if this were out of our control. But again, I am assuming in this and other cases that we are capable of intentionally acting together. While it might be out of my control whether we do B, it is not out of our control. So it seems inadequate merely to think of our failing to do so as unfortunate.¹⁵

11. *Ibid.*, 21–22.

12. Parfit makes this suggestion, *ibid.*, 22.

13. This objection also applies to cooperative utilitarianism, a theory that Regan proposes in order to address Gibbard's Case (*Utilitarianism and Co-operation*, chap. 8). Very roughly, this theory requires each of us to cooperate, with whoever else is cooperating, in producing the best outcome. In order to "cooperate," we must not only act in the right ways but also have the right motivations. If in Gibbard's Case we both do A because we are both uncooperative, then we will have failed to satisfy Regan's theory, since we will have failed to have the right motivations. However, we will not have done anything wrong: when you are uncooperative, Regan would agree, A is the right thing for me to do, and vice versa.

14. See Mark Schroeder, "Ought, Agents, and Actions," *Philosophical Review* 120 (2011): 1–41.

15. As Frank Jackson writes of different examples, "It is evident that something wrong *happens*" in these cases, "but more than that is evident: something wrong is *done*. (It would

Finally, reconsider the collective reasons account. Again, this account claims that we together ought to carry out the course of action in which we each do B in virtue of the fact that we would together be making things better. And we together have decisive reasons not to do A, in virtue of the fact that we would together be making things worse than they could be.

These claims offer straightforward and what seem to me to be very plausible explanations of both in what sense we should do B, and not A, and why. In addition, they do not have counterintuitive implications. As I will argue later, the claim that we together ought to act in some way does not imply that each of us ought to do his part, no matter what. So we can accept that we together ought to do B without claiming that I ought to do B even if you do A. Finally, this account offers claims about what we have decisive reasons to do now, rather than merely about our previous acts, our motives, or what ought to happen.

I conclude, then, that Gibbard's Case provides significant evidence in favor of the idea that we have not only individual reasons to make outcomes better, but also collective reasons to do so.

Harm

Consider next our reasons not to harm other people in certain ways, such as by killing them.

It might seem that there is some prima facie reason to be skeptical that reasons of this sort would apply equally to groups. Unlike our reasons to bring about good outcomes, our reasons not to harm others might derive from the relations in which we would stand to other people if we acted in these ways. And there might be important differences between the ways in which I can relate to someone and the ways in which a group of people can relate to someone.

However, I will now argue that groups, and not only individuals, have reasons not to harm people. I will support this claim with the following case. Consider

Firing Squad: You and I each shoot some innocent victim, who then dies. Our bullets both contribute to this person's death. However, either shot, by itself, would have killed.¹⁶

Finally, suppose again that we are each stubborn and would shoot regardless of what the other does.

be quite wrong to think of either case as being one of a natural misfortune, like a flood)" ("Group Morality," 100).

16. This is based on a case Parfit describes in *Reasons and Persons*, sec. 26.

Now, because any one of our contributions would be enough, it is true of each of us that whether our victim is injured or killed does not depend on what we do. Nevertheless, in this case, it seems obvious that, in some sense, we are doing something seriously wrong.

On the account I will defend, if we both shoot, we will together be doing something wrong, either in intentionally carrying out this course of action together, or in failing to stop ourselves from doing so individually. For we would together be killing this person. This fact gives us as a group a strong reason to avoid this course of action.

Again, it might be claimed that we can account for the wrongness of what we would be doing purely in terms of our individual duties and reasons for action. I will consider two possible accounts of this kind. I will then argue, again, that the collective reasons account enjoys significant advantages over these alternatives.

One purely individualistic account could go as follows. First, even if our victim's death does not depend on what either one of us does, it could be argued, at least one of us must still have killed this person—that is, at least one of us must in fact have caused this person's death. Other things equal, we might believe, it is wrong to kill someone, even if we do not hasten this person's death. So, we might claim, if I turn out to be the killer, then this makes what I have done wrong. And you may have acted wrongly insofar as you risked being the killer. Or we might each have caused this person's death, and so we might each be killers.

However, even if killing that does not hasten death is objectionable, it seems less serious than the other kind. Suppose that

two hospital patients are dying of some disease. There is only one drug that would have any effect. It would cure this disease, but it would also kill these patients painlessly at the exact moment that the disease would have done so. However, the drug would also make their organs safe for transplantation. While the patients have no objection to donating their organs, they refuse to take this drug. But if their doctor gives them this drug anyway, she could save someone else's life.

We might normally find killing unacceptable, even when it would prevent other deaths. But it at least seems open to question whether the doctor should give her patients this drug. So it seems that our reasons not to perform even two killings that do not hasten death must be weaker than our reasons not to perform a normal killing. And if this is so, then even if, in *Firing Squad*, we both turn out to be killers, the fact that we are committing two killings that do not hasten death does not seem to be enough to capture what wrong is being done. Again, what we are doing seems to be seriously wrong.

It might next be suggested that even if we want to explain this case solely in terms of individuals' reasons for action, we do not have to ground these reasons in facts that are only, in a narrow sense, about these individuals. Parfit uses a firing-squad case to support the following claim:

Even if an act harms no one, this act may be wrong because it is one of a *set* of acts that *together* harm other people.¹⁷

Parfit's proposal does not explicitly make claims about what groups ought to do. This proposal thus allows us to appeal to facts about what we together do without relying on the unfamiliar notion of collective reasons for action. In this way, we can think of this proposal as occupying a middle ground between a purely individualistic account of Firing Squad, and a collective-reasons account of the kind I will describe below.

However, it is unclear whether this middle ground is tenable. For it is unclear whether this proposal, unaided by an appeal to collective reasons not to harm, can offer a convincing explanation of why it would be wrong for each of us to shoot. On this proposal, if we both shoot, then:

- (1) I will be acting wrongly, because we will together be harming others.

If we accepted the claim that we have collective reasons not to harm others, we could expand this to become:

- (2) I will be acting wrongly, because we will together be harming others, and so will together be acting wrongly.

But if we reject (2), then (1) seems undermotivated. It seems unclear why my action should be criticizable in virtue of its part in some larger activity, if there were nothing wrong with that larger activity itself.

It might be replied that we can criticize our larger set of acts simply by saying that it is a bad thing that these acts are performed. On this suggestion, the wrongness of my shooting derives from the fact that what I did would be a part of something bad that happens, rather than of something wrong that is done.

However, this seems implausible. Suppose that

I suffer from a genetic condition that causes my body to produce too much of certain chemicals. My favorite food also contains these chemicals, and if I ate this food, the chemicals from it would combine with the chemicals from my body to produce painful symptoms. But because my body already produces more

17. *Ibid.*, sec. 26.

than enough of this chemical to cause these symptoms, eating this food would not make me feel any worse.

It seems very implausible that these facts could make it wrong for you to offer me this food.

This case suggests that the harms that we produce together with other people must be criticizable in some way that does not apply to the harms we produce together with natural or inanimate causes. And it seems to me that the most natural criticism we could make is that if we together harm other people, we will together be acting wrongly.

This leads us to the collective-reasons account of Firing Squad. Again, on this account, we together have a strong reason to avoid shooting. For if we both shoot, we will together be killing our victim. (This is true whether or not this killing will count as an intentional collective action.) Moreover, our combined actions will hasten this person's death.

Since our combined actions do hasten our victim's death, this account seems capable of capturing what is seriously wrong about what we are doing. In addition, unlike the Parfit proposal, this account offers a straightforward and plausible explanation of why what we are doing is wrong.

I conclude, then, that Firing Squad provides us with significant evidence in favor of the idea that we have reasons not to harm people in certain ways, not only as individuals, but also as groups.

Self-Interest

Consider finally our reasons to benefit ourselves. Here parity between individual and collective reasons might seem especially unlikely. Cases in which I benefit myself, and cases in which we together benefit ourselves, differ with respect not only to the agent but also to the beneficiary. And if groups can be said to have a welfare at all, it might be thought, this must be a very different sort of thing from the welfare of an individual.

Despite these differences, I will now argue that groups, as well as individuals, may have self-interested reasons. This claim, I will argue, is supported by the case of the *Prisoner's Dilemma*, where we must each choose whether to benefit the other at some lesser cost to ourselves. This case has the following structure:

		You	
		benefit me	harm me
I	benefit you	Second-best for each	Best for you, worst for me
	harm you	Best for me, worst for you	Third-best for each

Suppose finally that, unlike in Firing Squad, the relevant harms here would not involve killings or other kinds of actions thought to be specially prohibited.

Discussions of the Prisoner's Dilemma often assume that we cannot currently communicate and must each decide whether to harm the other. But suppose instead that we can communicate and agree to harm each other. This may seem, in some sense, crazy, at least to some of us. It may seem that, in some sense, we must have some reason not to harm each other, given that there is another course of action that would be better for both of us. If this is right, then how can we explain in what sense this is true, and why?

Note first that we could not explain why we have some reason not to harm each other in virtue of the reasons that we each have to do what is in our own interest. Whatever you do, benefitting you would make me worse off.

We might next claim that we each have altruistic reasons not to harm the other. If my acting in some way would make someone worse off, we might claim, this fact in itself gives me a reason not to act in this way.

This suggestion is very plausible. However, it does not seem to be enough to capture what is going wrong. What we are doing seems to be not merely unfriendly, but self-destructive.

Next, as in the discussion of Gibbard's Case, we might focus on prior acts, or claim that what happens is merely unfortunate, or reveals something defective about our characters or motives. However, we might think that in choosing to harm each other rather than benefitting each other, what we are doing is itself wrong. And as before, these suggestions do not support the idea that we are making the wrong choice.

Finally, we could claim that we together have a reason of some kind not to carry out the course of action in which we harm each other, since this would be worse for both of us. Such an account might offer a straightforward and plausible way to make sense of the idea that we have some reason not to harm each other, and why. In addition, such an account might make sense of the idea that we have some reason not to harm each other not only for altruistic reasons, but also because doing so would be self-destructive.

Now, to develop this proposal, it might be suggested that the same ultimate self-interested reasons apply to groups and individuals alike. In other words, it might be suggested, our underlying principle should be that each agent has a reason to do what benefits that agent. On this view, our collective reasons to benefit ourselves would be reasons to benefit us as a group.

However, philosophical theories of group welfare have not yet, to my knowledge, been developed as thoroughly as theories of individual

welfare. And again, given the important differences between individuals and groups, it might turn out that the welfare of a group is a very different sort of thing from the welfare of an individual. So until theories of group welfare are more fully developed, it seems premature to claim that groups and individuals have the same fundamental self-interested reasons.

But even if groups do not have the same reason to benefit themselves that individuals have, this does not mean that they do not have self-interested reasons at all. Instead, I propose that we can take the Prisoner's Dilemma to suggest that a similar reason applies to them. In particular, I propose, we can simply claim that when acting in some way would be better for all of us, this fact gives us a collective reason to act in this way.

Conclusion

I have now argued that there is significant evidence to support the idea that groups, like individuals, have reasons to make outcomes better, to benefit other people, to avoid harming other people, and to benefit themselves. This is strongly suggested, I have argued, by Gibbard's Case, by the Firing Squad, and by the Prisoner's Dilemma. Thus, I conclude, we have good grounds to think that to a significant, though limited, degree, the kinds of reasons that apply to groups are the same as those that apply to individuals.

IV. WHEN SHOULD I PARTICIPATE?

Introduction

So much for the question of what groups have reasons to do. We can now turn to our second question: If we together have reasons to act in certain ways, how, if at all, does this affect what each of us ought to do? Do individuals have any *group-based* reasons for action: reasons to do our parts in what we together ought to do, deriving from the group's reasons to perform this collective act?¹⁸

In what follows, I will first describe one intuitively appealing principle of group-based reasons and show that this principle faces several significant objections. These objections suggest that our principle should not apply in noncooperative contexts, should allow for the possibility that members have independent reasons for action, and should not allow for an objectionable kind of double-counting. I will then propose a principle

18. I take the phrase "group-based reasons" from Christopher Woodard, whom I discuss below. However, Woodard uses this phrase to refer to reasons an individual might have to do her part in actions that a group containing her could perform, in virtue of the good or bad consequences of those actions, rather than in virtue of the group's reasons for action, which might not all derive from consequences (*Reasons, Patterns, and Cooperation*, chap. 2, sec. 2).

that meets these conditions. Finally, I will list some remaining questions about when we have reasons to do our part.

Just as I do not claim to have offered a full account of collective reasons, I do not claim that the principle I propose represents a complete account of group-based reasons. My goal is only to make progress.

The Simple Principle

Here is an initial proposal about how what we together ought to do can influence what each of us should do. According to

the Simple Principle: If we together ought to act in some way, then each of us ought to do our part. And if we together ought not act in some way, then each of us ought not do our part.

This is a natural answer. As I will now show, however, it is open to significant objections. By revising this principle so as to avoid these objections, we can try to get closer to the truth.

What If Others Don't Cooperate?

First, note that Simple Principle tells me to do my part even when others will not do theirs. This makes this principle vulnerable to two objections.

First, if others won't cooperate, my contribution can seem pointless. This is especially apparent in cases where my own contribution is some mundane task with no special significance apart from the role that it would play in the group activity. For example, suppose that we ought to paint a house together, and my job is to mix the paints.¹⁹ The Simple Principle implies that, even if you abandon the project, and I can't do it by myself, I still ought to mix the paints. But it seems pointless for me to spend my time doing this.

Second, if others won't cooperate, my contribution could lead to disaster. Suppose again that we are in Gibbard's Case and ought to both do B, thereby producing the best outcome. The Simple Principle implies that, even if you will do A, I still ought to do B. But again, doing this would make the outcome very bad.

Christopher Woodard suggests that when others will not cooperate, while it may not always be true that we ought to do our part anyway, we may still have some reason to do so.²⁰ Thus, on

Woodard's Principle: Even when others will not do their parts in what we together ought to do, I may still have some group-based reason to do my part.

19. Michael Bratman discusses a similar example in "Shared Cooperative Activity," *Philosophical Review* 101 (1992): 327–41.

20. Woodard, *Reasons, Patterns, and Cooperation*.

While Woodard claims that the uncooperativeness of others does not prevent me from having some reason to do my part, he leaves open the possibility that there may be other constraints on when we have group-based reasons.²¹ However, he suggests that we do have group-based reasons in many actual noncooperative contexts.

Woodard argues that the idea that we have such reasons even in noncooperative contexts is the best explanation of the intuitive tension we feel between being “principled” and being “pragmatic.” For example, in Bernard Williams’s famous *Jim and the Indians* case, Woodard writes, Jim seems to have both reasons to accept Pedro’s offer to shoot one of his Indian captives, so that Pedro will let the other nineteen go, and reasons to decline the offer, and not shoot.²² These latter reasons, Woodard suggests, are the reasons that Jim has to do his part in the best group course of action, in which neither Jim nor Pedro shoots anyone. And Jim still seems to have these reasons even though Pedro is unwilling to cooperate in that course of action.

However, Woodard’s Principle faces two objections.

First, this principle implies that even if you abandon our house-painting project, I still have some reason to mix the paints, unless this situation would fail to satisfy some other constraint on when we have group-based reasons. Therefore, this principle implies that, if there is no independent constraint that applies here, and if I have nothing better to do, then this is what I ought to do. But again, it seems that I have no reason to do this.

Woodard’s Principle seems more compelling when my part consists in not killing anyone than when it consists in mixing paints. But this, I suggest, may be because the idea that I always have some reason not to kill anyone is independently plausible and supported by other plausible theories.²³ In contrast, it is not plausible that I always have some reason to mix paints.

Now, to avoid this objection, Woodard might restrict his principle to cases where our group action is morally required, rather than merely a useful project.²⁴ However, as I will now argue, even if he does make this restriction, Woodard’s proposal is still vulnerable to a second objection: the most natural ways to develop this proposal still imply that we should act in ways that lead to disaster, and it is not clear whether Woodard could avoid this implication in a principled way.

21. *Ibid.*, chap. 5, sec. 5.

22. Bernard Williams introduces this case in “A Critique of Utilitarianism,” in *Utilitarianism: For and Against*, ed. J. J. C. Smart and Bernard Williams (New York: Cambridge University Press, 1973), sec. 3.

23. Woodard offers forceful objections to several of these alternative theories. While I do not think these objections are decisive, it would take us too far afield to discuss them here.

24. I owe this suggestion to an anonymous reviewer.

To start, if I can have some reason to do my part even in noncooperative contexts, how strong would this reason be? It seems that the reason would have to be proportional to the strength of our collective reason to cooperate. If I have a group-based reason to mix paints, it seems this must be weaker than Jim’s group-based reason not to shoot, since the stakes are so much lower.²⁵ If this is right, then we might claim:

The strength of my reason to do my part in noncooperative contexts is some proportion of the strength of our collective reason to cooperate.

But even if this proportion is very small, this principle, together with other plausible normative views, can imply that I ought to do my part even when this would lead to disaster.

For example, suppose that we face another version of Gibbard’s Case. In this version, a million lives are in danger, and we have the following options:

		You	
		do A	do B
I	do A	One hundred saved	All die
	do B	All die	All saved

It seems plausible that, if we are in a position to save lives, either together or individually, then other things equal, the strength of our reasons to do so will be proportional to the number of lives that would be saved. Let’s say, then, that if we could save n lives, we have a reason of strength n to do so. If so, then we together have a reason of strength 1,000,000 to both do B. Now suppose that my reason to do my part in noncooperative contexts is only one thousandth as strong as our collective reason to act. Then if you do A, I will still have a group-based reason of strength 1,000 to do B. And since I only have a personal reason of strength 100 to do A, this means that I ought to do B, even though I will be letting one hundred people die, and saving no one.

Now, we might be skeptical that the strength of reasons can be quantified so precisely. But we can also state the problem without relying on precise numbers. There are some versions of Gibbard’s Case, we can claim, in which our reason to cooperate will be extremely strong, but in which you refuse to cooperate. So if the strength of my reason to do my

25. I should emphasize that Woodard does not himself make this claim.

part is some proportion of the strength of our collective reason, however small, this reason could still be very, very strong. It could thus be strong enough even to outweigh the personal reason I have against doing my part, deriving from the fact that this would lead to disaster.

To avoid this implication, Woodard might set an upper bound on the strength of our reasons to do our part in noncooperative contexts.²⁶ However, this seems ad hoc. In addition, where exactly should we set this bound? Setting the bound at any particular strength seems arbitrary. Therefore, while the idea that we have some reason to do our part even when others won't cooperate does not entail that we should act in ways that lead to disaster, it is not clear whether there is any principled way of avoiding this conclusion.

Finally, Woodard might again suggest that there may be some other constraint on when we have group-based reasons which applies to cases like the one described above. If some such constraint applies here, Woodard might claim, then I may have no reason to do my part. For example, Woodard might plausibly claim that I have no reason to do my part if I would thereby let one hundred people die. Woodard could then adopt some nonarbitrary proposal about the strength of the group-based reasons that we do have, in cases where no such constraint applies.

However, recall that Woodard's main argument in favor of the idea that we can have group-based reasons even in noncooperative contexts focused on the case of Jim and the Indians. Our intuitions about this case, Woodard argues, are best explained by the hypothesis that Jim has some reason to do his part in the course of action in which neither he nor Pedro shoots anyone. This is true, Woodard claims, despite the fact that if Jim does his part and refrains from shooting one of the captives, Pedro will kill them all. But if Jim has some reason to do his part even when he would be letting nineteen people die, it seems that the fact that I would be letting one hundred people die could not prevent me from having a reason to do my part in Gibbard's Case. Nor does it seem that the lives that we would be saving if we both did B could prevent me from having a reason to do my part. So it seems that there is no plausible constraint on group-based reasons that applies to this version of Gibbard's Case but does not apply to Jim and the Indians.

Thus, both the Simple Principle and the weaker view that we have some reason to do our parts even in noncooperative contexts can imply that we ought to act in ways that seem pointless or that would lead to disaster, and it is unclear if there is any principled way to avoid this. So we should restrict our principle to cases where others are willing to cooperate. Frank Jackson proposes one such principle. According to Jackson,

26. Again, I should emphasize that Woodard does not himself make this claim.

If a group act is right, *and it is in fact performed*, then each individual constituent act is right.²⁷

Or, equivalently, according to what we can call

Jackson's Principle. If we together ought to act in some way, and we will act in this way if I do my part, then I ought to do my part.

We might also add that if we together ought not act in some way, but will act in this way if I do my part, then I should not do my part.

Independent Reasons

Again, the Simple Principle claims that if we together ought to act in some way, then each of us ought to do our part. As we just saw, both this view and Christopher Woodard's weaker principle face objections arising from the possibility that others might not do their parts. Jackson's Principle avoids these objections, since it is restricted to cooperative contexts. But I will now argue that this principle faces another objection, which can be brought out in the following case.

Suppose that

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.

Here, we might think, I do have a special reason to save my child, stronger than my reasons to help save the strangers. So I ought to save my child. But we might think that we as a group don't have any special reason to save my child. After all, my reason for acting might derive from the fact that this is my child, and this fact might not apply to the group. So, we might think, we as a group only have reasons to do whatever would save the most lives, and so ought to save the two strangers. Jackson's Principle seems to imply, then, that I ought to do my part in our rescue mission, even though I would be letting my child die.

As this case illustrates, since Jackson's Principle claims that I always ought do my part in reasonable group acts, this view ignores the possibility that I may have other, stronger reasons for action that are not shared by the group.²⁸

Now, it might be suggested that what a group has reasons to do can in fact be affected by special individual reasons of this kind. For exam-

27. Jackson, "Group Morality," 107.

28. I owe this point to Jacob Ross.

ple, it might be claimed that a group has a special reason to save the child of any one of its members, just as strong as the member's individual reason.²⁹ More generally, it might be claimed that even if I have a reason to bring about some outcome in virtue of some fact that does not apply to a group to which I belong, this might always provide the group with a reason to bring about this outcome. And if there is this sort of "upward" transmission of reasons from members to groups, it might be argued, then cases like the one above do not in fact raise a problem for Jackson's Principle, since special individual reasons will already be reflected in the group's reasons.

However, while suggestions of this kind are worth considering, they also seem open to question. For example, while I may have strong reasons to make sure that I keep my promises, it is not clear that any group to which I belong has equally strong reasons to make sure that I keep my promises. In addition, "upward transmission" seems especially unlikely in groups whose members have no prior relationships with each other. But recall that the examples that I used to argue for collective reasons did not rely on any details about the prior relationships between group members. If these arguments were sound, then, it seems unlikely that upward transmission is a feature of all groups that have collective reasons for action.

Thus, our principle of group-based reasons should at least leave open the possibility that individuals have independent reasons for action that are not shared by the group, or that are otherwise reflected in the group's reasons. We can do this by switching to a weaker principle about what we merely have some reason to do, rather than about what we ought to do, all things considered. I therefore propose a fourth principle of group-based reasons,

P4: If some fact about our acting in some way gives us a reason (not) to act in this way, and we will act in this way if I do my part, then I have a reason (not) to do my part.

Double-Counting

I have argued that our principle of group-based reasons should apply only in cooperative contexts, and that this principle should be stated as a claim about reasons, so as to allow for independent individual reasons to act in other ways. Therefore, I argued, we should reject both the Simple Principle and Woodard's and Jackson's alternative principles, and instead consider accepting P4. However, I will now argue that P4 faces another objection and thus needs to be qualified with one additional claim.

29. I owe this suggestion to Julia Nefsky and to an anonymous reviewer.

The objection can be illustrated with another rescue case. Suppose that

two people are in imminent danger. I could help you to save one of these people or could save the other on my own. But there is not enough time to save both.

Suppose we believe that any person or group has some reason to act in such a way that a life is saved. Then the fact that a life will be saved if I help you gives me some reason to do so. But we will also be together acting in such a way that a life is saved. So P4 implies that I have not only a personal reason to help you, but also a group-based reason to help, because I would thereby be doing my part in our rescue mission. But it seems implausible that I could have stronger reasons to do my part in our rescue mission than to save the other person's life, merely on the grounds that our rescue mission would involve cooperation. I shouldn't have extra reason to do what will bring about some outcome merely because I would thereby be playing a part in a group act that would achieve the very same thing. So in this way, P4 seems to double-count our reasons for action.

This may seem to be a minor technical problem. But as I will argue later, the need to avoid double-counting has important implications.

To avoid this objection, I propose that we combine P4 with the following proviso:

If by acting in some way I would be doing my part in multiple acts, then the personal and group-based reasons I have deriving from any given feature of the outcome will not add up.

For example, by helping you to rescue someone, I am acting in such a way that this person's life is saved. And by carrying out our rescue, we are acting in such a way that this person's life is saved. But the reasons given to me by these acts should not add up.³⁰

In claiming that our group-based reasons should not add together with my personal reasons in such cases, I am not claiming that we do not have group-based reasons in these cases. Rather, we can think of these two kinds of reasons as overlapping.³¹ To justify our actions, we could appeal either to our individual reasons, or to our group-based reasons, or to both. We just shouldn't claim that these reasons combine to make this justification extra strong.

30. Mark Schroeder has also argued that there are other kinds of reasons which should not add up, because they are not independent in the right way. See *Slaves of the Passions* (New York: Oxford University Press, 2007), sec. 7.1. It seems plausible that two reasons given by the very same feature of an outcome might similarly not be independent enough to add up.

31. I owe this image to Abelard Podgorski.

Remaining Questions

P4, I have argued, avoids implausible implications in noncooperative contexts, allows for members to have independent reasons to act in other ways, and can avoid double-counting. I therefore offer this principle for your consideration.³²

There are at least three other important sets of questions about our group-based reasons which I will leave open.

First, how strong are our group-based reasons? Can we compare the strength of a collective reason with the strength of an individual reason? If so, is the strength of my group-based reason some function of the strength of the collective reason from which it derives? Is my reason as strong as the group's reason? Or is the strength of the group's reason divided among the group's members?

Second, do group-based reasons depend on whether the group will intentionally act in the relevant way? If we have not agreed to act together, but the others will just happen to "do their parts" for their own reasons, could this be sufficient for me to have a group-based reason to do my part?

Third, are any of our collective reasons essentially *contrastive*? That is, are there any collective reasons that are essentially reasons for the group to act in one way rather than another?³³ If so, how, if at all, do these transmit to the group's members?

These questions represent promising topics for further reflection. But for now, I will turn to the implications of P4 and my earlier claims about where there is parity between individual and group reasons.

V. IMPLICATIONS

Introduction

I earlier argued that there was limited, but significant, parity between individual and group reasons. In particular, I argued, if we have individual reasons to make outcomes better, to avoid harming people, and to benefit ourselves, then we also have collective reasons to do so. And in the previous section, I proposed

32. P4 requires at least one further qualification: when determining the group's reasons for the purpose of determining my group-based reasons bearing on some set of options, we should not take into account which of these options I will or would choose. This claim is needed in order to prevent what I ought to do from being indirectly affected by what I will or would in fact do. Also, Jacob Ross has pointed out in correspondence that, without this claim, there are some cases in which P4 could imply that I have unfulfillable obligations.

33. Compare Justin Snedegar, "Contrastive Reasons and Promotion," *Ethics* 125 (2014): 39–63.

P4: If some fact about our acting in some way gives us a reason (not) to act in this way, and we will act in this way if I do my part, then I have a reason (not) to do my part,

together with the proviso that

If by acting in some way I would be doing my part in multiple acts, then the personal and group-based reasons I have deriving from any given feature of the outcome will not add up.

Here, I will discuss some of the implications of these claims.

Again, as I will explain, one natural aspiration we might have for a theory of collective reasons is to help support the idea that we should contribute to good causes even when our own contribution would not make a difference. However, I will argue that our general collective reasons to do good might never affect what I ought to do. Still, I will argue, this does not mean that these reasons are unimportant. Finally, I will show that unlike our general reasons to do good, other kinds of collective reasons can have implications for what individuals ought to do.

Do Our Reasons to Do Good Ever Affect What I Ought to Do?

It might be thought that if we accept the idea that groups have reasons for action, and that individual members have reasons to do their parts, this must have implications for how individuals should act in one widespread and highly important type of case: cases in which we can together do good, or make things better in some way. After all, in many such cases, people are often troubled by the thought that their own contribution might not make any difference.³⁴ But it might be suggested that we can address this concern by appealing to the effects of what we together do.³⁵ As I have argued, we have reasons to do what would make outcomes better not only as individuals, but also as a group. While my contribution might make no difference, I might be participating in a joint act that makes a big difference. And if we have reasons to act in some way, I have argued, then this can give me a reason to do my part.

However, I will now argue that if my proposed principle of group-based reasons is correct, then at least when others are willing to agree to certain courses of action, our general collective reasons to make outcomes better will in fact never affect what I ought to do, whether or not

34. See Jonathan Glover and M. J. Scott-Taggart, "It Makes No Difference Whether or Not I Do It," *Proceedings of the Aristotelian Society, Supplementary Volumes* 49 (1975): 171–209; Shelly Kagan, "Do I Make a Difference?" *Philosophy and Public Affairs* 39 (2011): 105–141; and Julia Nefsky, "Consequentialism and the Problem of Collective Harm: A Reply to Kagan," *Philosophy and Public Affairs* 39 (2011): 364–95.

35. Kagan mentions such a suggestion in "Do I Make a Difference?," 112–13.

my contribution would make a difference. In contrast, I will argue in the next section that other kinds of collective reasons can affect what I ought to do. These arguments rely on a distinction between what I will call “agent- and act-neutral” reasons, which include our reasons to do good, from “agent- and act-relative” reasons. Whereas our collective agent- and act-neutral reasons will not affect what I ought to do, I will argue, our agent- and act-relative reasons will.

I will start by explaining what I mean by the distinction between these different kinds of reasons. Whenever someone has a reason to act in some way, I will assume, we can speak of this reason as deriving from some feature of the “outcome” of that act, very broadly construed. For example, an agent might have a reason to act in such a way that someone’s life is saved, that that very agent is benefitted, or that that act is the fulfillment of a promise. A reason is *agent-relative*, in my sense, when it derives from a feature that involves the agent as such, and can be specified only using terms like “that agent”; otherwise, it is *agent-neutral*. For example, an agent’s reason to act in such a way that that very agent is benefitted would be agent-relative, whereas a reason to act in such a way that someone’s life is saved would be agent-neutral. A reason is *act-relative*, in my sense, when it derives from a feature of the outcome of an act that involves that very act as such and can be specified only using terms like “that act”; otherwise, it is *act-neutral*. For example, a reason to act in such a way that that act is the fulfillment of a promise would be act-relative, whereas the other two examples would be act-neutral.³⁶

In what follows, I will assume that our general collective reasons to make outcomes better are agent- and act-neutral, such as reasons to act in such a way that someone’s life is saved. (We might also have more specific agent- or act-relative collective reasons to do good, such as reasons to benefit people to whom our group has special ties.) With this assumption in place, we can now explore how such reasons might bear on what individual members of the group ought to do.

Consider any feature of the outcome of our act which might provide us with a general collective reason to make outcomes better. Now, this feature will either depend on what I do, or it will not.

Suppose first that some such feature does depend on what I do. For example, suppose that someone’s life will be saved if we together undertake some rescue mission, and that this person’s life will not be saved if I fail to do my part. Now, if we act in this way, then, since this feature is

36. It is important to note that reasons can be act-neutral, in my sense, even when they derive from features that refer to act types. For example, a reason to act in such a way that lying is minimized would be act-neutral, since this feature does not refer to the very act under consideration. In contrast, a reason not to act in such a way that this act would be a lie would be act-relative.

agent- and act-neutral, it can be described not only as a feature of the outcome of our act, but also as a feature of the outcome of my own act.

However, as I have been assuming, and as many people believe, we already have our own individual reasons to do good. If so, then this feature will also give me an independent individual reason to act in this way. In that case, since my group-based reason to do my part and my personal reason to act in this way derive from the very same feature of the outcome, our proviso against double-counting will prevent these reasons from adding up.

Suppose next that some feature of the outcome of our act does not depend on what I do. For example, suppose that someone's life will be saved regardless of whether I do my part in a joint rescue mission, or act in any other way. In that case, we could defensibly claim that this feature will not already have given me a personal reason to act in the relevant way, since what I do will not affect this feature. So this feature will give me a nonredundant group-based reason to do my part in our rescue mission.

However, if this feature does not depend on what I do, then it will also be a feature of the outcome of what we will be doing if I act in any other way. So this feature will also give me an equally strong reason to do my part in these alternative courses of action, so long as others are willing to agree to these courses of action. (Note that such an agreement would not change what any of them would be doing, but would merely determine whether our activity will count as an intentional group act.) But this means that I will have a new group-based reason for each of the options open to me. Since these reasons will cancel each other out, they will not make a difference to which of these options I ought to do.

Thus, our general collective reasons to make outcomes better will never make a difference to what I ought to do, as long as others are agreeable.

Note that if it turns out that my reasons to do my part in making outcomes better do not depend on whether our group act would be intentional, then we could extend this argument even to cases where others are not willing to agree to alternative courses of action. In this case, we could conclude that our general collective reasons to do good will never make a difference to what I ought to do, period.

For many people, such a conclusion might seem to render everything that I have said irrelevant. Many people accept

Act consequentialism: Each person always ought to do whatever makes things go best.

I have argued that groups, as well as individuals, have reasons to do what makes outcomes better. Some act consequentialists might accept this claim. And since many act consequentialists also believe that our only

reasons for action are reasons to do what makes outcomes better, these people might adopt the stronger claim that

Each group always ought to do whatever makes things go best.

They could thus become *collective consequentialists*. But if it's true that our collective reasons to make outcomes better can never make a difference to what I ought to do, then these people would not need to revise their claims about what each person ought to do. And since my claims would then have no implications for what each of us ought to do, they might take these claims to have no practical significance.³⁷

However, even if my claims have no implications for what individuals ought to do, they will still have implications for what groups ought to do. And I can't see any reason to claim that what groups ought to do is important only insofar as it affects what individuals ought to do.

In addition, even if accepting collective consequentialism would not require us to revise act consequentialism, whether we follow this theory could make a big difference to how much good we will do. Once again, recall Gibbard's Case. Such cases may be very common. If we act in accordance with collective consequentialism, then we will do B in these cases, and together produce the best outcome. We will then also be satisfying act consequentialism, since we will each be doing the most good we can, given what the other is doing. But if we only want to satisfy act consequentialism, then we could instead both do A, even though this will only produce the second-best outcome. For again, if we both do A, it will also be true that each of us will be doing the most good we can, given what the other is doing. Therefore, if we only follow act consequentialism, we might make things much worse than they could be.³⁸

Agent- and Act-Relative Reasons

I have argued that our collective reasons to do good will never make a difference to what I ought to do, at least as long as others are agreeable. Even so, I claimed, it is important to recognize these reasons. However, I will now argue that there can also be cases in which collective reasons can make a difference to what I ought to do. In particular, I will argue that our agent- and act-relative collective reasons, unlike our agent- and act-neutral reasons to do good, can affect what individuals ought to do.

As I argued, when I have reasons to do my part in virtue of some good feature of the outcome that we would together produce, and when this

37. As Julia Nefsky writes, "What does the claim that the *group did wrong* amount to? If it does not say that any individual ought, or even had reason, to have acted otherwise, then it doesn't seem to be a normative claim at all" ("The Morality of Collective Harm" [PhD diss., University of California, Berkeley, 2012], 7–8).

38. Donald Regan makes this point in *Utilitarianism and Co-operation*, chap. 2.

feature will not be produced regardless of what I do, then I may already have personal reasons to produce the very same feature, and so the proviso against double-counting will prevent these reasons from adding up. This is because these reasons derive from the very same feature or features of the outcome.

However, some features of an outcome cannot be produced by a group and by an individual at the same time. Again, some features of the outcome of an act involve the agent, or the act, as such. For example, if we together lie, we will have acted in such a way that this very group agent will have lied, and that this very group act was a lie. But I could never act in such a way that this very group agent will have lied, or that this very group act was a lie, since I am not the same as the group, and since my acts are never the same as the group's acts.³⁹

Therefore, if groups have any reasons for action deriving from features of the outcome that involve the agent or the act as such, then these features will never apply to individuals. And since the proviso against double-counting only applies to reasons which derive from the very same feature of the outcome, this proviso will not prevent my reasons to participate deriving from a group's agent- or act-relative reasons from making a difference to what I ought to do.

For example, recall Firing Squad. As I argued earlier, it is hard to explain on purely individualistic grounds why, in participating in a firing squad, we are doing anything seriously wrong. After all, neither of us is hastening our victim's death. So this case, I argued, supports the idea that we have collective reasons not to kill. And if we agree to shoot our victim, and you will do your part, then my principle of group-based reasons implies that I will have a group-based reason not to shoot. And since this reason derives from a fact about our group act, and not my individual part, it will not be neutralized by the double-counting proviso. Therefore, the appeal to collective reasons might also help us to explain why we would each be doing something seriously wrong.

Next, suppose again that we are in a Prisoner's Dilemma. This case, I argued, supports the idea that we have collective reasons to do what would be better for each of us. If we agree to benefit each other, and you will do your part, then my principle of group-based reasons implies that I will have a group-based reason to benefit you. And since this reason derives from a fact about us, the members of the very group which is acting, it will not be neutralized by the proviso.

Now, as I mentioned earlier, it is very plausible that I already have some reason to act in this way merely in virtue of the fact that it would

39. To be clear, I am not claiming that an individual cannot affect what a group does. I mean only that reasons to perform a particular act deriving from features involving the agent or the act as such could not apply to two different agents at the same time.

make you better off. But many people also believe that our purely altruistic reasons to help other people are weaker than our reasons to take care of ourselves. Therefore, if the net cost to me of benefitting you would be significant, my altruistic reasons might not be enough to outweigh my own self-interested reasons. However, once we add my group-based reason to do my part in a course of action that would be better for both of us, this might tip the scales. It might be because of this group-based reason that I ought to benefit you, even though doing so would be worse for me.

VI. CONCLUSION

I have argued that there are things that we together have reasons to do. In particular, I have argued, groups of people have reasons of some of the same kinds as those which individuals have. As Gibbard's Case, the Firing Squad, and the Prisoner's Dilemma suggest, groups as well as individuals have reasons to make outcomes better, to avoid harming others, and to benefit themselves.

I next discussed what this means for what we each ought to do. I proposed that if some fact about our acting in some way gives us a reason to act in this way, and the others will do their parts, then I also have a reason to do my part. This view, I argued, can avoid unacceptable implications in noncooperative contexts, can allow for independent individual reasons, and can avoid an objectionable kind of double-counting.

I then discussed what these claims imply. I first argued that our general collective reasons to make outcomes better might never affect what each of us ought to do, but that they matter nonetheless. Finally, I showed that other collective reasons might never affect what individuals ought to do.

Many philosophers have argued that we should care less about ourselves and those who are closest to us, and more about other people, including the distant poor, and about animals. But it is not enough to rethink the limits of our concern. We should also rethink the limits of the unit of agency from which we assess what to do.