

The Argument from Normative Autonomy for Collective Agents

Kirk Ludwig

1. Introduction

This paper is concerned with a recent, clever, and novel argument for the need for genuine collectives in our ontology of agents to accommodate the kinds of normative judgments we make about them. The argument appears in a new paper by David Copp, “On the Agency of Certain Collective Entities: An Argument from ‘Normative Autonomy’” (henceforth ACE),¹ and is developed in Copp’s paper for this special journal issue, “The Collective Moral Autonomy Thesis” (henceforth CAT).² The argument goes as follows:

- (i) We correctly assign blame (or obligations) to collectives in circumstances in which it would be a mistake to assign any (relevantly related) blame (or obligations) to their members.
- (ii) If (i), then collectives are genuine agents over and above their members.
- (iii) Therefore, collectives are genuine agents over and above their members.

Following Copp, I will call (i) the Collective Moral Autonomy Thesis (CMA). Copp argues for CMA primarily by appeal to cases, but also offers two general arguments for it. In the cases that Copp describes, we are to judge that a collective act is blameworthy, though each member of the group that acts is blameless because he is merely following procedures appropriate for his participation, or because there are excusing factors, or because of overriding personal duties.

I will argue that the case for CMA has not been made. In particular, I will argue that, in each case in which we feel inclined to hold a group responsible for something but not its members, it is because

- (i) we have accepted a false dilemma, that when no one agent is fully responsible for the action of a group of which he is a member, the only entity that could be responsible is the group as such, or
- (ii) we have directed our attention to the wrong individual or individuals, or
- (iii) we have become confused about the commitments of the individuals, or
- (iv) we have mistaken ameliorating for excusing factors, or

- (v) we have mistaken moral blameworthiness and all-in rational blameworthiness, or
- (vi) a combination of these things.

In section 2, I consider six cases that might be thought to support CMA. I begin with two cases which do not figure in Copp's argument, but which raise puzzles of their own, the diagnoses of which help us see the cases Copp appeals to more clearly. I conclude in section 3 by responding to Copp's two general arguments for CMA and offer some general considerations in favor of seeing the responsibility of groups as tracing back to the responsibility of individuals who are their members.³

2. Arguments from Cases

(1) Suppose that a group of people kill someone by beating him to death, though the victim dies from the accumulation of the several blows and not from the blow or blows of any of the individual members of the group. They bear responsibility for his death, but it might be argued that no *individual* who is a member of the group is responsible for *the death*, and that, thus, we must admit the group over and above its individual members as an agent in order to have an appropriate entity to hold responsible for the death. Individuals might be held responsible for the harms their blows caused, but since their blows did not cause the death, they may not be held responsible for it.

In this case, the thought that we must assign responsibility to the group per se arises from overlooking the option of apportioning responsibility among the members of the group in accordance with their roles. That is, in addition to the option of an individual being wholly responsible and the option of the group as such being responsible, there is the option of saying that as, and to the degree to which, each contributed to the death each bears responsibility for it. When we have assigned responsibilities in this way, there is no need for any further assignment of responsibility to the group over and above the individuals. This case shows that the full measure of responsibility which would rest on an individual if an act were carried out by him may fail to rest on *any* agent when the same thing is done by a group. When we say that *they* are jointly responsible for the death, we mean that they and only they share responsibility for it. The attribution of blame to the group is an abstract from the more specific facts about how responsibility is apportioned among its members. When we say that a group is to blame for something, our first question should always be: How then is the responsibility to be shared among the members of the group?

(2) The second case varies one discussed by Copp. A three-member tenure committee is charged by the university with deciding whether a candidate exhibits excellence in teaching, research, and service. Excellence in all is required for tenure. The committee has some leeway in deciding on its procedures for making

the decision. (This marks a difference from the case that Copp discusses. The point will become clear later.) The committee determines that they will proceed by voting. They canvass opinions at the outset. Committee members A and B judge a candidate to have achieved excellence in teaching, while C does not. A and C think the candidate achieved excellence in research, while B does not. B and C think the candidate achieved excellence in service, while A does not. They have two options. They can vote by category on whether the candidate has achieved excellence in each category. Alternatively, they can cast individual votes for or against the candidate on the basis of the judgment of each on whether the candidate has achieved excellence in all the categories. If they vote by category, they vote by a majority that the candidate has achieved excellence in teaching, research, and service, and so award tenure. However, if they each vote individually on whether the candidate displays excellence in all categories, they vote unanimously that the candidate has not achieved excellence in all three categories, and so deny tenure.

	A	B	C	Vote by category
Teaching	Yes	Yes	No	Yes
Research	Yes	No	Yes	Yes
Service	No	Yes	Yes	Yes
Individual votes	No	No	No	

Suppose they agree to vote by category, and let the majority decision determine whether the candidate has achieved excellence in each category. Then the committee decides to recommend tenure, even though it seems no member of the committee would individually have voted for tenure. Let us suppose that the decision was a mistake and that the candidate should not have been awarded tenure. Then the committee is to blame, it seems, for the mistaken decision. Yet, it might be argued, the members cannot be blamed because none were in favor of tenure. Indeed, the committee recommends tenure intentionally, though this goes against the best judgment of each of its members. If that is right, then it seems that the committee is an agent over and above the members, which has its own autonomous intentional states, and which can be held to blame when its members cannot.

The first charge here is that the committee is responsible for a bad decision though the members are not, because they all had the right view of the case. The second is that they all opposed granting tenure, and so talk of the committee intentionally recommending tenure is not reducible to talk of intentions of its members. The response to the second shows the correct response to the first.

The committee members intentionally adopt a decision procedure. The recommendation is to be decided in accordance with separate votes on teaching, research, and service. The committee members do not intend flatly that they recommend for or against tenure. Rather, each forms an intention that they shall recommend in accordance with the outcome of the decision procedure they adopt. They each have certain conditional intentions, for example, that they recommend

tenure if voting on the various matters independently is uniformly in favor of excellence. They then each form the intention that they recommend tenure or not upon the result of the separate votes so indicating, for the conditional intention becomes unconditional when the condition is met. This in turn shows how the responsibility of the committee for the decision is factored out into the responsibilities of individual members of the committee, for they each share equal responsibility for the decision procedure that they adopt and so for a bad outcome that results from following the procedure.⁴

(3) The third case involves the same committee. I now focus on the variant that Copp uses. In this example, crucially, the committee does not have the leeway to decide what decision procedure to adopt. The committee is required to vote on excellence in teaching, research, and service, and independently on tenure, by the university's rules, and award tenure in accordance with the *final* vote. Furthermore, they are required to announce at each stage the results of the votes, but they are not allowed to inform the candidate of how the individual members voted. Suppose, as before, a majority vote in favor of excellence in teaching, research, and service, but each votes against recommending tenure. The decision is made by the committee, and we may suppose that the committee's decision in the institutional context constitutes the university's decision—so that there are no additional steps or reviews which are part of the University's procedures which would allow a correction to be made in light of the conflict between the vote by individuals and by areas.

In this case, the individual members of the committee vote as they are supposed to, in accordance with their best judgment, in accordance with the university's rules, and accordingly, it seems, cannot be blamed for their votes. However, the decision is arguably morally problematic, since a majority held that the candidate achieved excellence in each of the areas the university required, and by announcing the results of each vote as it occurred, gave the candidate a reasonable expectation of a favorable final vote. The university, it seems, is blamable for this, though none of the members of the committee are blamable.

The primary question I want to raise is whether there are any other individuals who share responsibility for the decision who are members of the relevant collective, namely, the university, who are not members of the committee.⁵ Reflection on the contrast with the previous case will help here.

In the present case, in contrast to the previous one, the committee is not responsible for the rules that it follows and not in a position to modify them. Furthermore, the individuals who are members of the committee are obligated, in virtue of their agreeing to play their institutional roles, to follow the university's procedures. In the previous case, we did not blame the members because of a failure to carry out correctly the decision procedure they had adopted but for adopting a flawed decision procedure. In this case, the committee members did not adopt the decision procedure they were required to follow. However, if in the previous case the committee members were to blame because they chose a flawed decision procedure, in this case some responsibility for the untoward decision

rests with those responsible for the decision procedure. The university is ultimately responsible, but the rules did not arise by magic. Rather, some person or group within the university set the rules by which the committee was to proceed. Those responsible for the rules share some blame for the inappropriate decision in this case. *They* are not let off the hook even if *the committee members* are not individually to blame. So it does not follow from the individual members of the committee being free of blame that the university *over and above any individuals* is to blame. Some individuals who have roles in the institution share responsibility. Consequently, this case does not establish that the university is blameworthy for the decision though none of its members is.

In ACE, Copp anticipates a response of this sort and suggests that it is “possible that no individual was at fault for the university having the rules it has” because “no-one anticipated the kind of problem” in question (pp. 217–18). In CAT, he suggests further, “. . . that the university is an ancient one, that it has had these rules for two hundred years, and that it has never before faced a problem of the kind that arose in Borderline’s case. . . . Moreover, the issue is whether someone was blameworthy for the unfairness of the tenure decision, not whether someone was blameworthy for failing to change the rules” (CAT, p. 380).

That the university has had the rules for a long time would allow that those responsible for the rules in the first place are not now members of the institution. But it would not show that no individuals whose activities constitute those of the university are responsible. We may imagine that no individual anticipated the problem, but it is not in general an adequate response to a charge that someone is responsible for something in a way that warrants blame for a bad outcome that he or she did not anticipate a kind of problem that arose as a result of something he or she did. When it is someone’s responsibility to design a fair procedure, for example, and he or she does not think through adequately the implications of the procedure which is adopted, she can be blamed for the result. For example, someone who designs interrogation techniques while overlooking the likelihood of their inducing false confessions may be blamable (not necessarily solely) for the outcome of instituting such procedures. Likewise, university administrators can be responsible for failing to notice problems with the procedures which they institute. They have an obligation to formulate fair procedures, and if they do not anticipate circumstances in which a procedure they design results in unfair results, that does not automatically resolve them of their responsibility for the failure.

Further, those who formulate unfair rules can be legitimately blamed for decisions made in following them and their consequences. In the case in which the committee adopted its own decision procedure, the responsibility traces to the decision about the decision procedure. The individuals did not make a mistake in following the procedure they had designed. That the decision to use the procedure and the decisions involved in the particular application of it are distinct does not show that they were not blameworthy for the morally incorrect decision in virtue of their decision to adopt the procedure. The same goes if others design the procedure and require or expect its implementation. The blame for the outcome

attaches to the design and implementation of the procedure rather than the instruments used to implement it. Suppose certain people are responsible for certain morally problematic interrogation procedures, which others are responsible for carrying out. For example, suppose that they result in permanent personality disintegration. In this case, it is easy to see that those responsible for the procedures and rules are also morally blamable for the morally problematic outcomes of their application.

That is not to say that there may not be untoward consequences that no one could have reasonably anticipated and so for which they should not be blamed. It might be argued that it would be unreasonable in this case to expect anyone to have anticipated the problem, and Copp suggests as much (CAT, p. 380). However, if an individual were to be absolved of responsibility on that score, a group agent would be due the same consideration, for a group cannot be expected reasonably to anticipate something its members cannot. So one cannot have it both ways and maintain that the institution is responsible though no individuals are. For if the institution does it, as it were, unwittingly, this must be because individuals do things unwittingly. But if that lets the individuals off the hook, it must let the institution off the hook too. But, conversely, if the institution is to be blamed for unwitting harm, individuals cannot be let off the hook for the harm they are unwitting agents of.

The university may still have a responsibility to offer redress to the candidate even if it is not blamable. Such responsibilities can have sources besides being at fault. For example, a legislature may pass laws in accordance with the best scientific advice they have to allow certain additives to foodstuffs to allow a longer shelf life. If further research shows that these additives are harmful over long periods of time, it is certainly unclear that the legislature, which acted responsibly in the light of the information available to them, is morally blamable for the harm that occurred, though they arguably have a moral responsibility to offer redress to people who have been affected and to change the law. So we should be careful not to confuse an intuition that the university owes redress with an intuition that the university is morally blamable. That the university owes redress does not show that it is morally blamable.

To sum up, the tenure committee case is not one in which a collective is to blame for something though none of its members are. The university is either to blame or not. Even if it is not, it plausibly has a responsibility to offer redress to the candidate for the unfairness of the outcome. So we cannot conclude from our sense that the university owes redress to the conclusion that it was to blame. If the university is to blame, then it must be because its procedures are unfair and its having the procedures it had, though the problem with them was unanticipated, was yet something that the university ought to have anticipated or noticed, though it did not. But the rules were designed by individuals whose activities constituted the activities of the university, and, being charged with the task for the university, they likewise had the responsibility of designing a fair procedure, and in their failing to anticipate the untoward results in certain circumstances they are blam-

able for the same reason that the university is. And they are blamable not just for the rules, but, as we saw in the case involving the design of morally problematic interrogation procedures, also for the outcomes of applying the rules.

The argument generalizes. The case is an instance of a type in which certain individuals are jointly assigned the task of carrying out a procedure designed by other individuals where both groups are part of a larger group. The procedures in certain circumstances have morally problematic outcomes. The members of the implementation group are absolved of responsibility for implementing the procedures in the circumstances. Either the larger group ought to have anticipated the possibility of the untoward outcomes in the relevant circumstances and so is blamable for them, or not. If not, then no support accrues to CMA. If so, however, since the procedures were designed by members of the group charged with doing it for the larger group, the same considerations that support the view that the group ought to have anticipated the untoward outcomes support the view that the members charged with designing the procedures ought to have anticipated the untoward outcomes. In no case do we find that the larger group is responsible though none of its members are.⁶

(4) A fourth type of case is illustrated by Copp's example involving Bob, Carol, Ted, and Alice. As members of the Board of Governors of a state prison, they vote against expenditures for additional security procedures recommended by the Warden and the Board's staff. They are subsequently blamed for bloodshed in an escape attempt that would have been prevented by the recommended security measures. Carol brought forward the motion. Bob missed the meeting because of a family emergency. Alice voted against the measure because she always agrees with Ted, they alternate on who does research, it was Ted's turn, and she thought Ted indicated that she should vote against the proposal. None of them are blameworthy. Ted, however, had been under enormous stress because of a serious illness in his family. He had trouble sleeping and became confused, thinking that it was Alice's turn to do research. When the time to vote came, "he thought that Alice was signaling him to vote against the motion, and since she did vote against the motion, he thought he was doing the right thing in voting against it as well" (ACE, p. 219). But he would have voted for Carol's motion if he had done the research himself. Copp says: "it seems to me that Ted has an adequate excuse and that he should not be blamed personally for the bloodshed. It seems that it would be a mistake to absolve the Board from fault in the case" (ACE, p. 219).⁷

We can see that something has gone wrong with the description of the case as one in which the Board is to blame though all its members have exculpatory excuses by considering a parallel case in which the decision rests with just one individual who has exculpatory reasons whose force is equivalent to the reasons Ted has in the case above. Suppose the decision rests with Ted alone. He generally relies on an assistant for research. However, he has given his assistant the week off to attend to a family emergency. Ted has been himself under great stress due to a serious illness in the family. He is getting very little sleep. He becomes confused and forgets that he had given his assistant the week off. At his office, he notices his

assistant is out and looks on his desk for the report which he, in his confused state of mind, expects to find with the provisional recommendation from his assistant. There is a report there. He picks it up, and, reflecting that his assistant always does a good job, decides that he need not this time read through the report and quickly flips to the report summary, where he reads on the last page, "It is recommended that the warden's request be denied." He makes his decision quickly, notifying the warden of the denial. However, the report he picked up was one pertaining to a previous request the warden had made, rather than the one that Ted was to make a decision on. Because of his state of exhaustion, Ted failed to notice this. Thus, Ted makes the wrong decision, though it is understandable. The rest of the story goes as before.

Surely we will say Ted is to blame in this case, if we said the Board was to blame in the previous case. There is *no one else to blame*. But it seems that Ted has just as good exculpatory reasons in this case as in the other. If we blame him here, then we can blame him in the other case. If we excuse him in the first case, we must excuse him in this case. But in neither case can we blame the Board and let Ted off. It might be objected that there is some relevant difference between the excuses Ted has in the first and in the second case so that they are not the same in exculpatory force. But even if this were so, one would have to maintain that there is no way to modify it so that we give Ted exculpatory reasons equivalent in force to those he has in the first case, while letting the decision rest solely on him, in order to avoid the difficulty. For it only needs to be that we can find such a case to give the argument, and it seems very implausible that an individual working alone could not have reasons the same in exculpatory force as an individual working with a group.

This provides reason to think that in general we will not be able to describe a case of this kind which gets the desired result that a group is to blame, though no individual member is, because each member either acted correctly or had exculpatory reasons. For any case we are given, it seems plausible that we will be able to construct a case in which the decision rests with a single individual, rather than a group, whose exculpatory reasons are equivalent in force to those of the individuals it is proposed that we let off in the case in question, though we blame the group of which they are members. In the individual case, there will be a fact of the matter about whether the individual is to blame or not. If he is, then his exculpatory reasons are not sufficient to insulate him from blameworthiness. The same then will go for the members of the group whose exculpatory reasons were equivalent in force. If he is not to blame, however, then it would be inconsistent to blame the group in the parallel case.⁸

If these considerations are correct, then we cannot get a coherent case in which Ted is let off but the Board blamed. But what is the right description? Plausibly, if it is just Ted who makes the decision, given its gravity, the blame will rest with him. While we will not think it reflects badly on his character, given the facts of the case, we will also think that it was his responsibility to make the decision on the basis of appropriate research on the question. He had the respon-

sibility to do the research, and he did not do so, and so is to blame for the result. Similar remarks can be made about Ted in our first case. He had a responsibility to do the research and did not. We will not think it reflects badly on his character, given the circumstances, but he failed to carry out his duties responsibly. He is insofar to blame for the bad result.

The Board is to blame because its members did not all do their jobs adequately. Ted has some excuses, but excuses do not always exculpate, but may simply ameliorate. Plausibly Alice bears some responsibility as well, both for entering into her arrangement with Ted, which is to fail to take fully seriously her obligation to be fully informed on the issues that the Board must make decisions about (evidently Carol does the work for herself, why not Alice and Ted?), and for her failing to have arranged with Ted an adequate means of communicating what decision the one of them who was to have done the research had reached. If this were correct, then Ted would bear responsibility in this respect also.

It tends to confirm this diagnosis that the blame we direct at the Board is sensitive to the excuses its members can offer for failing to do their parts adequately. If they were all misinformed through no fault of their own about the need for the security measures, the Board will not be blamed. If the presentation to them is merely confused, less blame will attach to the Board. In the case above, it is clear that knowing the details does moderate the degree of blame that attaches to the Board. We can easily imagine the Board having acted blatantly irresponsibly and thereby coming in for much harsher criticism.

A final factor may be at work in the view that the Board's being to blame is not just a matter of the blame attaching to individual members. For, as we noted in discussion of our first case of someone being beaten to death by a group of people, the full measure of responsibility that attaches to the group may not attach to any individual member. That seems to be so also in the case at hand. Ted bears some blame, but Alice does plausibly as well, and so blame may seem to be shared between more than one individual, thereby diluting the individual portions. This may give the impression that the group must be to blame over and above the individuals because no individual has the full measure of blame that rests on the group. But that is just to say that the blame is shared, not that there is an extra portion left over after it has been appropriately allocated to the dissenters in the vote.

(5) A fifth case, introduced in CAT, is that of the kidnapped Prime Minister, who alone has the authority to sign a release for a dangerous prisoner, which her captors demand in return for her life. Stipulate that they will do what they say, that the Prime Minister's kidnapping is known to no one else, and that the Prime Minister by herself can sign the form and execute the release of the prisoner. Stipulate it would be wrong, all things considered—the danger the prisoner represents and the danger of copycat crimes—for the government to release the prisoner even if it meant saving the life of the Prime Minister. And stipulate as well that the Prime Minister knows with certainty all of these things. Let us say that to save herself from certain death, the Prime Minister, “acting in self-defense” (CAT,

p. 375), signs the release, and secures the prisoner's freedom. Who among us would not do the same thing? Surely, as Copp says, for someone to sacrifice herself in such circumstances would be heroic, but we cannot blame someone for failing to be heroic. Thus, we seem to have a case in which someone does something for which she is not blamable, though it constitutes the Government's doing something for which it is morally blamable, for the excuse that the Prime Minister has, that she acted in self-defense, is one that the government lacks.

Self-defense, in this case, I think, is not appropriately invoked. Self-defense as it is understood in the law allows for the use of force which would otherwise be unlawful when one is defending oneself against an attack, but specifically as directed against the attacker(s). It would not be correctly invoked as a justification or excuse for killing someone who is not attacking one, under threat from the attacker. For example, one does not have a right to kill an innocent bystander to save one's own life derived from the right to self-defense. Like the legal concept, the extra-legal concept of self-defense as a justification or excuse for an act that would otherwise be prohibited is properly invoked to justify or excuse acts directed at the attacker. In the present case, the Prime Minister is not using force which would otherwise be prohibited against those who are threatening her. Rather, she is releasing a dangerous prisoner, the release of whom she is well aware will very likely lead to harm even to the extent of death to others, and she knows furthermore that what she does will invite copycat crimes to secure the release of other dangerous criminals.

The Prime Minister is, however, acting for self-preservation, if not in self-defense. Does this justify or excuse her action? Suppose it were someone else, the Minister of Justice, with whom the decision lay. The description of the case leaves no doubt that he would be wrong to release the prisoner even to save the Prime Minister's life. The harm that comes from the release of the prisoner then is greater than the harm that would occur from the death of the Prime Minister at the hands of her kidnappers. If there is a difference when we put the power to release the prisoner in the hands of the Prime Minister, it must derive from a general principle allowing one to cause harm to others to save oneself from harm. Is there then a moral right each of us has to cause greater harm to greater numbers who do not threaten one to save oneself from death, which the Prime Minister could invoke to justify or excuse her releasing the prisoner? This seems implausible to say the least.

But can we *blame* her for what she does? Maybe not. *Who among us would not do the same?* But if we say this, it must be because we have shifted from evaluating what she does from the perspective of what it is morally right for her to do to the perspective of what it is right for *her* to do all things, moral and otherwise, considered. When we say she is not blamable, we do not mean she is morally blameless, but that she is not blameworthy from the standpoint of practical deliberation. It was all in all the *rational* thing for *her* to do. For unless we think moral considerations are strictly overriding in rational practical deliberation, we will allow that someone is not rationally blamable for a decision that may be

morally blameworthy. The Prime Minister is morally blamable but she is not rationally blamable, for in practical deliberation it is one's own all-in goals that are relevant, and there merely numerical differences which are irrelevant to morality are enormously important.

(6) The sixth and last case, also presented first in CAT, is that of the kidnapping of the Prime Minister's infant daughter. The Prime Minister knows with certainty that the child cannot be rescued and will die unless the prisoner is released, and she has the power to do it and knows that no one else knows or can be informed until it is too late. Again we stipulate that all things considered it would be morally wrong for the government to accede to the kidnappers' demand, even granting that it is the Prime Minister's infant daughter whose life is at stake; and the Prime Minister knows this with certainty. The suggestion is that the Prime Minister's *special duty* to her infant daughter makes it morally right for her to free the prisoner, though it would not be morally right for her to free the prisoner if it were the infant daughter of someone unrelated to her.

We must distinguish, as in the previous case, between whether it is morally right all things considered for the Prime Minister to release the prisoner and whether it is all-in rationally the best thing for her to do. Only the first question is relevant to our concern. It is plausibly the case that for most people their commitment to their children in life-threatening situations would trump other considerations, but this shows what it is rational for *them* to do, not necessarily what is morally required of them.

The Prime Minister's special duty to her daughter marks an important moral difference between her and other officers of the government. The question is whether it *can* override the moral obligations she has as Prime Minister given that it is wrong for the government to release the prisoner even if it means the death of an infant girl held hostage as leverage to force the release. The problem with the suggestion that it can shows up when we ask what the duties of the Prime Minister entail. For in accepting the position, the Prime Minister accepted a responsibility to carry out the duties of the office. There has been no suggestion, in the description of the case, that there was anything problematic about her taking up the duties of the office. So I will assume, to begin with, that the Prime Minister has accepted the responsibilities of her office without its being morally problematic that she should do so.

If it is morally best for the government not to release the prisoner in the circumstances, then the Prime Minister's duty in her official capacity is not to release the prisoner, even if that means the death of her infant daughter. In accepting her office, she accepted the duties which come with it and the weight they are to have relative to special duties she has. If those require, in the event that her infant daughter is kidnapped, that she not release a prisoner to save her, then she has accepted a position that will require her to put aside her special duty to her own child. We are assuming that the duties she incurs as Prime Minister do require her to sacrifice her child in the circumstances described. Otherwise, the power of the government residing in her, and its moral obligations being expressed in the

obligations resting on the office, it would not be the case that the government should, all things considered, in the circumstances, release the prisoner. We are assuming for now that it was morally coherent for her to accept the office, but it would not have been so if it required accepting duties which were to override duties that cannot be properly overridden by occupying such an office. If her accepting the office is morally unproblematic, then it involves her incurring an obligation that overrides her special duty to her child. Thus, if she releases the prisoner to save her child, she is failing to do her all-in duty. Consequently, she is morally blamable for releasing the prisoner despite her special duty to her child.

Suppose, however, that we wish to insist that she is not morally blamable taking everything into consideration. Then we must give up the assumption that (i) it was morally coherent for her to accept the office of Prime Minister or that (ii) her accepting it required accepting obligations which were to overrule certain special duties she has. Suppose that we accept (i) and reject (ii). In this case, she violates no obligations she has in representing the government in her office as Prime Minister in releasing the prisoner, and the government has in consequence done nothing wrong, for its obligations are expressed in the obligations which its effective instruments have in their official capacities.

Suppose we accept (ii) and reject (i). In this case, there is something wrong with the institutions of the government. The government requires one of its officers, in certain circumstances, to do something that it would be wrong for her to do, and no morally sound institution should do this. This casts doubt on whether it is right, all things considered, in these very special circumstances, on these assumptions, for the government not to release the prisoner. It seems plausible that in any circumstances in which a type of institution is faced with a decision, a morally sound institution of that type could reach the morally correct decision. If the Prime Minister held office in a morally sound institution, then if her acts constituted an act of the government, and she would not be morally blamable all-in, what the government did would not be morally wrong all-in. Therefore, on the assumption that the right decision in the case of the prisoner is the one that would be taken by a government with morally sound institutions, it would not be right, even in the case of a government whose institutions were not sound, in the special circumstances, for the government not to release the prisoner, given that it is all-in right for the Prime Minister to do so.

If this reasoning is correct, then there appears to be no correct description of the case that allows the Prime Minister to be morally blameless while the government is blamable. Since the response relies solely on general features of the case, it is applicable to any case in which there is a conflict between individual duties and duties one has in virtue of institutional role. For a morally sound institution, the obligations one incurs by virtue of accepting an office will not require one to do something that it is morally wrong for one to do all things considered. As the institution's obligations are expressed in the obligations of the fillers of various institutional roles, this means that it cannot be right for the institution to do something if it means it requires one of its role fillers to do

something in her official capacity that it is morally wrong for her to do. If the right decisions on any case facing an institution can be made by a morally sound institution of the relevant type, then no case of the sort described above is possible.

3. General Considerations

In this final section, I consider Copp's two general arguments for CMA. I also offer some brief general considerations in favor of what I will call the factor model of group responsibility, according to which the blame or praise that a group may be due can be factored into the contributions to what the group does of the various members of the group who are participants in the group action.

The first argument, which aims to support the conceptual possibility of CMA, goes as follows. Consider a simple version of consequentialism according to which an agent has an all-in moral obligation not to do something if and only if there is an alternative open to him with better consequences, and stipulate that something is a consequence of an agent's action if and only if it would not have occurred if the agent had not done what he did. Take a case in which a group does something that has consequences that would not have occurred otherwise, and suppose that there were alternative actions the group could have undertaken which would have had better consequences. Then, on this view, the group has an obligation not to do it and its doing it is wrong all things considered. Suppose, however, that for no member of the group is it true that if he had not done what he did, the outcome would not have occurred. Perhaps the participation of any n members of a group consisting of m members, where $m > n$, would have been sufficient, as in the case of a firing squad. Or perhaps others are ready to step in for each member of the group who does something to do the same thing if he should not. Copp says that, given our stipulation about "consequence," "the theory implies that it is possible for a collective to have an all-in obligation even if no member has a relevantly related all-in obligation," and, therefore, "if the conjunction of this consequentialist theory with our stipulation is not conceptually false," it is at least possible for a collective to have an all-in moral obligation though none of its members do (and so can be in the corresponding way blamable though none of its members are) (CAT, p. 372). We are invited to think that the conjunction is not conceptually false.

I will construe the theory as the claim that it is conceptually necessary that, for any x , x has an all-in obligation not to do something y if there is a range of actions r open to x among which is y and at least one action r has better consequences than y .⁹ If the theory-cum-stipulation (so understood) is not conceptually false, then we can readily grant that it is conceptually possible that CMA is true. Since the theory makes a claim about what is conceptually necessary, it is conceptually false if what it claims is conceptually necessary and is not true in every conceptually possible world.¹⁰ Therefore, one way to test whether it is conceptually false is to describe possible cases and ask whether the results of applying the theory-cum-stipulation accord with the intuitive application condi-

tions of our moral concepts. Consider a possible situation in which A and B hold C down while D rapes C. Suppose either that if A had not done his part, E would have or that only one person need hold the victim down so that neither A's nor B's contribution is necessary. Suppose that there are no other differences between the cases with respect to morally relevant consequences of A's actions. Does A *fail to have* an all things considered obligation not to participate? No. But according to our theory-cum-stipulation, the answer is "Yes," because D would have raped C even if A had not done what A did. We therefore have good reason to think that the theory-cum-stipulation is conceptually false.

The second argument relies on the distinct moral positions that groups and their members may occupy. Individuals have different sets of obligations than the groups of which they are members and may have available different excuses (special duties, for example). Copp argues that:¹¹

[a] Any given member of a collective that has an all-in obligation in a situation might have an overriding obligation that conflicts with any obligation he has to contribute to the collective's fulfilling its obligation. [b] In the absence of a reason to think that this could not be the case for all members, it appears possible that no member has an all-in duty to contribute to the collective's fulfilling its duty. (CAT, p. 373)

I take [a] to be the claim that: [any-*m*] for a given collective C and any member *m* of it, that C has an all-in obligation to A is compossible with *m* not having an all-in obligation to participate in C's A-ing. I take [b] as follows: if we do not have a reason to think that it is not the case that [every-*m*] a given collective C's having an all-in obligation to A is compossible with every member *m* of it not having an all-in obligation to participate in C's A-ing, then it is possible that no member has an all-in duty to contribute to the collective's fulfilling its (all-in) duty.¹² Since the consequent is equivalent to [every-*m*], the conditional has the form: if we do not have a reason to think that \sim [every-*m*], then [every-*m*]. I think we are invited to think that the antecedent is plausible and to draw the consequent as a further conclusion.

I do not see that [any-*m*] provides support for [every-*m*]. It does not do so in virtue of its form. The claim that for any member *m* of a given collective C, that C's doing A is compossible with *m*'s not participating in doing A, for example, clearly gives us no reason to think that for a given collective C, C's doing A is compossible with every member *m* of C not participating in doing A. Moreover, when we consider actual cases, we turn up reasons to think that the fact that, in a given case, any one of us could be excused while the group was still obligated to do something would not suggest that every one of us could be excused while the group was still obligated. For the fact that *we* have the all-in obligation has to be considered in deciding for each of us what our individual all-in obligations are. If the thing in question can be accomplished by some subgroup, then, if someone knows this, he may in virtue of other conflicting obligations be relieved of his obligation to participate. However, if the all-in obligation *we* have will not be met without a

particular person's contribution, and the obligation *we* have to do the thing in question is not overridden in the circumstances, it is certainly not clear that the individual's other duties can override his duty to participate. How could *we* have an all-in obligation to do something if one of us whose participation is necessary does not have an all-in obligation to participate in doing it? Similarly, suppose that n many of m members of a group ($m > n$) must participate to fulfill a group obligation. Then, while any $m-n$ number of members may be excused compatibly with their all-in obligation, no more may be. Clearly members of such a group must take into account how many others will participate if they do not in deciding whether they can be let off. Take a case in which a surgical team is removing a brain tumor in an emergency procedure, and all members of the team are essential for success. If a member of the team must miss her marriage anniversary dinner with her husband in order to participate, is she thereby excused from participating? No. Suppose her help not essential—then maybe she can be excused. Suppose something similar can be said for every member; that is, if he or she alone were excused the surgery could be completed successfully, so it is compatible with the group's having the all-in obligation that any particular member be excused. That clearly gives us no reason to think that *everyone could be simultaneously excused* while the group still has the all-in obligation to perform the surgery.

Does [any- m] though give us reason to think that *we do not have reason to think that it is not the case that* [every- m] or to accept the conditional [b]? I do not see that the one bears on the other in any straightforward way. Independently of [a], do we have reason to accept [b] and its antecedent? This depends on how we understand the antecedent. If "have a reason" is interpreted as "have an explicit argument," then we do not have a reason if we are just beginning to investigate the matter. On this reading, we have reason to accept the antecedent, but not the conditional. Suppose then we interpret it as equivalent to the claim that [every- m] is negatively conceivable in the sense that after considerable effort we have not been able to discover a contradiction in it, but have not yet positively conceived of a situation in which it is true. Negative conceivability in this sense does not always provide a reason to believe a proposition is (even) possible. For example, Goldbach's Conjecture and the Riemann Hypothesis seem to be negatively conceivable, but this does not by itself constitute a reason to accept that they are possible (and so true, since if true, they are necessarily true). Still, plausibly for some range of propositions, negative conceivability is a reason for accepting that a proposition is possible. If the proposition is about something being possible, this would provide reason to think that it was true. Construing the conditional in this way, we have: if that [every- m] is a proposition of a type for which negative conceivability is reason to think it possibly true and it is negatively conceivable, then (it is reasonable to think that) [every- m]. But even if we accept this conditional, we do not yet have a reason to accept the antecedent.

In conclusion, I describe briefly the factor model of collective responsibility and provide some reasons in its support. On the factor model, claims to the effect that a group or collective is responsible for some outcome are to be factored into

claims about the responsibilities of its members. Thus, in my first case in section 2 above, the claim that a group of people are responsible for beating someone to death and so blamable for it is to be factored into a claim about the responsibilities of individual members of the group and the blame each of them bears for the result. If their contributions were the same, then they share in equal measure responsibility for the death and are equally blamable for it.¹³ If the factor model is correct, then when a group is to blame for something, it must be the case that some of its members are to blame, for the blame for the outcome is simply to be distributed over the members on this model. What is to be said for this conception?

First, the source of all moral agency is individuals. The agency of groups is grounded in the agency of individuals, and there would be no group actions without individuals in them doing their parts.¹⁴ It makes sense then to think of the facts about the blameworthiness and praiseworthiness of groups as reducing to the facts about the blameworthiness and praiseworthiness of the members of those groups, and, indeed, this is our standard practice.

Second, a central role of activities like blaming and praising in our moral lives is to modify behavior. We aim to blame and praise where appropriate, and we expect that if the blame and praise are recognized as appropriate, then others (including those involved) will take it that the relevant conduct is to be avoided or emulated in accordance with whether the evaluation was negative or positive. That is, that an action is blameworthy is taken *in itself* to provide a reason to avoid it, and that it is praiseworthy is taken *in itself* to provide a reason to do it. This carries over also to praise and blame of groups. We blame groups for things and praise groups for things because in part we wish the behavior to be avoided or emulated, and we expect that the praise and blame will have a role in effecting this. Moreover, we expect it to have this role, as it does in the case of individuals, in virtue of the recognition of the action in question being blameworthy or praiseworthy by the group in question. Since the agency of groups is grounded in the agency of individuals, this means that we must suppose that the blame or praise the group comes in for will be seen as relevant to the behavior of individuals who are members of the groups by those individuals because of their recognition that the action in question is blameworthy or praiseworthy. The mechanism by which the blame or praise gives reasons to the group against or for action must go through the reasons of its members. If members of a group (or potential members of a group) see blame and praise that the group may come in for in its activities as translating into blame and praise for them, then we have a straightforward way of understanding how the practice of blaming and praising groups is to affect the behavior of groups in virtue of its recognition of the blameworthiness or praiseworthiness of group action by the group. If, however, a group could be blameworthy or praiseworthy without the members of the group being blameworthy or praiseworthy, then the members of the group would not *thereby* have a reason to regulate *their* behavior so as to avoid the blameworthy behavior of the group or to promote the praiseworthy behavior of the group simply through recognition of the group's behavior as being blameworthy or praiseworthy.¹⁵ This would disconnect

the practice of blaming and praising groups from a central role it plays in the regulation of moral conduct. I take that role to be central to our understanding of a practice as one of blame and praise. If that is right, then we should reject any view that allows blame and praise of groups to become disconnected from blame and praise due their members.

I am grateful to David Copp for providing me with a copy of his paper for this journal issue in time to take into account his latest views, as well as for extensive and helpful conversation and correspondence.

Notes

¹David Copp, "On the Agency of Certain Collective Entities: An Argument from 'Normative Autonomy'," *Midwest Studies in Philosophy* 30 (2006): 194–221.

²David Copp, "The Collective Moral Autonomy Thesis," *Journal of Social Philosophy* 38, no. 3 (2007): 369–88.

³I will concentrate on issues about blameworthiness. If I am right about the blameworthiness of groups reducing to that of individuals, the same will go for obligations and responsibilities. I am also concerned to argue that whatever sort of blameworthiness a group has (all things considered or pro tanto) it has in virtue of the same sort of blameworthiness attaching to at least some of its members. I am concerned throughout also with conceptual possibility, and I deny that possible cases of moral autonomy of collectives have been described.

⁴I assume following the procedure in light of the conflict between individual and category votes is in itself problematic and contributes to the untoward outcome. Nothing hinges on this, but it highlights a way they can be responsible without it resting on mistakes in their individual votes, which brings out a point relevant to the next case.

⁵Could the committee be said to be blamable though its members are not? No, for the committee's acting appropriately just comes to its members doing so. Are the committee members free from any blame for the outcome? One is not always absolved of responsibility on the grounds that one was obeying orders, as in the case of a soldier being ordered to torture a prisoner. However, I will assume the case fleshed out so as to absolve the members of the committee of responsibility.

⁶Copp suggests that "the fact that a moral argument can be given for the claim that the university was all-in blameworthy although no individual was even pro tanto blameworthy, and the fact that this argument can intelligibly be debated, are some reason to think it is at least *conceptually* possible for the university to be all-in blameworthy even if no individual is even pro tanto blameworthy" (CAT, p. 381). That an argument for whether it is the case that *p* can be intelligibly debated does not give us a reason to think it is at least conceptually possible that *p*. Suppose the question is whether the Continuum Hypothesis is true. We can intelligibly debate this. But this does not give us reason to think that the Continuum Hypothesis is at least conceptually possible, for that would mean that the fact that we were unclear about it was in itself a positive reason for its possibility, and hence truth, since if false, it is necessarily false.

⁷I think the Board is not responsible for the bloodshed, as that was caused by the actions of the prisoners, but they are arguably responsible for failing to upgrade as needed the security procedures in the light of a good case for doing so.

⁸It might be charged that this begs the question because the thesis is that a group can be blamed when no individuals can be, so that the fact that an individual would not be blamed if the decision rested solely with him, does not count against the thesis. However, the considerations which govern when moral agents in general are to blame should be applied in the same way to group and individual agents. The group may have its excuses and these come down to the excuses its members have. If

those do not excuse the group, then when the group shrinks to one they do not excuse the one who remains; and if they do excuse the one, then they likewise excuse the group.

- ⁹If we express the theory just as the biconditional, then its not being conceptually false is insufficient for the conceptual possibility of CMA because its being true in a conceptually possible world with no agents would show nothing about the possibility of CMA. If we treat the biconditional as conceptually necessary, then it is clear why its conceptual possibility would lead to the right result. It would be possible to hold a weaker position, namely, that the theory and a circumstance in which a group but none of its members are responsible for something is compossible. But if the theory is not expressing anything about the nature of right and wrong, we could have reason to think this possible only if we had independent reason to think CMA was.
- ¹⁰If $\Diamond \neg p$ then $\neg \Box p$. Conceptual necessity has an S5 modal logic. So, if $\neg \Box p$ then $\Box \neg \Box p$.
- ¹¹The considerations are the same for the second parallel argument, so I will consider only the first.
- ¹²I assume that “appears” in [b] is a remark about the epistemic status of the claim rather than strictly a part of it.
- ¹³Yet the appropriate *response* need not be less than each would receive if he had been solely responsible for the death. If the penalty for murder is hanging, it does not follow that the appropriate penalty for equal participation in beating someone to death with nine others is one-tenth of a hanging. Cooperation in killing someone does not lessen one’s liability to punishment.
- ¹⁴My “Collective Intentional Behavior from the Standpoint of Semantics” (*Noûs* 41 (2007): 355–93), argues that the logical form of plural action sentences shows that they do not involve commitment to group agents. The difference between the distributive and collective readings of “We built a boat” is a matter of whether the event quantifier introduced by the verb takes wide or narrow scope with respect to a quantifier over members of the group introduced by “We”: “[$\forall x: x \in us$]($\exists e$)(building(e , a boat) & agent(e , x) & [only $y = x$](agent(e , y)))” gives us the distributive reading and “($\exists e$)[$\forall x: x$ is one of us](building(e , a boat) & agent(e , x) & [only $y \in us$](agent(e , y)))” gives us the collective reading. It follows that collective praise and blame sentences in respect to collective action attributed using plural action sentences do not attribute praise and blame to group agents.
- ¹⁵They might have other reasons to modify their behavior, but the point here is that recognition of blameworthiness or praiseworthiness in itself provides a reason for modification of behavior and that that internal connection between an agent’s recognition of blame or praise as attaching to him and the agent having reasons for action is cut in the case of groups if we allow the group to be blameworthy or to be liable for praise though none of its members are.

Copyright of Journal of Social Philosophy is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.