

The Collective Moral Autonomy Thesis

David Copp

Suppose that a government wrongfully ordered the bombing of a civilian population, and suppose it was blameworthy for doing so. The question is whether it is possible that no member of the government was blameworthy or acted wrongly in the matter. More generally, the question is whether it is possible for a collective entity such as a corporation or a state to have an agential moral property even if no member of it has a relevantly corresponding agential moral property. According to the collective moral autonomy (CMA) thesis, this is possible.

In this paper, I restrict myself to considering two kinds of agential moral properties, “obligation” and “responsibility.” Accordingly, for my purposes, we can view the CMA thesis as the disjunction of two claims. First is the “forward-looking” claim (O), that it is possible for a collective to have a moral obligation to do something even if no (natural) person who is a member of that collective has a relevantly related moral obligation—an obligation, as we might say, to take part in bringing it about that the collective fulfills its obligation. Second is the “backward-looking” claim (R), that it is possible for a collective to be morally responsible for having done something (that is, to deserve a specific moral response for having done it, such as praise or blame) even if no (natural) person who is a member of that collective is morally responsible (or deserves the same kind of response) for having done something relevantly related—for having taken a part, as we might say, in bringing it about that the collective did the relevant thing.

The CMA thesis might seem dubious for at least two reasons. First, it seems clear that any collective entity is, roughly, “constituted” by its members.¹ This seems clearly to be true of unorganized or minimally organized collectives, such as mobs, discussion groups, and “arbitrary” groups, such as the passengers in an airplane. It seems also to be true of highly organized or institutionalized entities such as armies, universities, and criminal gangs. It might seem to follow that the properties of a collective must be accounted for by properties of its members. Yet even if this is so, it does not follow that a collective has a property only if at least some member has precisely that property. Some states are democracies, for example, but no person is a democracy. Second, collectives are not “independent agents”; that is, a collective can perform an action only if some person or persons act in a way that, in the circumstances, constitutes the collective’s performing that action.² But, as I will explain, even if the actions of certain persons constitute an action of a collective, it does not follow that these persons must have acted

wrongly if the collective did or be blameworthy if the collective is. It does not follow that the CMA thesis is false.³

The CMA thesis is the disjunction of claims O and R, and, as I will explain, each of these claims can also be viewed as a disjunction of several claims. I speak of these claims as “sub-theses” of O and R and hence as sub-theses of the CMA thesis. In section 1, I distinguish eight sub-theses and briefly present an argument for one of them. In section 2, I propose an argument for an additional five. In the central section of the paper, I discuss a series of examples that seem to support all but one of the sub-theses. In section 4, I discuss objections. My arguments are not decisive, but I believe they show that the CMA thesis cannot easily be dismissed. I conclude by cautioning against drawing naive political conclusions from my arguments.

1. Eight Sub-theses

To begin, I need to explain my terminology. I shall not distinguish among moral obligations, moral duties, and moral requirements, nor shall I distinguish between what a person is morally required to do and what she ought morally to do. I shall say that one has an obligation “all things considered” or “all-in” to do A just in case one has an obligation simpliciter to do A, all relevant moral considerations being taken into account. I stipulate that one has a “pro tanto” obligation to do A just in case there is a moral reason for one to do A where, in the absence of conflicting reasons, this reason would be sufficient to make it the case that one has an obligation all things considered to do A.⁴ A pro tanto obligation may be outweighed by conflicting considerations.

I speak of an agent as “morally responsible” for having done A just when the agent deserves or is liable to an appropriate negative or positive moral response for having done A. For example, on the negative side, the appropriate response might be blame or a demand for compensation or for an apology, and, on the positive side, the appropriate response might be praise.⁵ I stipulate that an agent is responsible “pro tanto” for having done A just in case there is a moral basis for holding the agent to be deserving of a negative or positive moral response of some kind for having done A, where, in the absence of a countervailing consideration, such as an excuse or justification, this reason would be sufficient to make it the case that the agent deserves the response all things considered. Accordingly, a person who has done something that is pro tanto wrong, and for which she is pro tanto blameworthy, might not have done something that is wrong all things considered, and she might not be blameworthy all things considered.

Both O and R have strong and weak versions on two different dimensions, a “stringency” dimension and a “modal” dimension. There are therefore four versions of each.

On the dimension of stringency, we can distinguish “all things considered” and “pro tanto” versions of both O and R. The all things considered or all-in version of O says it is possible for a collective to have an all-in moral obligation

even if no member has a relevantly related all-in moral obligation, and the corresponding version of R says it is possible for a collective to be all-in morally responsible for something even if no member is all-in morally responsible in the matter. The pro tanto versions say it is possible for a collective to have an all-in moral obligation even if no member has even a relevant pro tanto moral obligation, and it is possible for a collective to be all-in morally responsible for something even if no member is even pro tanto morally responsible.

Someone might accept the all things considered versions on the ground that the fact that a collective has a duty or bears moral responsibility all things considered does not rule out the possibility that its members have excuses or that countervailing moral considerations give them conflicting and overriding duties. Yet she might reject the corresponding pro tanto versions. For she might think that the fact that a collective has a duty or bears moral responsibility all things considered is itself a consideration that gives at least some members a pro tanto duty to do their part in bringing it about that the collective does its duty, and, if they fail, it gives them at least pro tanto responsibility in the matter. She might therefore insist that if a collective has an all-in moral obligation or if it is all-in morally responsible, then at least some member must have a related pro tanto moral obligation or be pro tanto morally responsible.

On the modal dimension, we can distinguish conceptual versions and substantive versions of both O and R. A substantive version says that the correct moral view allows the possibility of a collective's having an agential moral property even if no member has a relevantly corresponding agential moral property. Collective moral autonomy is "morally possible." A conceptual version says merely that it is conceptually possible or conceivable that a collective has an agential moral property even if no member has a relevantly corresponding agential moral property. That is, the negation of this thesis is not a conceptual truth.⁶ Collective moral autonomy is "conceptually possible." To see the distinction I have in mind, notice that we may think that consequentialism is false, without thinking that it is conceptually false—without thinking that its negation is a conceptual truth.

I am taking the CMA thesis to be the disjunction of O, the thesis about obligation, with R, the thesis about responsibility. For each of these, we have now distinguished four sub-theses: the conceptual/all things considered thesis, the conceptual/pro tanto thesis, the substantive/all things considered thesis, and the substantive/pro tanto thesis. Accordingly, each of O and R can be viewed as the disjunction of four sub-theses. The CMA thesis is the disjunction of all eight:

- (O-C-all): It is conceptually possible for a collective to have an all-in moral obligation even if no member has a relevantly related all-in moral obligation.
- (O-C-PT): It is conceptually possible for a collective to have an all-in moral obligation even if no member has a relevantly related pro tanto moral obligation.

- (O-S-all): It is morally possible for a collective to have an all-in moral obligation even if no member has a relevantly related all-in moral obligation.
- (O-S-PT): It is morally possible for a collective to have an all-in moral obligation even if no member has a relevantly related pro tanto moral obligation.
- (R-C-all): It is conceptually possible for a collective to be all-in morally liable to a moral response of a given kind for having done something even if no member is all-in morally liable to that kind of response for having done something relevantly related.
- (R-C-PT): It is conceptually possible for a collective to be all-in morally liable to a moral response of a given kind for having done something even if no member is pro tanto morally liable to that kind of response for having done something relevantly related.
- (R-S-all): It is morally possible for a collective to be all-in morally liable to a moral response of a given kind for having done something even if no member is all-in morally liable to that kind of response for having done something relevantly related.
- (R-S-PT): It is morally possible for a collective to be all-in morally liable to a moral response of a given kind for having done something even if no member is pro tanto morally liable to that kind of response for having done something relevantly related.

The following argument seems to support O-C-all. A simple consequentialist theory says that an agent has an all-in moral obligation not to do A just in case there is an alternative open to the agent that would have better consequences than doing A. Stipulate that E is a “consequence” of an agent’s doing A just in case E would not have occurred if the agent had not done A. The combination of the consequentialist theory with this stipulation does not seem to be conceptually false. Imagine now that a state is at war and that it bombs an enemy city, and suppose that a certain bad upshot would not otherwise have occurred. It is possible, nevertheless, that there is no individual who did anything relevant such that the bad upshot would not have occurred if she had not done what she did; perhaps many people were ready and able to drop the bombs in place of the persons who actually did drop the bombs. In this case, given the above stipulation, the bad upshot is a “consequence” of the state’s bombing the city, but it is not a “consequence” of any member’s action. Hence, given this stipulation, our theory implies that even if the state had an all-in obligation not to bomb the city, no person need have had an all-in obligation to avoid the thing she did that contributed to the bombing. Given our stipulation, then, 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. Hence, if the conjunction of this consequentialist theory with our stipulation is not conceptually false, then O-C-all is true,

and since O-C-all is one disjunct of the CMA thesis as I understand it, it appears that the CMA thesis is true.

In what follows, I shall argue for the plausibility of all four conceptual theses as well as three of the substantive theses. One of the substantive theses, O-S-PT, strikes me as false, and I shall present an argument to undermine it.

2. An Argument for Six Sub-theses

It is intuitively plausible that collectives can have duties and that they can bear responsibility when they fail to fulfill their duties. A corporation ought to treat its employees fairly. A country ought not to wage aggressive or preemptive war. The members of a collective also have duties, of course, including pro tanto duties to do their part as appropriate in bringing it about that the collectives they belong to meet their duties. But members of a collective also have duties arising from their personal situations, including family duties. Such personal duties could conflict with, or get in the way of a person's fulfilling, her duties as a member of a collective. Hence there can be situations in which the duties of a collective are simple, but in which a person who is in an official position in the collective has complex personal duties that get in the way of her fulfilling her duty to do her part in bringing it about that the collective meets its duties. Or the person might have excuses for doing what she does that stem from her personal situation, excuses of a kind that are not available to the collective. Given this, we should not be surprised by cases in which a collective has agential moral properties that differ from those of its members.

A collective and the individual persons who are its members are distinct agents and can face different moral situations. In a case in which the moral situation of a collective is simple and one-dimensional, its members might face a variety of conflicting obligations. 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. 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. Similarly, any given member of a collective that is liable to blame or to some other negative moral response in a situation might be able to excuse or to justify his failing to contribute to the collective's doing its duty in the situation. He might have an excuse that is not available to the collective or he might have had an obligation that conflicted with and overrode any obligation he had to contribute to the collective's fulfilling its obligation. In the absence of a reason to think that this could not be the case for all members, it appears possible that no member is all-in blameworthy even though the collective is all-in blameworthy. It appears, then, that it could be that no member has an all-in obligation to contribute to a collective's fulfilling its all-in obligation in a situation and that no member is all-in blameworthy for what he did in a situation in which a collective is all-in blameworthy.

This argument supports O-C-all, O-S-all, R-C-all, and R-S-all, the conceptual and substantive versions of the all-in sub-theses about both obligation and responsibility. It does not support the pro tanto versions of the theses. For nothing in the argument undermines the claim that if a collective has a duty, this is itself a moral consideration that gives at least some members a pro tanto duty to do their part in bringing it about that the collective does its duty, and if the collective fails to do its duty and is blameworthy, this is itself a consideration in light of which those who failed to do their part are at least pro tanto blameworthy in the matter. Call this claim the “transfer doctrine.” The argument does not undermine the transfer doctrine. If the transfer doctrine is either a conceptual truth or a substantive moral truth, we must reject the corresponding pro tanto sub-theses, for it will not be possible that a collective has an all-in duty or that it is all-in blameworthy unless some members have at least a relevant pro tanto duty or are pro tanto blameworthy.

The transfer doctrine does not strike me as a conceptual truth, however. I think that a person could grasp the doctrine, consider it carefully, yet coherently deny it. A person with a highly individualistic moral view might deny that the duties and responsibilities people assume when they take roles in collectives have any moral significance. Such a position appears to be coherent. If it is, then the conceptual pro tanto doctrines O-C-PT and R-C-PT might be correct.

Yet even if the transfer doctrine is not a conceptual truth, it does seem substantively plausible. A person in a given role in an institution may thereby have the authority to act on behalf of the collective, or otherwise to contribute to the collective’s doing something. If a person has voluntarily accepted such a role in a collective and if she has the authority to contribute to the collective’s doing its moral duty, it seems plausible that she has a pro tanto duty to do so. And it seems as well that a person in such a role will be pro tanto blameworthy if she fails to act in accord with her pro tanto duties. Of course there are cases in which a person does not voluntarily accept her role in a collective and cases in which a collective is not organized in a way that involves assigning roles to members. Even in these cases, I am inclined to say that if a person can play a part in bringing it about that a collective of which she is a member meets a duty, she has a pro tanto duty to do so; it may also be the case that she would be pro tanto blameworthy if she did not. Perhaps, then, we should reject both O-S-PT and R-S-PT, the substantive sub-theses about pro tanto obligation and responsibility.

As we shall see, one of the examples that we will consider in the next section seems to support R-S-PT. This leaves the substantive sub-thesis about pro tanto obligation, or O-S-PT, which I am inclined to reject.

I hold in general that people have moral reason to assist others to meet their duties. If a neighbor has a moral duty to do something, and if I can help, I have a moral reason to assist, and this means that I have a pro tanto duty to assist. It follows that people have a pro tanto duty to do their part in bringing it about that collectives to which they belong meet their duties. It may be that this duty is easily overridden, but it is there.

In the end, therefore, I find it plausible to accept all of the CMA theses with the exception of O-S-PT: when a collective has an all-in duty, some members have at least a pro tanto duty to do their part in bringing it about that the collective meets its duty.

3. Examples

3.1 *The Kidnapping: The Prime Minister*

I begin with an example that supports R-C-all and R-S-all—the sub-theses about all-in responsibility. In the next subsection, I modify the example so that it supports O-C-all and O-S-all—the sub-theses about all-in obligation.

Suppose a person is taken hostage by an outlaw group. The kidnapers threaten to kill the hostage unless the government releases a certain prisoner. Suppose the circumstances are such that, all things considered, the government ought not to release the prisoner. The prisoner is dangerous. But more important, let us suppose, there is a danger that if the government acceded to the demand of these outlaws, other groups would adopt similar tactics. It would be best for the government to stand up to the kidnapers in this first case to avoid giving outlaw groups the idea that it can easily be manipulated by threats of various kinds. For this reason, let us say, all things considered the government ought not to release the prisoner. If it does release the prisoner, and if it does not have an exculpatory excuse, it is blameworthy for doing so.

Under the circumstances, let us suppose, only the Prime Minister (PM) can authorize the release of the prisoner. If she signs an official order, and if the prisoner is released in accord with the order, the government will have released the prisoner. In this case, it seems plausible that the PM has an obligation not to sign the order. If she does sign, and if she lacks an excuse, she will be blameworthy. As I shall explain, however, it is possible for her to have an excuse without the government's having an excuse. Because of this, it is possible for the government to be all-in blameworthy for releasing the prisoner even though the PM is not all-in blameworthy for signing the order.

Suppose it is the PM who is the hostage and whose life is threatened by the kidnapers. The kidnapers are in her office and are holding a gun to her head. The PM realizes that her predicament is not known and that she will not be rescued. In this situation, the balance of considerations still is such that the PM ought not to sign the release order, it seems to me, and it remains true that the government ought not to release the prisoner. The fact that a powerful citizen is threatened with death does not permit the prisoner's release. But suppose that the PM does sign the order, acting in self-defense. A person can at least sometimes be excused for acting wrongly in order to save his own life. The fact that one feared for one's life is a powerful exculpatory excuse in all but the most extreme circumstances.⁷ The situation I have imagined is not so extreme that this excuse is not available to the PM. The prisoner is not that dangerous, I say, and the risk is not that great. It seems

then that she has the excuse that she acted in self-defense. But the government cannot claim that it acted in self-defense. The government will continue to exist even if the PM is killed. Apparently, then, in the circumstances I have imagined, the PM is not all-in blameworthy for signing the order, since she acted in self-defense, but the government is blameworthy, for it lacks an adequate excuse.

One might object that the government has the excuse that it acted to save the life of one of its members. But even if we grant that this is an excuse, it is different from the PM's excuse that she feared for her life. Since the excuses are different, they might well have different force, so there is no reason to suppose that the government's excuse would be exculpatory in every case in which the PM's excuse would be exculpatory. I think, moreover, that we should not grant that the government had an excuse. The government would have had no excuse if some other powerful person had been the hostage, and the fact that the hostage is the PM does not change matters relevantly. We expect a government not to let itself be overwhelmed by fear and we expect it not to act out of favoritism to protect its members. Individual persons are in a different position. We excuse individual persons when they act out of well-grounded fear for their own lives. If the PM had chosen to let herself be killed rather than sign the release order, we would have viewed her as heroic. This reflects the fact that we would not blame her for acting to save her life. We do not blame people for failing to act heroically.

3.2 The Kidnapping: The Prime Minister's Daughter

Let us now modify the example so that it illustrates the possibility of a collective's having an all-in obligation even though no individual person has an all-in obligation to do his part in bringing it about that the collective fulfills its obligation. Suppose that the situation is just as I described it in the preceding subsection except that it is the PM's infant daughter who has been kidnapped and is threatened with death unless the government releases the prisoner.

No one but the PM knows about the kidnapping and the child cannot be rescued. The balance of considerations still is such that the government ought not to release the prisoner. It is a difficult case since the hostage is a child, but let us suppose that, despite this, the government still ought not to release the prisoner. The PM is in a different position from the government's, however, because the child is her daughter; obviously the child is not the government's daughter. The PM has an obligation toward her children that she does not have toward other citizens. For her, then, let us say, the balance of considerations permits her to do what she can to rescue her daughter. In certain cases, and with certain prisoners, the obligation she has toward her daughter would be overridden by her duty not to sign. But in this case it is not. She signs the order and the government releases the prisoner. In this case, the government has violated an all-in obligation not to release the prisoner. However, neither the PM nor any other individual has violated an all-in obligation.

One might think that there must have been certain officials who carried out the PM's release order. Even though they knew nothing about the kidnapping, one might think that objectively they had an obligation not to carry out the order since objectively the government had an obligation not to release the prisoner. The PM's order was a legal order, however, and it is plausible that officials in a morally sound state have an objective duty to carry out orders that are legally given to them except in extraordinary circumstances of a much more serious nature than are involved in the kidnapping case. So it seems to me plausible that the officials did not have an *all-in* duty not to carry out the order. In any event, to avoid this objection, let me stipulate that the PM has the ability to release the prisoner without depending on the cooperation of any officials.

One might object that the PM has an obligation not to take into account the special relationship she has to members of her family in deciding what to do in her official capacity. Yet in some cases an official's duty to his family may override his official duties. For example, a PM might rightly think that her duty to attend a family funeral outweighs her official duty to meet a visiting dignitary. There is no reason to think that in a case of the kind I have imagined the PM's duty to her daughter could not override her duty not to release the prisoner.

It seems to me that the kidnapping examples support the idea that it is conceptually possible for a collective to be blameworthy all-in for something even if no member is blameworthy all-in for the role she played in the matter, and that it is also conceptually possible for a collective to have an all-in duty even if no member has an all-in duty in the matter. That is, the examples support R-C-all and O-C-all. It also seems to me that the examples support R-S-all and O-S-all. That is, it seems plausible that the government ought not to release the prisoner. And, in the first example, it seems plausible that the PM is excused because she acted in self-defense, and, in the second example, it seems plausible that her duty to her daughter overrides her duty not to release the prisoner. In the next subsection, I provide another example that seems to support R-C-all and R-S-all. In the final subsection I propose an example that, if I am correct, supports all of the doctrines except the substantive pro tanto thesis about obligation, O-S-PT.

3.3 *The Prison Board*

Suppose that Bob, Carol, Ted, and Alice are the members of the Board of Governors of the state prison. A month ago, the Board voted against strengthening security procedures at the prison, and it is now being blamed for this decision in light of bloodshed resulting from a recent attempted escape. The Governor of the prison had been worried and had asked the Board to authorize new expenditures to strengthen security. The Board's staff had supported the Governor's request and had made available to the Board extensive research that strongly supported the need to increase funding. This information had been available to the Board for several weeks before Carol brought forward a motion calling for the new expenditures on security. Despite the strong representations of the staff, the Governor,

and Carol, the Board voted the motion down. If the expenditures had been authorized, the Governor would have made the improvements in security well in advance of the escape attempt. This is why the Board is being blamed for its decision.

It would be unreasonable to blame Carol since she brought forward the motion to increase the funding, she argued for the motion, and she voted for it. Bob missed the meeting. He had a very good excuse, since he was out of town dealing with a family emergency. Alice voted against the motion, but she had an excuse for doing so. She and Ted agree on most issues, and there are too many difficult issues that come before the Board for both of them to do thorough research on every issue. They have a standing agreement to take turns doing the necessary research, and it was Ted's turn to do the research. Alice voted against the motion, seeing that Ted was also voting against it, and believing that Ted had done the research needed to make a reasonable decision. She believed on good grounds that she and Ted would have agreed if both were to have done the research, and she would have supported Carol's motion if she had done the research herself. But, under the circumstances, she cannot reasonably be blamed for voting against the motion.

As I described the case, it seems reasonable all things considered to hold the Board to be blameworthy, but all things considered three of the four Board members are not blameworthy. The question is whether it follows that Ted is blameworthy. Suppose we learn that Ted had an excuse. Would accepting his excuse as freeing him from blame entail that the Board also had an excuse that excused it from blame? Or could it be instead that although the Board is all-in blameworthy, no member is all-in blameworthy?

Suppose that Ted has been under enormous stress because of serious illness in his family and as a result he has been unable to sleep. Given his exhaustion, he mis-remembered whose turn it was. He believed it was Alice's turn. When the time came to vote, he was confused as a result of his exhaustion, and he thought that Alice was signaling him to vote against the motion. Since she did vote against the motion, he thought he was doing the right thing in voting against it. He would have supported Carol's motion if he had done the necessary research. In this case, it seems to me that Ted is not blameworthy for voting against the motion. It also seems, however, that it would be a mistake to absolve the Board from blame for the decision.

One might object that Ted's excuse is not sufficient to remove all blame. But then we can imagine a case where his excuse is even stronger. Or one might object that if the members of the Board are not blameworthy, the Board cannot be blameworthy. However, this objection would be question-begging. One might object instead that if the members are not all-in blameworthy, this by itself excuses the Board from blame. This, however, is not a familiar kind of excuse and one might wonder why we should accept it. It would be question-begging in this context to insist there is such an excuse. One might argue that it would be wrong to blame the Board because doing so would harm the members' reputations. But

it does not follow that the Board is not blameworthy. Blaming someone can sometimes affect other people unfairly, and because of this it may be wrong to blame someone even if he is blameworthy.

I think the example supports R-C-all for I think it shows it is conceptually possible that the members of the Board are not all-in blameworthy even if the Board is all-in blameworthy. To reject R-C-all, one would have to claim that my example and all similar examples are ruled out on conceptual grounds. One would have to claim that all such cases are conceptually incoherent. One would need an argument against the conceivability of the case, not merely an argument that the case is morally implausible.

I also think that the example supports the substantive thesis R-S-all, for I think the example is morally plausible. I think that the Board is appropriately blamed, all things considered. Moreover, it seems to me that (in a well-constructed example of this kind, even if not in the example precisely as I have described it) the individual members of the Board have excuses and do not deserve blame, all things considered. They do not need to resign in shame, for example. To be sure, it would be appropriate for each of them to feel bad about the Board's decision. The reason for this is that they are members of the Board, and the Board made a decision for which it is blameworthy. Given Ted's role in events, it would be especially appropriate for Ted to feel bad. Assuming, moreover, that Ted and the others remain on the Board, each of them ought to accept responsibility for dealing with the security problem in the future. That is, they have a "role-responsibility" that they must accept, assuming they keep their roles.⁸ Nevertheless, as I have described the case, they are not blameworthy.

3.4 The Tenure Committee

My final example builds on a case that Christian List has used to illustrate the so-called discursive dilemma.⁹ Suppose that a university's standard for tenure says that excellence in the three areas of research, teaching, and service is necessary and sufficient for tenure. The university's policy provides for a three-member committee. The committee is required to vote on a candidate's qualifications in each of the three areas before voting on the overall question of whether to grant tenure. It is required to inform candidates of its decisions as they are taken, but without revealing how the members voted. The policy stipulates that decisions taken by the committee are final and constitute decisions of the university.

In the case of Mr. Borderline, in each of three areas, a majority of the committee votes that he has achieved the required standard of excellence. The members are A, B, and C, and the votes are as follows: A and B vote in favor of Borderline in the area of research but C votes against; B and C vote in favor in the area of teaching but A votes against; A and C vote in favor in the area of service but B votes against. All three then vote against awarding tenure to Borderline because all three think that he has failed to achieve excellence in all three areas. These decisions of the committee constitute decisions of the university, so the

university thereby denies tenure to Borderline despite having decided that he had achieved excellence in all three areas. The university's decision to deny tenure to Borderline therefore violates its own standard for tenure. Because of this, the university's decision was unfair to Borderline. It frustrated his legitimate expectation that the university would comply with its own standards. It was also unfair because he was told of the committee's decisions in the three areas, and this gave him a reasonable expectation that the committee would decide in favor of tenure. Despite this, however, each member of the committee voted in accord with the university's rules and standards, and no member decided or voted in a morally problematic way. Moreover, we can assume that each member reasonably believed that the university's standard for tenure was appropriate. Each believed that a candidate deserves tenure if and only if she has achieved excellence in all three areas. Each voted in a principled way. Each reasonably believed that the result was the right one.

It seems plausible in this case, however, that Borderline was treated unfairly and that the university is all-in blameworthy or liable to a negative response. It is liable at least to apologize to Borderline and perhaps to compensate him. But I think that no member of the committee acted unfairly and that no member is blameworthy; no member owes Borderline an apology and no member is liable to compensate him. Perhaps they ought to feel bad about what happened, but this is a different matter.¹⁰ It seems to me that they are not even pro tanto blameworthy. There is no moral charge against them that they need to answer by presenting an excuse. They acted reasonably and in accord with their duties as members of the committee and their individual votes were fair and accurately reflected their judgments about Borderline's deserts.

To be sure, the university could have had different procedures. It could have required tenure to be decided by a single vote on the tenure issue while prohibiting separate votes on candidates' achievements in the three areas. It also could have required that the decisions of tenure committees be confidential, except of course for the final decision. Given that the university could have had different rules, one might think that someone must be blameworthy for the faultiness of the rules it actually had.¹¹

Let me therefore stipulate 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. Let me also stipulate that the Borderline case arose long before paradoxes and problems of voting procedures came to be widely known and to be studied in universities. Given these stipulations, and given the familiar technical problems in designing voting procedures,¹² it would be unreasonable to hold that someone in the university must be blameworthy for the faultiness of the university's rules. Moreover, the issue is whether someone was blameworthy for the unfairness of the decision in Borderline's case, not whether someone was blameworthy for the faultiness of the rules. A successful objection must show either that the university was not blameworthy for the decision or that some person was blameworthy for his role in that decision.

To me it seems plausible that the university was blameworthy for the decision or at least liable to apologize or to compensate Borderline even though no person was blameworthy in the matter or liable to apologize or to compensate Borderline. Accordingly, I think the example supports both of the substantive theses about responsibility, R-S-all and R-S-PT. It therefore also supports the corresponding conceptual theses, R-C-all and R-C-PT. To argue against the conceptual theses, one would have to argue that the case as I described it is not even conceptually possible. One would have to argue that the example is incoherent. It would not do to argue on moral grounds that the university had an excuse or that if we look far enough we will find some individual who is blameworthy. Indeed, 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. This is reason to accept R-C-all as well as R-C-PT.

I think in addition that the example supports the all-in doctrines about obligation, O-C-all and O-S-all. For it is plausible, it seems to me that, given its standards for tenure, and given that the tenure committee decided in favor of Borderline in the three required areas, the university had an all-in obligation not to deny tenure to Borderline. Yet none of the members of the committee had an all-in obligation to vote in favor of tenure since each reasonably believed that Borderline had failed to achieve excellence in all three areas and since each reasonably believed that the university's standard for tenure was appropriate. To vote in favor of tenure, each of them would have had to vote to award tenure to someone they believed did not deserve it. Their vote would also have been in violation of the university's procedures.¹³ Hence it is plausible that although the university had an all-in obligation not to deny tenure to Borderline, each member of the committee was at least permitted, all-in, to vote against tenure—even though their unanimous vote constituted the university's deciding against tenure and thereby constituted its violating its obligation. The university had an all-in obligation even though no member had the relevant corresponding all-in obligation.

I think, finally, that the example supports O-C-PT, the conceptual doctrine about pro tanto obligation. For I think it is coherent to hold that the members of the committee did not even have a pro tanto obligation to vote for tenure. It is coherent to think that their only obligation was to vote on the basis of their assessments of whether Borderline achieved excellence in all three areas.

Nevertheless, I believe that the example does not support the substantive doctrine about pro tanto obligation, or O-S-PT. There was a reason for the university to award tenure to Borderline, and I believe this was also a reason for the members of the committee to vote in favor of tenure. The committee decided that Borderline had achieved excellence in all three required areas, and the university's standard for tenure requires that a candidate be awarded tenure if he achieves excellence in all three areas. Hence, to avoid unfairness, the university

had an obligation to award tenure to Borderline and this was also a reason for the members of the committee to vote for tenure. They therefore had a pro tanto obligation to vote for tenure. Despite this, of course, it seems to me that each of them was at least permitted, all things considered, to vote on the basis of her own assessments of Borderline's accomplishments. To vote in favor of Borderline, each of them would have had to vote to award tenure to someone whom she thinks does not deserve it. So although it seems to me that the members were permitted all things considered to vote against tenure, it also seems to me that they had a pro tanto duty to vote for tenure. I therefore do not think that the example supports O-S-PT.

One might object that if the members had a pro tanto obligation to vote in favor of tenure, then they were at least pro tanto blameworthy for not doing so. But as I said, it seems to me that the members of the committee need no excuse or justification for what they did. There is no moral charge that they need to answer since they voted in a principled way and in accord with their duties under the university's policies. So I do not believe that they were pro tanto blameworthy. This suggests that their pro tanto duty to vote in favor of tenure was exceptionally weak, since acting contrary to it does not even create a pro tanto blameworthiness. Their pro tanto duty was based in considerations of fairness, given the committee's prior votes and given the university's standard for tenure, but, given their beliefs, each of the members voted in a way that was fair to Borderline and each voted in a principled way. So it seems plausible that although they had a pro tanto duty to vote for tenure, they were not even pro tanto blameworthy for not doing so.

To summarize, it seems to me that the tenure example supports all of the CMA sub-theses except the substantive doctrine about pro tanto obligation, or O-S-PT.

4. Objections

In this section, I consider four general objections to the CMA thesis. I should mention that there is a tendency for people to bring forward question-begging objections. For example, someone might argue that if no member of the government was blameworthy for releasing the prisoner in the first kidnapping example, the government as such must not have been blameworthy. Objections of these kinds are unhelpful, and I will set them aside.

There are sophisticated theories that entail the negation of at least some of the CMA sub-theses. For example, Raimo Tuomela's theory seems to entail the falsity of at least the pro tanto sub-theses about responsibility, R-C-PT and R-S-PT.¹⁴ But I think my examples are intuitively plausible. If they provide intuitive support for a given CMA sub-thesis, they give us intuitive reason to doubt or to modify theories that imply that this sub-thesis is false.

The first objection I want to discuss begins with the contention that collectives are not capable of intentional action. It argues on this basis that collectives are not moral agents capable of having obligations or of being blameworthy. I am unable

to address this worry in detail here.¹⁵ First, however, as I said before, I agree that collectives are not “independent agents.” It might even be the case that any sentence that (putatively) ascribes agency to a collective is logically equivalent to some (perhaps very complex) sentence that ascribes agency only to individual persons.¹⁶ But it does not follow from this claim that collectives are not agents. On the contrary, the claim tells us under what circumstances a sentence that ascribes agency to a collective is true. It tells us under what circumstances the actions of individual persons constitute an action of a collective. Second, the thesis that no collective is a moral agent has implausible implications. It implies that no state has ever wrongly bombed a city and that no state has ever deserved blame for imprisoning an innocent person. These claims are implausible, and this is an objection to theories that imply that collectives are not moral agents. Nevertheless, my arguments can be viewed as conditional. If collectives can have moral obligations or be blameworthy, then, I am arguing, the CMA thesis is true.

The second objection turns on a distinction between the moral situation of a person in her own right, so to speak, and the moral situation she might be in *qua* occupant of a role in a collective. The idea is that a person with such a role might have duties, or she might be blameworthy for something, *qua* the occupant of that role. For example, one might distinguish the duties of the person who happens to be PM from the duties that this person has *qua* PM. One might hold that although the fact that she acted in self-defense would excuse the woman, it would not excuse her *qua* PM. Following this approach, one might argue that it is not possible for an institutional collective, such as those in my examples, to have a duty or to be blameworthy unless some member has a corresponding duty or blameworthiness *qua* occupant of some role.¹⁷

A role in an institution typically does bring with it certain institutional duties or role-responsibilities, and these might include a duty not to take into account personal considerations in discharging these responsibilities. Moreover, if an institution is morally sound, and if a person has voluntarily taken on a role in the institution, the person would have moral reasons to discharge her role responsibilities. The collectives in my examples seem to be morally sound. I therefore concede, for example, that the role responsibilities of the PM give her a moral reason not to release the prisoner and that they might also give her a moral reason not to take into account the risk to herself or her daughter. Nevertheless, once the role responsibilities of the people in my examples are taken into account, the question arises as to what these people themselves ought *all-in* to do. And the question arises as to whether they deserve blame or any other negative response for their actions once *all* excuses and justifications are taken into account, and not merely excuses that are institutionally sanctioned. The fact that the role of PM does not allow a PM to show favoritism toward her family does not entail that it is *all-in* morally wrong for the PM to release a prisoner in order to save her daughter. At most it means that there is a moral reason for the PM not to do this, which leaves it open that the PM might have a stronger moral reason to save her daughter. In general, I agree that a person’s role responsibilities in morally sound

institutions give her pro tanto moral duties. My examples show that these pro tanto duties can be overridden.

The CMA theses are claims about the possibility of a collective's having a duty or being blameworthy when no member has a duty or is blameworthy. It is no objection to argue that a collective cannot have a duty or be blameworthy unless some member has a duty or is blameworthy "qua member" if this is compatible with no member's having an all-in duty or being all-in or even pro tanto blameworthy "in her own right."

The third objection turns on theoretical considerations about the action-guiding nature of morality combined with the thesis that collectives are not independent agents. Since collectives are not independent agents, no collective can fulfill a duty unless relevant persons perform actions that constitute its doing so. Given this, and given that the point of a moral requirement is to guide action, one might think that it would be incoherent to suppose that a collective has a moral obligation unless the relevant persons in the collective have an obligation to act in ways that would constitute the collective's fulfilling its duty.¹⁸ Similarly, one point at least of blaming agents is to give them reason to do their moral duty. Given this, and given that no collective is an independent agent, one might think it would be incoherent to suppose that a collective is blameworthy for something unless at least some member is blameworthy. For it is only by giving reasons to individual persons that morality can guide collectives to fulfill their duties and thereby avoid blameworthiness.

A putative reason for a collective to do something that did not provide a reason for members to bring it about that the collective acts appropriately would have no bearing on the individual actions that would constitute the collective's action—given, again, that collectives are not independent agents. Members would have no moral reason to bring it about that the collective does what it has moral reason to do unless, by good fortune, they happened to find their own personal reasons so to act. This is a problematic picture. It suggests that if a collective could have a moral obligation without any members' having relevant corresponding obligations, then the moral obligations of collectives are not action-guiding. For in this case, since collectives are not independent agents, the moral obligation of a collective might not provide reasons to the agents whose actions could bring it about that the collective fulfills its obligation.

It is important to remember, however, that I have argued against the substantive version of the pro tanto thesis about obligation, or O-S-PT. That is, I have argued for the *pro tanto duty thesis*: when a collective has a moral obligation, relevant members have a pro tanto moral obligation to act in ways that would bring it about that the collective meets its obligation. It is in virtue of this that the moral obligation of a collective is able to guide the actions of its members by providing moral reasons to them to perform actions that would bring it about that the collective fulfills its obligation. To be sure, if I am correct, a collective can have an all-in obligation without any member's having a relevant corresponding *all-in* obligation. Yet it is enough if, as I argue, some members must have relevant

corresponding pro tanto obligations, for this means that they must have moral reasons to act accordingly. As I have explained, to have a pro tanto obligation to do something is *inter alia* to have a moral reason to do it.

One might think that the pro tanto duty thesis is a conceptual truth rather than merely a substantive normative truth. For one might think it is a conceptual truth that moral obligations are action-guiding. I am not convinced of this. I am confident that any plausible morality would be consistent with the pro tanto duty thesis. But a skeptic might deny that moral obligations are action-guiding in the sense at issue here, and I doubt that this skeptical view is conceptually false. Of course I would be delighted if someone can show that it is conceptually false. I would be delighted if the pro tanto duty thesis is a conceptual truth.

The fourth objection holds that the CMA thesis is implausibly paradoxical. First, a collective is (roughly speaking) “constituted” by its members. Because of this, it might seem that it would be paradoxical if a collective could be blameworthy without some member’s being blameworthy. But to be blameworthy is simply to be deserving of blame. It is to deserve, as it were, a negative mark in one’s moral record.¹⁹ Since a collective is distinct from any member, there should be no difficulty understanding the distinction between the moral record of a collective and that of a member and between directing blame to a collective and directing it to a member. Similarly, one can add Clearwater beach to the list of things that deserve admiration without adding any grain of sand from the beach to that list.

Second, it might seem that it would be paradoxical to suppose that a collective could have an all-in obligation without some member’s having a relevantly related all-in obligation. For, given that a collective is not an independent agent, the collective could not fulfill its obligation unless some persons acted in a relevant way. But there can be cases in which an *individual person* has an all-in obligation that she is unable to fulfill unless some other person does something that he is not all-in obligated to do. Suppose that Barbara has an obligation to return a library book and suppose that Allan promises Barbara to do this for her. Suppose, however, that Allan is prevented from keeping his promise because of an unexpected conflicting duty of greater importance, such as a duty to rescue a person who is drowning in the canal in front of the library. Given that Allan has the book in his pocket and that it will be destroyed if he jumps into the canal, it might be impossible for Barbara to fulfill her duty to return the book if Allan fulfills his duty to rescue the person. One all-in obligation or the other will not be fulfilled. The result may be that Barbara is liable to compensate the library even though Allan is not and even though Barbara could not have fulfilled her duty to the library once Allan had ruined the book. This is an unfortunate situation, but not an impossible one.

5. Conclusion

One might initially find it mysterious that a collective could have an agential moral property in a case in which no member has that property. However, my

examples show how such cases can arise as a result of moral complexity. Individual persons who are members of a collective can be in more complex moral situations than the collective, because their personal situations can give them duties that are not incumbent on the collective, or excuses that are not available to it. As a result, any pro tanto duty they have to contribute to the collective's fulfilling an all-in obligation may be overridden, or they may not be blameworthy even if the collective is. Once this is understood, I believe the CMA thesis will not seem mysterious.

It seems to me that my arguments make a strong case for three of the four conceptual CMA sub-theses, namely, O-C-all, R-C-all, and R-C-PT, and I think that the remaining conceptual thesis, O-C-PT, is also plausible. I think that my arguments also make a strong case for three of the four substantive sub-theses, namely, O-S-all, R-S-all, and R-S-PT. These theses are more controversial than their conceptual twins because the case for them rests on moral judgments about the examples I presented, and because there is room for reasonable disagreement about these judgments. For this reason, it is important to think of the examples as revisable. In order to reject the substantive sub-theses, one must be prepared to argue that *no* example can be constructed in which it is morally plausible that a collective entity has an agential moral property although no member has a relevantly corresponding agential moral property.

Of course I grant that it is generally the case that if a collective is blameworthy for something then some member or members are blameworthy. Moreover, I have argued for the pro tanto duty thesis, that if a collective has a duty then at least some members have a corresponding pro tanto duty. It is therefore sensible to have a rule of thumb that when a collective bears moral responsibility for something, some member or members bear moral responsibility in the matter. Where a collective is blameworthy, it is sensible to look for members who had a duty they failed to fulfill. The point of my argument is not to absolve executives and government officials of liability in cases in which corporations or governments or states act badly. The argument concerns what is possible, both morally and conceptually, not what is likely or plausible in the everyday world. The import of my argument for the real world is that we should avoid automatically drawing conclusions about the moral status of the members of a collective from our assessment of the moral status of the collective. To reach reliable conclusions about such matters, we need to consider in detail the moral situation faced by each of these individuals.

Versions of this paper were presented in August 2006 to the Conference on Collective Intentionality, at the University of Helsinki, and in April 2007 to the workshop on Collective Responsibility and Group Rights, at the University of Groningen. I am grateful for the comments of those in the audiences on these occasions and, especially, for detailed commentaries by Kirk Ludwig and Seumas Miller. I first presented some of the arguments in this paper, and the examples involving the Prime Minister, long ago in chapter 4 of my doctoral dissertation,

Individuals, Collectives and Moral Agency (Cornell University, 1976); the other examples first appeared in my paper, "On the Agency of Certain Collective Entities: An Argument from 'Normative Autonomy'." I presented this paper in October 2005 to the Kline Conference on Collective Rationality, at the University of Missouri. I would like to thank the participants in that conference, especially Robert Johnson, for their helpful comments and suggestions. Many years ago, David Lyons advised me in the writing of my dissertation, and I owe him special thanks for his judicious comments.

Notes

- ¹ See David Copp, "What Collectives Are: Agency, Individualism and Legal Theory," *Dialogue* 23, no. 2 (1984): 249–69.
- ² See David Copp, "Collective Actions and Secondary Actions," *American Philosophical Quarterly* 16 (1979): 177–86. There can be cases of "automated actions" such as cases in which a country's "doomsday machine" is programmed to launch missiles automatically against enemy cities when an attack is detected. Even in such cases, however, I would argue that a country acts only in virtue of the actions of some persons. I owe this example to Dan Hausman, in personal conversation.
- ³ An agent is sometimes said to be "morally responsible" for having done something where what is meant is that the agent *performed* the action and had the kind of *control* over it that is a necessary condition of being deserving of a positive or negative moral response for having done it. Call "responsibility" in this sense "agency responsibility." The thesis that collectives are not independent agents seems to entail that if a collective has agency responsibility for some action then some persons have agency responsibility for having done whatever they did that constituted the collective's so acting. The CMA thesis is compatible with this claim.
- ⁴ See Shelly Kagan, *Normative Ethics* (Boulder, CO: Westview Press, 1998), 180–81.
- ⁵ There are other uses of the expression "morally responsible." For example, a person can be said to be "morally responsible" for doing something where what is meant is that the person has a duty with respect to that thing. This sense of "morally responsible" is relevant to thesis O rather than R. Below I mention the idea of "role responsibility." In note 3 above, I mention "agency responsibility." For useful discussion, see Joel Feinberg, *Doing and Deserving* (Princeton, NJ: Princeton University Press, 1970).
- ⁶ That is, I stipulate, a proposition p is "conceptually possible" just in case it is not the case that not-p is a conceptual truth. I discuss the notion of conceptual truth in chap. 4 of my *Morality in a Natural World: Selected Essays in Metaethics* (New York: Cambridge University Press, 2007).
- ⁷ In some cases, self-defense justifies acting in a way that would otherwise be wrong. I am assuming that, in this case, because of what is at stake, the fact that she would be acting in self-defense would not justify the PM in signing the order, but it would instead be an excuse. Nothing in my argument turns on whether self-defense is a justification rather than an excuse.
- ⁸ This point was made by Robert Sugden at the Kline Conference.
- ⁹ Christian List, "The Discursive Dilemma and Public Reason," *Ethics* 116 (2006): 362–402, at 367. See also Philip Pettit, "Groups with Minds of Their Own," in *Socializing Metaphysics: The Nature of Social Reality*, ed. Frederick Schmitt (Lanham, MD: Rowman and Littlefield, 2003), 167–93. Also, Philip Pettit, "Responsibility Incorporated," *Ethics* 117 (2007): 171–201.
- ¹⁰ Robert Sugden made this suggestion at the Kline Conference.
- ¹¹ This objection was suggested by Kirk Ludwig in his commentary at the Helsinki Conference. Seumas Miller offered a similar objection in "Collective Moral Responsibility: An Individualist Account," *Midwest Studies in Philosophy* 30 (2006): 176–93, note 24.
- ¹² See Kenneth Arrow, *Social Choice and Individual Values*, 2nd ed. (New York: John Wiley and Sons, 1963). See also the literature on the discursive dilemma cited above in note 9.

- ¹³Kirk Ludwig objected in his commentary at the Helsinki Conference that it is not in general a justification for what a person does that she was simply acting in accord with rules governing her role. Torture cannot be justified in this way. But the injustice of torture is intrinsic to what torture is, whereas the wrongness in the Borderline case is procedural. It can in principle be corrected after the fact by the university. Moreover, the committee members reasonably believed that the university's standard for tenure was appropriate.
- ¹⁴See Raimo Tuomela, *The Philosophy of Sociality: The Shared Point of View* (New York: Oxford University Press, 2007), chap. 10.
- ¹⁵But see David Copp, "On the Agency of Certain Collective Entities: An Argument from 'Normative Autonomy,'" *Midwest Studies in Philosophy* 30 (2006): 194–221.
- ¹⁶See Kirk Ludwig, "Collective Intentional Behavior from the Standpoint of Semantics," *Noûs* 41 (2007): 355–93.
- ¹⁷Seumas Miller pursued this line of argument in his commentary at the Helsinki Conference.
- ¹⁸Kirk Ludwig seemed to express this worry in his commentary at the Helsinki Conference.
- ¹⁹For this idea, see Feinberg, *Doing and Deserving*.