[13] Kirk 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.

[14] See Raimo Tuomela, *The Philosophy of Sociality: The Shared Point of View* (New York: Oxford University Press, 2007), chap. 10.

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

[16] See Kirk Ludwig, "Collective Intentional Behavior from the Standpoint of Semantics," *Noûs* 41 (2007): 355–93.

[17] Seumas Miller pursued this line of argument in his commentary at the Helsinki Conference.

[18] Kirk Ludwig seemed to express this worry in his commentary at the Helsinki Conference.

[19] For this idea, see Feinberg, *Doing and Deserving*.

Against the Collective Moral Autonomy Thesis

Seumas Miller

David Copp has argued for what he calls the Collective Moral Autonomy Thesis (CMA) in his paper of the same name.[1] His argument for this thesis consists essentially in putting forward a number of examples. In this paper I want to argue that, contra Copp, his examples do not demonstrate the truth of the thesis.

What Is the Collective Moral Autonomy Thesis?

According to Copp, the CMA breaks into two parts: The first part is the thesis (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"; the second part is the thesis (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."[2]

Copp nowhere offers an account of what he terms "collectives," but he includes organized collectives such as armies, states, and universities, and unorganized collectives such as mobs, discussion groups, and the set of passengers in an airplane. Moreover, there is a certain unclarity in Copp as to the precise relationship between collectives and the individuals who (in part) constitute them; the unclarity pertains to both moral and mental properties. Surely, whatever else collectives might be, if they are to have moral obligations, and be the bearers of moral responsibility, they must be capable of deliberation and the performance of intentional actions based on their deliberations, that is, they must be fully fledged rational *agents*. Elsewhere I have argued against the strong collectivist conception that ascribes both mental and moral properties to collective entities (as opposed to ascribing them to the individual human beings who comprise them); this conception is both undesirable and unnecessary.[3] Copp, however, seems to occupy a kind of halfway house between strong collectivism and individualists such as myself who reject the ascription of either mental or moral properties to collectives per se. For Copp evidently believes that a collective could have obligations and be the bearer of moral responsibility yet not be an agent in the sense of being capable of

JOURNAL of SOCIAL PHILOSOPHY, Vol. 38 No. 3, Fall 2007, 389–409

deliberation and the performance of intentional actions based on their deliberations, that is, not be (in Copp's terminology) an independent agent. This is, at the very least, highly paradoxical; surely a substantial diminution in agency compromises the ascription of moral responsibility.

In this context it is important to stress that, in order to sustain the CMA, Copp must provide clear examples of collectives that possess *moral* obligations and are the bearers of *moral* responsibility; after all, it is uncontroversial that some collectives (e.g., corporations) have *institutional*, including legal, obligations and are the bearers of *institutional*, including legal, responsibility. Moreover, the collectives in question must not themselves be susceptible of individualist analysis such that it turns out that some individual human beings have the moral obligations and/or moral responsibility in question, albeit not the individuals Copp had initially nominated. For example, in Copp's opening example of individual bombers who (he claims) do not have an all things considered obligation not to bomb the city, there is a collective who does have such an obligation, namely, the government. However, the government is a collective susceptible of individualist analysis, for example, in terms of the Prime Minister and his/her Cabinet. Moreover, once this analysis is performed, then individual human beings (e.g., the Prime Minister) are available as possessors of the relevant all things considered moral obligations and bearers of the associated moral responsibility. Naturally, the claim here is not that with respect to any example put forward by Copp there *must be* some such individualist analysis (and, therefore, the *necessary* availability of individual human beings as possessors of the relevant moral obligations and bearers of the associated moral responsibility). Rather my claim is that with respect to each of Copp's examples, there is in fact an appropriate individualist analysis and, accordingly, no collective possessed of moral properties not possessed by its constitutive individual members (indeed, properly speaking, no collective possessed of any moral properties).

Copp's notion of moral responsibility is essentially a backward-looking notion to be applied to past and present morally significant actions—and as such associated with praise and blame—and not to be confused with the essentially forward-looking notion of moral responsibility in play in examples such as, "Parents are morally responsible for seeing to it that their children are well-fed," and, "The mechanic is morally responsible for seeing to it that the car brakes are fixed properly." Perhaps the latter forward-looking sense of moral responsibility is a special case of moral obligation, for example, "Parents have a moral *obligation* to see to it that their children are well-fed." This forward-looking notion of moral responsibility is akin to and, in some cases, derived from institutional responsibility, for example, police officers have an institutional and a moral responsibility to arrest offenders.

Further, I take it that the notion of moral responsibility in play here is pre-theoretic in character; no particular theory of moral responsibility is being presupposed. Of course, some theorists identify the concept of moral responsibility with the general concept of deserving of a positive or negative response, including but not restricted to blame or praise).[4] Doubtless, *in general* if someone is morally responsible for some act then the person will be blameworthy (in the case of a morally wrong act) and praiseworthy (in the case of a morally right act). However, it is controversial—and in my view false—that for someone to be morally responsible simply consists in their being blameworthy or praiseworthy. On the face of it a person can be morally responsible for an action without being either blameworthy or praiseworthy. Consider the case of someone who is asked if he will murder his mother and father in return for a payment of $1,000. He refuses to do so. He is morally responsible for having done the morally right thing; so he is not blameworthy. On the other hand, he is surely not worthy of praise merely because he refused to murder his own parents for a relatively small amount of money. So he is morally responsible for his action, but neither blameworthy nor praiseworthy. Now consider a different but related example. In this second example the man is not offered any money to kill his mother and father, rather he will be killed unless he kills two innocent strangers. He is under a moral obligation not to kill the two innocent strangers and, presumably, this overrides his moral right to preserve his own life; so all things considered he ought not to kill the two strangers. On the other hand, if in fact he did kill the two strangers he will be morally excused for doing so, given his moral right to preserve his own life. Self-evidently, he is not praiseworthy; after all, he did the wrong thing. But nor is he blameworthy, given that he is morally excused for performing the action. Yet the man is morally responsible for his action of killing the two strangers and, thereby, preserving his own life. For he correctly believed that the action was morally wrong all things considered, and he could have chosen to do what was morally right; he was not, for example, overcome with fear. It is just that he decided on reflection to act in accordance with his own self-interest rather than his moral obligation. At any rate, my main point is that Copp is not entitled to assume that the concept of moral responsibility is to be identified with the concept of blameworthiness/praiseworthiness.[5] So his arguments in favor of the CMA cannot rely on this assumption.

Naturally, Copp is entitled to recast the second part of the CMA in terms of blameworthiness/praiseworthiness. If he did so then he would be making the following claim: It is possible for a collective to be morally blameworthy for having done something even if no (natural) person who is a member of that entity is morally blameworthy for having done something relevantly related. Here I have omitted any reference to praiseworthiness, since in fact all of Copp's examples pertain to blameworthiness rather than praiseworthiness. In what follows I shall generally speak in terms only of blameworthiness, and refer to my recast version of the second part of the CMA as the B thesis.

Copp distinguishes between the versions of the thesis according to that which collective moral autonomy is a "moral possibility" (substantive version [S]) and which is a "conceptual possibility" (conceptual version [C]).[6] One might hold that because there are actual situations (or non-actual situations that are derived from actual situations and that are, nevertheless, realistic for us as human beings) in which consequentialism morally obliges us to do x when we are morally obliged, all things considered, to do y (where it is not possible to do x and to do y).

The precise meaning of the conceptual version of the CMA is unclear; this is in part because of the array of possible renderings of the term "conceptual possibility," both in general and in this context, for example, logical possibility, notional possibility, conceivable, intelligible, imaginable, logically consistent with all/most/some moral concepts, logically consistent with some concept of agency, and so on. Moreover, there are at least some renderings of the conceptual version of the CMA in which it is plausible, but uninteresting. Consider the following one. It is conceivable, because conceived by some theorists, that morality is such that, for example, consequentialism is true or that, for example, some collective entity has a moral obligation to do something even if no (natural) person who is a member of that entity has a relevantly related moral obligation. Given the unclarity surrounding the conceptual version of the CMA, and the consequent possibility that it is a plausible but uninteresting thesis, I will concern myself in this paper only with the substantive version of the CMA.

Copp further distinguishes between more and less stringent versions of the CMA and, specifically, between "pro tanto" (PT) and "all things considered" (All) versions of the thesis.[7] According to the less stringent "all things considered" version, it is possible for a collective to have an all things considered moral obligation (or moral responsibility) even if no member has a relevantly related all things considered moral obligation (or moral responsibility). According to the more stringent "pro tanto" version, it is possible for a collective to have an all things considered moral obligation (or moral responsibility) even if no member has even a relevantly related pro tanto moral obligation (or moral responsibility).

The distinction between a pro tanto and an all things considered moral *obligation* is reasonably clear. However, the distinction between a pro tanto and an all things considered moral *responsibility* is much less clear. Certainly, there is a distinction between pro tanto and all things considered moral responsibilities in the forward-looking sense of responsibility. The police officer has a pro tanto moral responsibility to arrest the young offender, but perhaps an all things considered moral responsibility not to do so given that the youth will end up in jail and become a hardened offender. However, as we have seen, this forward-looking notion of responsibility is not the one in question in Copp's paper.

Assume that with respect to two mutually exclusive actions, x and y, person A has a pro tanto moral obligation to do x and an all things considered moral obligation to do y, and that A does x and does not do y. A is morally responsible for doing x and failing to discharge his all things considered moral obligation to do y. This is the backward-looking notion of responsibility that is in question. However, A is neither pro tanto nor all things considered morally responsible for doing x and not doing y (or for doing x or for not doing y); rather he is just morally

pro tanto and all things considered does not apply to moral responsibility. At the very least, Copp owes us an account of how this distinction is to be understood in relation to moral responsibility.

Now let us consider "responsibility" in the sense of blameworthiness/praiseworthiness. Agent A is blameworthy, let us assume, for doing x and failing to do y. However, A is neither pro tanto nor all things considered blameworthy for doing x and failing to do y (or for doing x or for failing to do y); rather he is just blameworthy for doing x and failing to do y. It seems that the distinction between pro tanto and all things considered does not apply to blameworthiness/praiseworthiness.[8] At the very least, Copp owes us an account of how this distinction is to be understood in relation to blameworthiness/praiseworthiness (and the related, more general notion of deserving of a negative/positive response that he sometimes makes use of).

It might be that what Copp has in mind in relation to the distinction between pro tanto and all things considered as it applied to blameworthiness/praiseworthiness is certain moral features of the relevant moral agent (as opposed to the rightness or wrongness of the agent's actions). For example, an agent is pro tanto blameworthy for performing a morally wrong act (an action that is morally wrong all things considered), but is not blameworthy all things considered if he or she turns out to have performed the morally wrong action under duress. If so, this needs to be made a good deal clearer than it is.

In the light of the three sets of distinctions between obligation/responsibility, substantive/conceptual and all things considered/pro tanto, Copp generates eight versions of the CMA.[9] They are as follows: (1) O-C-All; (2) O-C-PT; (3) O-S-All; (4) O-S-PT; (5) R-C-All; (6) R-C-PT; (7) R-S-All); and (8) R-S-PT. Copp argues for seven of the above sub-theses of the CMA, and against one of them, namely, O-S-PT (sub-thesis [4]). It might be thought that if Copp rejects O-S-PT he must also reject the other substantive pro tanto R-S-PT (sub-thesis [8]). That is, it could be argued that if it is not possible for a collective to have an all things considered obligation unless at least one of its members has a relevantly related pro tanto obligation, then it is not possible for the collective to be morally responsible (and blameworthy) in relation to its discharging or failure to discharge its all things considered obligation unless at least one of its members is morally responsible (and blameworthy) in relation to his/her discharging or failure to discharge his/her relevantly related pro tanto obligation. However, as we shall see, Copp's thought is that the member's pro tanto obligation (but not the collective's all things considered obligation) might be overridden by some other moral consideration, thereby negating the moral responsibility (and blameworthiness) of the member's failure to discharge his or her pro tanto obligation.

Since I have already indicated my general concern in relation to the conceptual version of the CMA, I will not consider sub-theses (1), (2), (5), and (6), these all being variations on the conceptual version of the CMA. Moreover, since Copp himself rejects sub-thesis (4), I will not consider sub-thesis (4)

(7) (R-S-All), and (8) (R-S-PT). However, by my lights there are two versions of each of sub-theses (7) and (8); one version corresponding to moral responsibility and one to blameworthiness. Let me fully and explicitly state sub-theses (3), (7), and (8) before proceeding to consider the examples offered by Copp in support of them.[10]

Sub-thesis (3) (O-S-All): It is morally possible for a collective to have an all things considered moral obligation even if no member has a relevantly related all things considered moral obligation.

Sub-thesis (7) (R-S-All): It is morally possible for a collective to be morally responsible (or blameworthy/praiseworthy) all things considered for having done something even if no member is morally responsible (blameworthy/praiseworthy) all things considered for having done something relevantly related. Here I reiterate that it is unclear to me how the distinction between pro tanto and all things considered is to be understood in sub-thesis (7). Moreover, as noted above, Copp's examples all pertain to blameworthiness; none are concerned with praiseworthiness. Accordingly, I offer the following two renderings of sub-thesis (7); the first pertains to moral responsibility, and the second to blameworthiness.

(7a) (R-S-All): It is morally possible for a collective to be morally responsible for having discharged or failed to discharge an all things considered obligation even if no member is morally responsible for having discharged or failed to discharge that all things considered moral obligation (or otherwise relevantly related all things considered moral obligation).

(7b) (B-S-All): It is morally possible for a collective to be morally blameworthy for having failed to discharge an all things considered obligation even if no member is morally blameworthy for having failed to discharge that all things considered moral obligation (or otherwise relevantly related all things considered moral obligation).

Sub-thesis (8) (R-S-PT): It is morally possible for a collective to be morally responsible (blameworthy/praiseworthy) for having done something even if no member is pro tanto morally responsible (blameworthy/praiseworthy) for having done something relevantly related. Here I reiterate that it is unclear to me how the distinction between pro tanto and all things considered is to be understood in sub-thesis (8). Moreover, Copp's examples all pertain to blameworthiness; none are concerned with praiseworthiness. Accordingly, I offer the following two renderings of thesis (8); the first pertains to moral responsibility, and the second to blameworthiness.

(8a) (R-S-PT): It is morally possible for a collective to be morally responsible for having discharged or failed to discharge an all things considered obligation even if no member is morally responsible for having discharged or failed to discharge a relevantly related pro tanto moral obligation.

(8b) (B-S-PT): It is morally possible for a collective to be morally blameworthy for having failed to discharge an all things considered obligation even if no member is morally blameworthy for having failed to discharge a relevantly related

Example 1: The Kidnapping: The Prime Minister

Copp offers this example in support of sub-thesis (7). In this example the Prime Minister (PM), say, Gough Whitlam (a former PM of Australia), is taken hostage by an outlaw group who threaten to kill the PM unless the PM authorizes the release of a dangerous prisoner.[11] The government morally ought, all things considered, to refuse to release the prisoner, since if this demand is acceded to many other groups will adopt similar tactics. Moreover, according to Copp, the PM has a pro tanto moral obligation not to release the prisoner. However, notwithstanding the PM's (claimed) pro tanto moral obligation not to release the prisoner, the PM has an excuse for doing so, namely, self-preservation. As Copp plausibly claims, "A person can at least sometimes be excused for acting wrongly in order to save his own life"; and this is one of those times.[12]

Elsewhere I have considered this sort of case and invoked the distinction between, relativizing it to this example, Gough Whitlam qua PM and Gough Whitlam qua some other role occupant or Gough Whitlam simpliciter, that is, Whitlam the individual human being.[13] (See below for more on these distinctions.) Qua PM in a government obliged to pursue the public interest in the context of present and possible future hostage situations, Whitlam has an *all things considered* moral obligation not to release the prisoner; on the other hand, as a human being whose life is presently threatened, Whitlam has a moral right to self-preservation. Now Whitlam is both PM (and as such a member of the government) and a human being whose life is threatened; accordingly, the question arises as to what he Whitlam should do, all things considered. Now let us assume (as in effect Copp does) that whereas Whitlam's right to self-preservation does not override his obligation not to release the prisoner, it is morally excusable for him to release the prisoner in these circumstances.

On the analysis I have just offered, contra Copp, the example does not support sub-thesis (7) (in the sense of [7a]). Assume that Whitlam in his capacity as PM in fact releases the prisoner (by signing the appropriate release orders). Given that Whitlam correctly believed that the all things considered moral obligation of the government—and, therefore, his all things considered moral obligation qua PM—was to refuse to release the prisoner, Whitlam qua PM did what he knew to be morally wrong. Moreover, he could have chosen to do what was the morally right thing for him as PM to do. So Whitlam is morally responsible (as opposed to blameworthy) for committing what it was morally wrong for him as PM to do, that is, release the prisoner.

What about the government of which the PM is a member? Is the government morally responsible? In Copp's example, aside from the PM, no member of the government (or anyone else for that matter) could have done anything to prevent the prisoner being released; aside from anything else, they did not know, and could not reasonably have known, what was going on. Accordingly, the individual members of the (executive arm of the) government, that is, the Cabinet Ministers, they are not morally responsible (and, therefore, not blameworthy). (For simplicity, I will at this point treat the government as comprised of the PM and Cabinet rather than, for example, all the government members of the legislature.) So, contra Copp, the government per se is not morally responsible (and, therefore, not blameworthy); rather the PM alone bears moral responsibility for the act. Accordingly, perhaps the PM should be dismissed from office by, say, the Queen or, in the case of Australia, the Governor General (as indeed happened in the case of Whitlam, albeit for very different reasons). At any rate, the example as analyzed does not support sub-thesis (7) (in either the [7a] or the [7b] version). For, contra the assumption in (7) (in both the [7a] and the [7b] version), the collective (i.e., the government) is not morally responsible (and, therefore, not blameworthy) for having released the prisoner and, again contra sub-thesis (7) (the [7a] version), a member of the collective (the PM) *is* morally responsible (as opposed to blameworthy) for having released the prisoner, that is, for having done something relevantly related.

This analysis is consistent with Whitlam's action of releasing the prisoner being morally excusable (and, therefore, Whitlam not being blameworthy), given that it was an act of self-preservation. (Again, note the distinction between being morally responsible for an action and being morally blameworthy for performing it.) As we have seen, Whitlam failed to discharge his all things considered moral obligation not to release the prisoner and was morally responsible for failing to do so. However, that he acted to preserve his life constitutes a moral excuse. So, pace Copp, Whitlam committed a morally excusable morally wrong act, that is, Whitlam is not blameworthy.

The analysis is also consistent with the government, that is, the PM and the members of the Cabinet, being held *institutionally responsible*—as opposed to morally responsible—for the release of the prisoner. It might be that there is an institutional arrangement in place to the effect that in such cases the PM is necessarily to be taken to be acting on behalf of the government, that is, the Cabinet including the PM; accordingly, the entire membership of the Cabinet as well as the PM might be required to resign (or otherwise held to account for failing to discharge their collective institutional responsibility). Here I note that institutional responsibilities and liabilities do not necessarily track moral responsibilities and liabilities. Consider in this connection strict liability in law; an employer might be held strictly liable for the behavior of her employee, notwithstanding the morally exemplary behavior and attitude of the employer.

Indeed, even if we make the quite different (and implausible) assumption, namely, the assumption that the government simply is the PM (Copp seems at one point to be suggesting this), that is, Whitlam qua PM is the government, then the analysis is also consistent with the government being morally responsible for failing to discharge its all things considered moral obligation.[14] For the government, that is, Whitlam qua PM, is morally responsible for failing to discharge Whitlam qua PM's all things considered moral obligation. Is Whitlam qua PM

be held liable for his failure to do his duty, for example, tried and sentenced to imprisonment. However, Whitlam the PM is also Whitlam the man whose life is under threat, and it is unclear whether the concept of blame can sufficiently disaggregate these aspects of a single moral agent; blameworthiness, unlike moral obligation, qualifies agents not actions. But for the same reason it is equally unclear whether the government (i.e., Whitlam qua PM) is blameworthy. At any rate, the general point to be made here is that on this rendering of the example there is no collective entity possessed of any moral property not possessed by an individual human person; this is because the collective entity, namely, the government, *is* an individual person, namely, Whitlam qua PM.

A further issue raised by the example is whether the PM (some member of the collective) has in fact a pro tanto moral obligation not to release the prisoner, or perhaps not to sign the release order. Copp explicitly claims that "it seems plausible that the PM has an obligation not to sign the order."[15] Pace Copp, I also hold that the PM has a pro tanto obligation not to release the prisoner.[16] However, by contrast with Copp, I have suggested that the obligation of Whitlam qua PM not to release the prisoner is an all things considered obligation; so in the example as described I believe it is *both* a pro tanto and an all things considered obligation. But this brings us to the question of the acceptability of the distinction I have invoked between Whitlam the human being and Whitlam qua PM.

What is required now is to make clear the notion of someone, say Whitlam, having a moral obligation qua PM. Let me explain the notion of moral obligation (including moral responsibility in the forward-looking sense) qua role occupant by first introducing the related notion of acting qua member of a group.

Example 2: The Kidnapping: The Prime Minister's Daughter: Moral Obligations and Responsibilities Qua Role Occupant

Copp offers his second example in support of sub-thesis (3). Moreover, this second example conveniently exemplifies the notion of an obligation (including responsibility in the forward-looking sense) qua role occupant. This second example is a version of his first one. However, in this second version it is not the life of the PM that is under threat but rather that of his daughter; it is the PM's daughter who has been kidnapped.

The moral obligation of Whitlam qua father is to release the prisoner and save his daughter. (Here the institutional *role* of father might be invoked; alternatively, one might invoke a father's particularistic and emotionally held moral obligation to his daughter.) On the other hand, the institutionally based moral obligation of Whitlam qua PM is not to release the prisoner, but rather to sacrifice the life of the citizen who also happens to be his daughter. If the PM opts to release the prisoner, he is morally responsible for failing to do his duty as PM; presumably, at the very least he would need to resign. (And if he is acting on behalf of the government they are all jointly institutionally and morally responsible; accordingly, the other members of the Cabinet should also resign.) If Whitlam chooses not to release the prisoner, then he is morally responsible for failing to do his duty qua father to save the life of his daughter. He has to choose between two evils, and whichever option he chooses, he is morally responsible for it. But I am getting ahead of myself. The challenge at this point is to articulate the notion of moral obligation (including responsibility in the forward-looking sense) qua role occupant and do so following on an elaboration of the related notion of acting qua member of a group.

The notion of acting qua member of a group is often quite straightforward since the group can be defined in terms of the collective end or ends which the group of individuals is pursuing. Consider a group of individuals building a house.[17] Person A is building a wall, person B the roof, person C the foundations, and so on. To say of person A that he is acting qua member of this group is just to say that his action of building the wall is an action directed toward the collective end that he and the other members of the group are seeking to realize, namely, a built house.

Notice that the same set of individuals could be engaged in different collective projects. Suppose persons A, B, C, and so on in our above example are engaged not only in building a house but also—during their holidays—in building a sailing boat. Assume that A is building the masts, B the cabin, C the bow, and so on. To say of A that he is acting qua member of this group is just to say that his action of building the masts is an action directed toward the collective end that he and the other members of the group are seeking to realize, namely, a built boat. Accordingly, one and the same person, A, is acting both as a member of the "house building group" (G1) and as a member of the "boat building group" (G2). Indeed, since A, B, C and so on are all members of, and indeed the only members of, each of these two groups, the membership of G1 is identical with the membership of G2.

Moreover, when A is building the wall he is acting qua member of G1, and when he is building the mast he is acting qua member of G2. For A to be acting qua member of G1 is for A to be pursuing—jointly with B, C, and so on—the collective end of building the house; for A to be acting qua member of G2 is for A to be pursuing—jointly with B, C, and so on—the collective end of building the boat.

Further, let us suppose that G1 and G2 each have to create and comply with a budget; G1 has a budget for the house, and G2 has a budget for the boat. The members of G1 and G2 know that they must buy materials for the house and the boat (respectively) and do so within the respective budgets. Assume that A, B, C, and so on allocate $50,000 to pay for bricks for the house. This is a joint action. Moreover, this joint action is one that A, B, C, and so on have performed qua members of G1. G1 is individuated by recourse to the collective end of building the house, and the proximate (collective) end of buying bricks is tied to that group, G1, and to its ultimate end of building a house. Accordingly, A, B, C, and so on

the collective end of building a boat, and A, B, C, and so on do not qua members of G2 have any plans to build their boat from bricks!

Naturally, many groups are unified in more complicated ways than in our above example. Consider the members of an organization such as a government; each person occupies an organizational role. There is the PM, the Minister for Defense, the Minister for Education, and so on. Moreover, each of these roles sits at the apex of an (hierarchical) organizational structure comprised of a variety of public service roles. Elsewhere I have elaborated an account of the joint actions performed by such organizational role occupants.[18] Specifically, I have introduced the concept of a *layered structure of joint actions*. Roughly speaking, a layered structure of joint actions is a set of joint actions, each of which is directed to a further collective end; so it is a macro-joint action comprised of a set of constituent micro-joint actions. I have also supplemented this account by recourse to concepts of conventions, social norms and the like, and especially by recourse to the explicitly normative notions of rights, obligations, and duties that are attached to, and in part definitive of, many organizational roles.[19] It is not simply that organizational role occupants *regularly* jointly act in certain ways in preference to others, or in preference to acting entirely individualistically; rather they have institutional duties to so act and—in the case of hierarchical organizations—institutional rights to instruct others to act in certain ways. At any rate, the point to be made here is that my account of the notion of acting qua member of a group in terms of acting in accordance with collective ends can be, and should be, complicated and supplemented by the normative notions of rights and duties in order to accommodate various different kinds of acting qua member of an organized group, including acting in hierarchical roles such as that of PM. So role occupants such as Whitlam take on the tasks definitive of the role. More specifically, they take on the institutional rights and duties definitive of the role, and some of these institutional rights and duties are also moral rights and obligations. Accordingly, it makes sense to say of Whitlam that he has this and that moral obligation qua PM, but not qua father of the bride.

Moreover, it makes sense to ascribe to a role occupant (qua role occupant) moral obligations some of which are pro tanto obligations and others of which are all things considered obligations. Indeed, such is the case with Whitlam qua PM (in both versions of Copp's example).

So there are moral obligations and (forward-looking) moral responsibilities that an institutional role occupant might possess qua role occupant. However, in addition to these there are moral obligations and (forward-looking) moral responsibilities (both pro tanto and all things considered ones) that the occupant of the role might possess, but that do not arise by virtue of his/her role occupancy; rather they arise by virtue of some other institutional role the person happens to occupy or by virtue of something other than occupancy of an institutional role, for example, by virtue of being a human being. Accordingly, we need the notion of a bearer of moral obligations such that the obligations in question include, but are idea of an individual human being who might or might not occupy any given institutional role. If we use proper names, such as "Whitlam," to refer to individual human beings (names as so-called rigid designators), then we can ascribe moral obligations (both role-based and non-role-based obligations) to Whitlam *simpliciter* (and to other thus named individual human beings). Thus, it is Whitlam simpliciter—as opposed to Whitlam qua PM or Whitlam qua father—who has a moral obligation not to kill the innocent. And it is Whitlam simpliciter who might have a role-based moral obligation at time t1 but not at time t2, since Whitlam simpliciter is not *defined* as the occupant of this (or these) institutional role(s).

A moral obligation all things considered that a person has qua role occupant might conflict with a moral obligation that the person has qua some other role occupant, or with some non-role-based moral obligation. Accordingly, there might arise the further question as to what the moral obligation of the person all things considered is. Indeed, this is precisely what is taking place in Copp's "The Kidnapping: The Prime Minister's Daughter" example (as well as in his "The Kidnapping: The Prime Minister" version of this example). So Whitlam qua PM has a moral obligation all things considered not to release the prisoner. However, Whitlam qua father has a moral obligation to save the life of his daughter. Accordingly, Whitlam *simpliciter* now has a further moral question in relation to what he ought to do all things considered. Copp claims that the answer to this further question—what ought Whitlam to do, all things considered—is that Whitlam ought to release the prisoner. I don't want to dispute this. (Although it is not entirely clear to me why when Whitlam's own life is threatened [as in version 1 of the example], his obligation to refuse to release the prisoner is not overridden, whereas when the life of his daughter is threatened his obligation to refuse to release the prisoner is overridden.) Rather I want to point out that this is different from the answer to the same question addressed to Whitlam qua PM. Moreover, I want to insist that there are two distinguishable moral questions in play, both of which are questions about what ought to be done all things considered. However, the first question is addressed to Whitlam qua PM, the second to Whitlam simpliciter.

Let me now explicitly turn to sub-thesis (3), since Copp claims that the "The Kidnapping: The Prime Minister's Daughter" example supports this sub-thesis.

The reply to Copp is that the government has an all things considered moral obligation not to release the prisoner, but Whitlam qua member of the government (i.e., qua PM) *has* a relevantly related all things considered moral obligation; indeed, precisely the same all things considered moral obligation, namely, the obligation not to release the prisoner. So the example does not support sub-thesis (3).

Copp offers the next scenario, "The Prison Board," in support of sub-thesis (7).

Example 3: The Prison Board: Joint Institutional Mechanisms

Copp's Kidnapping examples have indirectly raised the issue of the collective

governments. His final two examples, namely, the "The Prison Board" example and the "The Tenure Committee" example, raise this issue in a more direct form. Accordingly, prior to consideration of these two examples I want to elaborate my account of what I take to be one of the key required notions in relation to these matters, namely, the notion of a joint institutional mechanism.[20]

Examples of joint mechanisms are the device of tossing a coin to resolve a dispute and voting to elect a candidate to office.

Joint mechanisms consist of: (i) a complex of differentiated, but interlocking, actions (the input to the mechanism); (ii) the result of the performance of those actions (the output of the mechanism); and (iii) the mechanism itself. Thus, a given agent might vote for a candidate. He will do so only if others also vote. Additionally, there is the action of the candidates, namely, that they present themselves as candidates. That they present themselves as candidates is (in part) constitutive of the input to the voting mechanism; voters vote *for candidates*. So there is interlocking and differentiated action (the input). Further, there is some result (as opposed to consequence) of the joint action; the joint action consisting of the actions of putting oneself forward as a candidate and the actions of voting. The result is that some candidate, say, Jones is voted in (the output). That there is a result is (in part) constitutive of the mechanism. That to receive the most number of votes is to be voted in is (in part) constitutive of the voting mechanism. Moreover, that Jones is voted in is not a collective end of all the voters. (On the other hand, it is a collective end of those who voted for Jones.) However, that the one who gets the most votes—whoever that happens to be—is voted in is a collective end of all the voters.

Some joint mechanisms are *institutional* mechanisms.[21] One such example is voting for public officials. Another has also been mentioned, namely, a parliamentary cabinet making a decision. Here the Cabinet makes a decision that each member has agreed to. Naturally, some members of Cabinet may privately disagree with the decision. However, ultimately they all agree to publicly accept that decision, which is to say that each agrees qua member of the Cabinet. Here the notion of acting qua occupant of a role is simply that of performing the tasks definitive of the role (including the joint tasks), conforming to the conventions and regulations that constrain the tasks to be undertaken, and pursuing the purposes or ends of the role (including the collective ends). Accordingly, each member of the Cabinet is *institutionally* responsible, jointly with the others, for, say, the Cabinet's decision to raise taxes. Moreover, given the moral significance attaching to raising taxes, each member of the Cabinet is *morally* responsible, jointly with the others, for the decision to raise taxes. This is so, notwithstanding the fact that some members might believe that it is immoral to raise taxes and, as a consequence, might have strenuously argued against this course of action. These members face a moral dilemma. They can resolve the dilemma by making the decision to set aside their privately held belief in favor of the majority view of the Cabinet; in effect, qua Cabinet member, each individually endorses the decision that taxes be the Cabinet, accept the decision to raise taxes, then her position qua member of the Cabinet may become untenable; if so, then she must resign from the Cabinet.

Indeed, it might be that most, or even all, of the members of the Cabinet believe that it is immoral to raise taxes, and yet all, or most, vote for taxes to be raised. However, even if this is the case, each Cabinet member remains, qua Cabinet member, *institutionally* responsible, jointly with the other members, for the decision to raise taxes; and, given the moral significance of taxes being raised, each Cabinet member is likewise *morally* responsible, jointly with the others, for this outcome.

David Copp's "The Prison Board" example constitutes an interestingly different version of this kind of case. In his example, all (Bob, Alice, and Ted) but one member (Carol) of the Board of Governors of a prison vote against a proposal that, morally speaking, they ought to have endorsed. The twist is that each member of the Board had a supposed excuse for voting as they did. Copp's principal concern here is with moral responsibility; specifically, he wants to endorse the conclusion that the Board is morally responsible all things considered for the outcome, but none of the individual members are morally responsible (all things considered) for it.[22] That is, Copp takes this example to support sub-thesis (7). (Recall that there are two versions of sub-thesis [7]: the moral responsibility version [7a] and the blameworthy version [7b]).

By my lights, if the members of the Board meet the conditions for being jointly institutionally responsible for the decision, and the decision is morally significant, then—other things being equal—they are, qua Board members, jointly morally responsible for the outcome. Assume that the decision is morally significant. What of their institutionally based moral responsibility?

Here there are a variety of putative excuses. Suppose a member argues strongly against the joint decision (Carol), or was not in attendance (due to an emergency) (Bob). If so, then the member can dissociate herself from the Board's joint decision, for example, by resigning or by stepping down pending the reversal of the decision or issuing a public statement dissociating herself/himself from the decision. If the member does not dissociate herself from the decision then, in effect, she has accepted it and, therefore, the institutional and moral responsibility that flows from it. In the example, as described, Carol and Bob each accepted the Board's decision. Accordingly, Carol and Bob are each institutionally and morally responsible for the adverse outcome of the Board's decision.

Second, suppose the Board's decision-making process has a design fault, for example, members vote by pushing a button on a panel, but the recording device does not count the votes correctly. Presumably, the members of the Board are not institutionally responsible for ensuring the good working order of such equipment and, therefore, not morally responsible for voting outcomes that are incorrectly generated in this way. This kind of design fault is not involved in this example as it is described.

Third, suppose two members (Ted and Alice) have an arrangement such that research relies on the other person's research. Suppose that this arrangement goes awry, and each of these members makes a bad decision based on the false belief that the other had done the research. Presumably, this is not an *institutionally* acceptable excuse, since it is up to the member to see to it that her decision is based on the actual results of good research. Moreover, each could have done so, for example, if each had done his/her own research or checked to ensure what the actual research results were. So each is institutionally and morally responsible for the Board's decision. Accordingly, Alice and Ted are each institutionally and morally responsible for the adverse outcome of the Board's decision.

I conclude that in the "The Prison Board" example, as described, the Board is morally responsible for the adverse outcome and, contra Copp, each member is morally responsible for the adverse outcome. So the example does not support sub-thesis (7) in its version (7a). What about sub-thesis (7) in its version (7b)? Is the Board blameworthy but none of its members morally blameworthy? Presumably, *some* of the members of the Board are blameworthy. For example, Alice and Ted are blameworthy in relation to the Board's voting outcome since they should not have had their informal practice of relying on the research of the other. Accordingly, the example does not support sub-thesis (7), even in its version (7b), for it is not the case, as required by sub-thesis (7b), that *none* of the members of the Board is blameworthy.

It is consistent with this to hold that Carol, in particular, is not blameworthy; after all, she voted in favor of the increased security measures and it might be that all things considered she ought not to dissociate herself from the decision; for example, perhaps in the scheme of things the decision in question is not all that important and she cannot reasonably dissociate herself from every decision the Board makes that she voted against.

Naturally, the example could be re-described in such a way that each member of the Board did not have moral responsibility all things considered for the adverse outcome. One such re-description is that offered by a design fault of the kind mentioned above. However, in this kind of example the collective, the Board, is not institutionally or morally responsible. So it does not support sub-thesis (7) (in either its [7a] or [7b] version). A second such re-description is one in which each of the members of the Board has a valid *moral* excuse for his or her actions, but none escapes institutional responsibility. On my account, the Board and its members would not be morally blameworthy, even though they were institutionally responsible. Would they, nevertheless, be morally responsible? Perhaps, perhaps not; this depends on the nature of the excuse. Some excuses, such as in the case of the PM in Copp's earlier examples, render a person blameless without negating the person's moral responsibility; other excuses negate moral responsibility and, therefore, render the person blameless as well.

My general suggestion here in relation to blameworthiness, in particular, is that to the extent that we are inclined not to blame any of the individual members of collective entities for some untoward outcome of their actions, we are inclined if we really believe each has an acceptable moral excuse—even Ted and Alice—then we will ascribe institutional and moral responsibility to the Board, but stop short of blaming it. At any rate, my initial point stands, namely, that Copp has not provided, as he needs to, a clear example in which the Board is morally blameworthy, but none of its members are.

Copp's final scenario, "The Tenure Committee," is in some ways the most important since not only does he offer it in support of one or other of sub-theses (3) and (7) (as was the case with his previous examples) but he offers it in support of sub-thesis (8).[23] That is, he argues that the example supports all three sub-theses.

Example 4: The Tenure Committee

In Copp's scenario a university's tenure committee consists of three persons, A, B, and C, and it has to determine whether or not candidate Borderline should be granted tenure. The university's standard for tenure is excellence in each of the three areas of research, teaching, and service. The university's procedure (P1) for determining whether this standard has been met is for each member of the committee to vote on whether she/he believes the candidate is excellent in all three areas; if there is a majority in favor then tenure is granted, if not then it is denied. However, in this instance A, B, and C each believe there are only two areas in which Borderline is excellent; so each votes to deny tenure and, consequently, the university denies tenure. The twist is that with respect to each area (e.g., research) a majority believes Borderline is excellent. Accordingly, if the procedure had been to vote on each area (P2), then Borderline would have been deemed to be excellent in all three areas and would have been granted tenure. Let us assume, as Copp does, that P2 is a good procedure for ensuring the university's standard for tenure is met and that P1 is flawed.

Copp argues that the example, as described, supports sub-theses (3), (7), and (8). I will argue that it does not.

A preliminary point here is that there is some confusion as to what the collective in question is, there being two collectives in the example. The tenure committee is one collective having A, B, and C as its members. The university is a second collective of which the tenure committee is merely a component, for example, perhaps it is a sub-committee of the governing body of the university, that is, of, say, the university senate. This is important since sub-theses (3), (7), and (8) are cast in terms of the moral obligation, moral responsibility (versions [7a] and [8a]), and blameworthiness (versions [7b] and [8b]) of the collective, on the one hand, and the moral obligation, moral responsibility, and blameworthiness of the membership of that *very same* collective, on the other. Given that these collectives are not identical with one another, one might well expect that there would be differences between the moral obligation, moral responsibility, and blameworthiness of the university and the moral obligation, moral responsibility,

to accept what Copp has to say about these respective moral obligations and attendant responsibility and blameworthiness, it is far from clear that the example would provide any support for sub-theses (3), (7), and (8).

As it happens, I do not accept a number of Copp's key claims in relation to the moral obligations, and attendant responsibility and blameworthiness of the university, and those of the members of the tenure committee.

Since P1 is a flawed procedure that is delivering an unjust outcome to Borderline, the university has a pro tanto moral obligation to adopt a new procedure (e.g., P2) and overturn the tenure committee's decision to deny tenure to Borderline. For the same reason the tenure committee and the individual members of the tenure committee, have a pro tanto moral obligation not to deny tenure to Borderline. On the other hand, since at the time in question P1 is the officially sanctioned procedure, the university, the tenure committee and its members have a pro tanto institutional obligation—and consequent moral obligation—to comply with it.

Let us assume, pace Copp, that the university has a moral obligation, all things considered, not to refuse tenure to Borderline; the pro tanto obligation to avoid an unjust outcome based on a flawed procedure overrides the pro tanto obligation to comply with the existing official procedure. What of the members of the tenure committee? Other things being equal, each has a moral obligation, all things considered, not to refuse tenure to Borderline; the members each have the same pro tanto moral obligations as the university and there is, thus far, every reason to believe that these moral obligations have the same moral weight for the members of the committee as they have for the university. In particular, the fact that the P1 is the officially sanctioned procedure does not absolve the members of the tenure committee from any moral responsibility for acting in accordance with that procedure, given that the procedure is a flawed one that is producing an unjust outcome, that is, given the assumption made by Copp in constructing "The Tenure Committee" scenario.

The conclusion is that the moral obligations, both pro tanto and all things considered, of the university and of the members of the tenure committee mirror one another. So, contra Copp, the example does not support sub-thesis (3). It is not the case that the all things considered obligation of the university is *not* to deny tenure *and* that some member(s) of the tenure committee has an all things considered moral obligation to deny tenure.

Both the university and the members of the tenure committee have a moral obligation, all things considered, not to deny tenure to Borderline. Consequently, if either did so, then each would be morally responsible for failing to discharge their respective moral obligations. Accordingly, on this analysis, the example does not support sub-thesis (7) (specifically, version [7a]). For while the university is morally responsible for denying Borderline tenure, some members—indeed, all members—of the tenure committee are also morally responsible in this respect. The same point holds in respect of the blameworthiness of the university and of the members of the tenure committee, that is, the example does not support

What of sub-thesis (8)? Copp's claim here is that the university is morally responsible (all things considered) for denying Borderline tenure, but no member of the tenure committee is morally responsible (pro tanto) for denying Borderline tenure. As suggested above, this claim needs to be recast in terms of sub-theses (8a) and (8b). I take it that Copp, himself, explicitly endorses version (8b), as opposed to version (8a) of his tenure committee example. If so, then he ought to believe that the example also supports version (8a), for if someone is blameworthy for some adverse outcome then he is surely also morally responsible for it. On the other hand, if the example does not support version (8a) then it cannot support version (8b), for if someone is not morally responsible for some adverse outcome then presumably they cannot be morally blameworthy for it. At any rate, for the sake of completeness I will consider the example in relation to both versions.

In relation to version (8a) there are two salient claims that need to be true: (i) the university denied Borderline tenure when it could have done otherwise, and was thereby *morally responsible* for failing to discharge its all things considered moral obligation not to deny Borderline tenure; and (ii) the members of the tenure committee are *not* morally responsible for denying Borderline tenure since, although they did so and could have done otherwise, they did not have a pro tanto obligation not to deny tenure. On my analysis of the example, (ii) is false since the members of the tenure committee had a pro tanto obligation not to deny tenure by virtue of P1 being a flawed procedure producing an unjust outcome. Therefore, the example does not support version (8a) of sub-thesis (8). What of version (8b)?

In relation to version (8b) there are two salient claims that need to be true: (i) the university denied Borderline tenure when it could have done otherwise, and was thereby *blameworthy* for failing to discharge its all things considered moral obligation not to deny Borderline tenure; and (ii) the members of the tenure committee are *not* blameworthy for denying Borderline tenure since, although they did so and could have done otherwise, they did not have a pro tanto obligation not to deny tenure. As already indicated, on my analysis of the example, (ii) is false since the members of the tenure committee had a pro tanto obligation not to deny tenure by virtue of P1 being a flawed procedure producing an unjust outcome. Therefore, the example does not support version (8b) of sub-thesis (8). And there is the further point that if it does not support version (8a) of sub-thesis (8), then it cannot support version (8b).

Now it might be that the two conflicting pro tanto obligations in question do not have the same moral weight for the members of the tenure committee than they have for the university. For example, the university might have an obligation to overturn the results of its procedures when they are revealed to be flawed; whereas this might not be an obligation, or might not be so strong an obligation, for the members of the tenure committee. However, if so this is because the university and the tenure committee are for certain purposes different collectives with different, albeit overlapping, sets of obligations and responsibilities; and this is so for the purposes of the example. Specifically, procedures for awarding tenure are not

responsibility for overturning unfair decisions properly made in accordance with the procedures; rather these matters are the responsibility of other committees of the university. So the part of the university that is responsible for the procedures being what they are—or for overturning unfair decisions—is not the same part of the university that is responsible for making individual decisions as to tenure on the basis of those procedures. But in that case the example is not an example of the kind that Copp needs because it is not an example which involves a collective and the members of that very same collective. On the other hand, if there is no committee or other sub-element that is responsible for establishing the procedures for tenure or for overturning the procedurally correct but unfair decisions of the tenure committee, then presumably the university leadership, for example, the Senate or the Vice-Chancellor, is morally responsible for failing to establish such committees or mechanisms. Moreover, in these circumstances the tenure committee is morally required to subject the procedures it is using to close moral scrutiny and also to see itself as the court of last appeal; after all, if it does not, who will?

As it happens, Copp focuses on a particular alleged moral difference between the university and the members of the tenure committee.[24] He argues that each member of the tenure committee considered individually did not have an all things considered obligation not to deny tenure, notwithstanding that each had a pro tanto reason not to deny tenure to Borderline (the same pro tanto reason the university had, namely, that this was an unjust outcome based on a flawed procedure). Copp introduces an additional alleged moral consideration in favor of this claim, namely, that "none of them believed that Borderline had achieved excellence in all three of the required areas." That is, each member of the committee considered individually had a reason to deny tenure above and beyond the failure of the candidate to meet the conditions for tenure under procedure P1, namely, that each believed that Borderline had not achieved excellence in all three areas; accordingly, each member considered individually did not have an all things considered obligation not to deny tenure. However, this fact should not be given any moral weight by the members of the tenure committee, assuming that all believe that P1 is a flawed procedure that has delivered an unjust outcome to Borderline. For it is the unacceptability of the procedure (P1) (by virtue of its inability to deliver a just outcome) that is in question; not the judgment of any individual member of the tenure committee in relation to Borderline's lack of excellence in all three areas.

Perhaps Copp intends here to be giving weight to the fact that *all* of the members of the tenure committee judged that Borderline was not excellent in all three areas, that is, each member considers not simply that he does not think that Borderline is not excellent in all three areas but that all the members of the committee think as he does. But that would be tantamount to retracting the proposition that P1 was in fact deeply flawed; it would be to suggest that P1—the procedure that delivered its verdict on tenure on the basis of whether each member of the committee believed Borderline was excellent in all three areas—was actually a reasonable procedure after all. Clearly this maneuver is not acceptable, since it

A final point to be made here is that even if it was allowable for each of the members of the tenure committee to give moral weight in their decision to their knowledge that it was their common belief that Borderline was not excellent in all three areas, this would not necessarily generate a different all things considered moral obligation for the committee members than that the university is under. For presumably the university should now likewise attach moral weight to the fact that none of the members of the tenure committee believed that Borderline was excellent in all three areas; and if the tenure committee is actually constitutive of the university with respect to the latter's decision making on tenure—there being no appeals process, for example—then the university would in fact be doing so (assuming the members of the committee decided to accord such weight). And if the university ought *not* to attach such weight, then this must be because it is a different collective—say, a committee of appeal or a Vice-Chancellor with the power to overrule the tenure committee's decision—from the collective constituted by the members of the tenure committee, that is, for this purpose the university is not identical with the tenure committee, and the two collectives have, to some extent, disparate obligations and responsibilities because the scope of their purposes is somewhat different. So we return yet again to the conclusion that Copp's example is not an example of the kind he needs, because it is not an example which involves a collective and the members of that very same collective for the purposes at hand.

Thanks to Andrew Alexandra, Dean Cocking, and Daniel Cohen for helpful comments on this paper.

Notes

[1] David Copp, "The Collective Moral Autonomy Thesis," unpublished paper dated December 29, 2005, prepared for the Fifth International Conference on Collective Intentionality, Helsinki, Finland, September 1, 2006, now published in *Journal of Social Philosophy* 38, no. 3 (2007): 369–88. See also David Copp, "On the Agency of Certain Collective Entities: An Argument from 'Normative Autonomy'," *Midwest Studies in Philosophy* 30 (2006): 194–221.

[2] Copp, "Collective Moral Autonomy Thesis," 369.

[3] Seumas Miller, *Social Action: A Teleological Account* (Cambridge: Cambridge University Press, 2001), chap. 5.

[4] Copp opts for this more general concept in the revised version of his paper. However, I want to maintain a distinction not only between moral responsibility and blameworthiness/praiseworthiness, but also between moral responsibility and moral fitness for a negative/positive response.

[5] Or, for that matter, the related but more general notion of fitness for (of deserving of) negative/positive response.

[6] Copp, "Collective Moral Autonomy Thesis," 371.

[7] Ibid., 370.

[8] I owe this point to Daniel Cohen.

[9] Copp, "Collective Moral Autonomy Thesis," 370–73.

[10] Ibid.

[12] Ibid., 375.

[13] Seumas Miller, "Institutions, Collective Goods and Moral Rights," *Proto Sociology* 18–19 (2003): 184–207; Seumas Miller, "Collective Moral Responsibility: An Individualist Account," *Midwest Studies in Philosophy* 30 (2006): 176–93; and Seumas Miller, "Joint Action: An Individualist Account," in *Intentional Acts and Institutional Facts: Essays on John Searle's Social Ontology*, ed. Savas L. Tsohatzidis (Dordrecht: Springer, forthcoming).

[14] I am assuming that the government is an *embodied* role, that is, if the PM's job is vacant then, properly speaking, there is no government (although there might be a caretaker government, for example). Accordingly, on this admittedly implausible view that the government simply is the PM, then it would in fact be critical to avoid the PM from being killed, since strictly speaking there would be no government unless and until a replacement PM was found. This is, of course, one of the problems of absolute monarchies and the like. Copp at one point appears to countenance the possibility of a government being an unembodied role (or perhaps set of roles). This really would be an abstractionist view of institutions! At any rate, it is to say the least implausible that abstract entities can have moral properties, not to speak of agency.

[15] Copp, "Collective Moral Autonomy Thesis," 375.

[16] Copp stipulates 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" ("Collective Moral Autonomy Thesis," 370). This has the appearance of defining pro tanto obligations by recourse to the notion of an all things considered obligation whereas, arguably, the notion of an obligation all things considered actually relies on the notion of a pro tanto obligation (or analogous notion, e.g., prima facie duty [objectively understood]). At any rate, in this paper of mine I assume that a *pro tanto* obligation (including forward-looking responsibility) is understood in the intuitively appealing sense of a *real* (as opposed to merely apparent) and *non-derived* (in the context in question) moral obligation (including forward-looking responsibility). A pro tanto obligation is non-derived in the sense that in the context in question it is not derived from any other obligations the agent has in that context. So a police officer has a pro tanto moral obligation to apprehend an offender in some crime context. This obligation is not in that context derived from other obligations he might have in that context. On the other hand, the standing obligation to arrest offenders (an instantiation of which is the obligation to arrest this offender in this crime context) may derive from some more relatively context-free set of moral obligations that has been abstracted from particular crime contexts, including the obligation to protect citizens and to discharge his oath of office. Given that a pro tanto obligation is non-derived (in the context in question), it might be overridden by other moral considerations in play in that context (hence the contrast with moral obligations [including forward-looking responsibilities] *all things considered*).

[17] Miller, "Joint Action."

[18] Miller, *Social Action*, chap. 5; and Miller, "Institutions, Collective Goods and Moral Rights."

[19] Ibid.

[20] Miller, *Social Action*, chap. 5; Miller, "Institutions, Collective Goods and Moral Rights"; Miller, "Joint Action"; Seumas Miller, "Social Institutions," *Stanford Encyclopaedia of Philosophy* (Spring 2007), available at http://plato.stanford.edu/archives/spr2007/entries/social-institutions; and Seumas Miller, "Collective Responsibility and Information and Communication Technology," in *Moral Philosophy and Information Technology*, ed. Jeroen van den Hoven and John Weckert (Cambridge: Cambridge University Press, forthcoming).

[21] Miller, "Institutions, Collective Goods and Moral Rights"; Miller, "Joint Action"; and Miller, "Social Institutions."

[22] Copp, "Collective Moral Autonomy Thesis," 378.

[23] Ibid., 379–82.

[24] Ibid., 380–81.