



The Weak Collective Agential Autonomy Thesis

David Botting
Instituto da Filosofia da Linguagem, Lisboa

Disputatio Vol. 4, No. 31

November 2011

DOI: 10.2478/disp-2011-0009

ISSN: 0873-626X

The weak collective agential autonomy thesis

David Botting

Instituto de Filosofia da Linguagem, Lisbon

Abstract

Can a collective be an agent in its own right? Can it be the bearer of moral and other properties that we have traditionally reserved for individual agents? The answer, as one might expect, is 'In some ways yes, in other ways no.' The way in which the answer is 'Yes' has been described recently by Copp; I intend to discuss his position and defend it against objections. This describes a fairly weak form of autonomy that I will claim does not require the abandonment of methodological individualism or our commonplace intuitions about individual responsibility. I will also discuss, and reject, a stronger conception of autonomy suggested especially by the work of Pettit that would result in methodological individualism.

Keywords

Collective responsibility; collective agency; methodological individualism; supervenience.

Introduction

Are there circumstances under which collectives should be considered as agents? There is a certain amount of common-sense to say that there are, purely from facts about our ordinary linguistic usage. It is common to attribute some action, attitude, or moral quality to a collective, e.g., 'The Simpsons went to church', 'England defeats Germany on penalties', 'Russell and Whitehead wrote the *Principia Mathematica*'. Let us suppose that each Simpson, with no thought of other family members, decide individually to go to church. The attribution 'The Simpsons went to church' can straightforwardly be

Disputatio, Vol. IV, No. 31, November 2011

Received on 14 September 2010

explained away in a distributive fashion by making the same attributions to each member of the collective, so in this example 'The Simpsons went to church' is just a paraphrase of 'Homer went to church', 'Marge went to church', 'Bart went to church' and 'Lisa went to church'. It does not seem that we would have to give the Simpsons, as a group, autonomy in this example. Not all collective attributions can be analyzed this straightforwardly, but the important question is whether they can still be analyzed distributively. The tradition in analytical philosophy has tended to answer 'yes' and endorsed methodological individualism even in complex cases where the members of the collective, unlike the Simpsons, co-operate, share goals and intentions. More recent philosophy has thrown this into doubt and argued that collectives have some form of autonomy that cannot be captured by an individualistic analysis. In this paper I will argue that there is collective autonomy, but only in a thin or weak form that does not require us to abandon distributive analysis or methodological individualism.

Although the data of our ordinary linguistic usage supports ascription to collectives, methodological individualism also seems to be endorsed by our common-sense: only individuals can think and act, and to suppose that collectives are appropriate bearers of such agential properties immediately conjures in the minds of some philosophers the metaphysical extravagances of Geists, group wills and the like, but these fears may be misplaced, and modern philosophers talk about 'plural subjects', 'pools of wills', and 'we-attitudes' whilst remaining modest in their ontological commitments. I will not be considering here the various ways in which philosophers have done this. What I am concerned with is the question whether all such attributions can be explained away as figures of speech such that it is only in a metaphorical sense that a collective can be considered as an agent: we must be able to explain away how we can attribute something to a collective without attributing it to all of its members, since in our second example, 'England beat Germany on penalties' is not a paraphrase of each English national beating each German national on penalties, but only some, namely the respective football teams, and; how we can attribute something to a collective without attributing it to any of its members, since to say 'Russell and Whitehead wrote the *Principia Mathematica*' does not seem equivalent to either a conjunction or a disjunction of the statements 'Russell wrote the *Principia Mathematica*' and 'Whitehead wrote the *Principia Mathematica*'. If the

term referring to the collective figures irreducibly in any analyses of these attributions, then we have a non-distributive analysis, and I will be considering this kind of analysis as a mark of a ‘thick’ conception of autonomy sufficient to result in the abandonment of methodological individualism and for collectives of a certain type to be considered on a par with individual human subjects.¹ Can examples be found of this thick conception? The claim that there can will be called *the strong collective agential autonomy thesis* and comes in two species.

One species concerns the ascription of *moral* properties, and can be further subdivided depending on the specific moral properties at issue, but I will be concerned with two basic kinds: a stronger and a weaker. The weaker thesis claims that we can ascribe moral properties – such as responsibility and obligation – to collectives, where we cannot ascribe the same properties to any of its members. In other words, collective *C* can be attributed moral property M_1 even though no member of *C* can be attributed moral property M_1 , but at minimum some moral property M_n must be attributable to some member of *C* and it is in virtue of this (at least in part) that *C* possesses M_1 . I will defend this claim, and argue that this is sufficient to achieve the result that collectives can be considered as having some moral autonomy, but not enough to conclude that they should be considered as agents in their own right. The stronger thesis claims that we can ascribe moral properties to collectives where we cannot ascribe *any moral properties at all* to any of its members such that the moral property of the collective is constituted by or supervenes upon the moral properties of the members. This result is held to be shown by discursive dilemmas, but I hold that these dilemmas show only that the

¹ Velleman (1997: 29-30), from whom I take this last example, notes that such attempts to explain away collective attributions using only individualist concepts fail to incorporate features of groups such as, when a group is asked to make a decision about something, it is being asked to make a decision as a group. This may be true, but does not necessarily entail a non-distributive analysis or a thick conception of autonomy. Note also that I am not saying that a non-distributive analysis is part of the definition of the thick conception of autonomy, only that it is a mark of this conception, such that a failure to instantiate one of these is also a failure to instantiate the other. I suppose that it is not inconsistent to hold to methodological individualism while still holding that collectives are genuine metaphysical entities, but I do not see what could motivate such a view since it seems a blatant violation of Occam’s razor.

relation between the properties of the collective and its members can be exceedingly complex, and reject this stronger thesis.

Another species of this genus concerns the ascription of *intentional* properties, i.e., the having of attitudes. Again, discursive dilemmas are used to show that there can be radical discontinuity of a particular kind between the beliefs of a collective and those of its members, and that we are faced with a choice between the deductive closure of the collective's attitudes and closure of its members' attitudes. The argument is made that if we opt for closure at the group level then we should consider that group as an agent in its own right. I reject this as based on a faulty criterion; the proper mark of thick autonomy is a non-distributive analysis, and not any kind of closure.

Any conclusion that we should treat collectives as autonomous agents is premature. All that we are entitled to conclude from discursive dilemmas is that a distributive analysis has to take the particular procedures of the collective into account; they are sufficient only for a *weak collective agential autonomy thesis*. Collectives are mereological sums of its members, and their agential properties supervene, in possibly complex ways, on agential properties of its members. By ascribing these properties to those members, we can produce a distributive analysis of the ascription of the agential property to the group.

1. The Collective Moral Autonomy Thesis

Copp (2007) argues that there are cases where we would like to say that a group is responsible for a particular act but we would not like to say that any of its members are responsible for that particular act. If so, then we are entitled to consider the group as having moral autonomy. On this definition, a group can have moral autonomy even if its moral properties depend on the moral properties of its members. Thus it is satisfied by the relatively weak claim that although none of the members are responsible for the particular act in question, they are responsible for some other act related to that act, or at least are the bearers of some relevant moral property or other. This is the *weak collective moral autonomy thesis*. Even if this is true, an ascription of responsibility for act A to collective G can still be reduced to the ascriptions of responsibility for acts related to A to the members of G, which is to say, it remains within the orbit of a distributive analysis.

Copp also argues for the stronger claim, that we might call the *strong collective moral autonomy thesis*, that a group can be responsible even when none of its members have any responsibility, even partially, connected to the act in question. Although Copp does not mention this, it seems that if this condition is satisfied, then it requires a non-distributive analysis.

Copp (2007: 370) introduces a distinction between *all things considered* obligations and responsibilities and *pro tanto* obligations and responsibilities. A *pro tanto* obligation is an intrinsically but defeasibly motivating reason for action with moral force. Thus an agent may have a *pro tanto* obligation to A but may not have an all things considered obligation to A, e.g., because of a conflict of duties out of which the duty to perform A is less pressing than some other incompatible duty. Similarly, the agent may be *pro tanto* responsible for A but not all things considered responsible for A. Whilst it is intuitively obvious what it means to have an all things considered or *pro tanto* obligation, what does it mean to be all things considered or *pro tanto* responsible? Copp (2007: 370) defines it as follows:

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 considerations² 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.

This definition is clearly utilizing a reactive attitude theory of responsibility. An all things considered responsibility derives its categorical force because it occurs after all relevant considerations have been taken into account.

We can now specify a strong collective moral autonomy thesis for responsibility by saying that a collective may be all things considered responsible for a particular act A even if none of the members is even *pro tanto* responsible for A. The weak collective moral autonomy thesis claims that a collective may be all things considered responsible for a particular act A when some of the members are only *pro tanto*

² It is not completely clear whether these countervailing considerations or defeaters must be specifically moral considerations, but I think that we should assume that they must.

responsible but not all things considered responsible for A. If the strong thesis is true, obviously the weaker thesis is also true.

Copp (2007) proceeds to give us some examples to argue for these theses.

i) The Kidnapping: The Prime Minister

The PM is held at gunpoint and told that she must release a particular prisoner or be killed. According to the rules of government, only the PM has this authority, so when she exercises it, she acts as a representative of the government. The PM has two *pro tanto* obligations, one not to give in to the demands and set a potentially dangerous criminal loose, and one to preserve her own life. In her deliberations the latter wins out, and she releases the prisoner. Although the PM has a *pro tanto* obligation not to release the prisoner, she does not have an all things considered obligation. However, the government does not have this excuse, but only the *pro tanto* obligation not to release the prisoner, which becomes an all things considered obligation in the absence of any other conflicting obligations or justifying reasons to do otherwise. So the collective is all things considered responsible for releasing the prisoner but the PM is not all things considered responsible, but only *pro tanto* responsible, for releasing the prisoner; we would not hold the PM to be morally criticizable in this scenario. This example seems to support the weak thesis.

Miller responds to this example by stating that an agent can have obligations *qua* group member and *qua* individual. The PM *qua* PM has an all things considered obligation not to release the prisoner, and since the PM acts in the name of the government, the government has the same all things considered obligation. *Qua* individual, it is morally excusable for her to release the prisoner (Miller 2007: 395).

It is not at all clear to me that it is intelligible to talk about all things considered obligations of S *qua* X, since the addition of '*qua* X' falsifies the maximal specificity condition on which the notion of 'all things considered' is based. If S considers something *qua* X that he does not consider *qua* Y, then any obligation can only be a *pro tanto* obligation and not an all things considered obligation. Perhaps we can partition considerations into X and Y and give a different meaning to all things considered such that it means only that S's obligation *qua* X cannot be defeated by another consideration within X, but I do not see the motivation or philosophical interest of doing this. What we

are ultimately interested in is the agent's moral situation overall, or as Copp (2007: 383) says: 'Once the role responsibilities . . . are taken into account, the question arises as to what the people themselves ought all-in to do.'

Miller (2007: 395-96) goes on to consider the responsibilities of the government:

Aside from the PM, no member of the government . . . could have done anything to prevent the prisoner being released. . . . they did not know, and could not reasonably have known, what was going on. Accordingly, the individual members . . . could not have prevented the prisoner being released and, therefore . . . they are not responsible. So, contra Copp, the government *per se* is not morally responsible.

Miller is here arguing from the non-responsibility of the members of the government to the non-responsibility of the government itself. But this clearly assumes what is at issue, namely whether the collective can be responsible while its members are not. The contributions of other members are not relevant here. It is by virtue of the fact that the PM's powers derive from the rules and norms of government that when those powers are exercised, the government is deemed to have acted.³ Miller seems to accept this kind of account, but says that although the government may be *institutionally responsible* it is not morally responsible. In contrast, he holds that the PM is morally responsible though not morally blameworthy. Copp, on the other hand, says that the PM is not morally responsible. If the government is not morally responsible either then no-one is morally responsible, and yet a reactive attitude does seem appropriate.

Ludwig (2007: 410) claims that this is a false dilemma – it is not true that 'when no one agent is fully responsible for the action of a group of which he is a member, the only entity that can be responsible is the group as such.' Instead, we should consider responsibility as distributed over the members of the group in proportion to their contribution to the performance of the act. Building on the responsible/blameworthy distinction, he carves out of those blameworthy cases some which although deserving of blame, we should not actively blame in the sense of intervening in the agent's habits. This is a rea-

³ We must assume what seems unlikely that these rules and norms are not contravened by the fact that the PM is held at gunpoint.

sonable point and is familiar from utilitarian literature, where we have to consider the utility of our moral practices of blame as well as the utility of the actions that we are blaming, and where it is tacitly assumed that coercion has a negative utility that must be overridden in the hedonic calculus, e.g., by deterring the repetition of the blameworthy actions. In Ludwig's opinion, the PM is morally blameworthy (and presumably responsible), although releasing the prisoner is the rational thing to do. I think that he would say that Miller, who we have seen says that the PM is not blameworthy, has misclassified the case, and that the relevant difference lies not in the moral status of the action but in the practical value of punishing it.

In summary, Copp says that the PM is not morally responsible. Miller says that the PM is morally responsible but not morally blameworthy. Ludwig says that the PM is morally responsible and morally blameworthy, but should not be blamed on rational grounds.

Although there are cases in which Ludwig's idea of distributing responsibility proportionately among the active participants applies, it seems to me that it only applies here if we presuppose that self-preservation is not a moral principle. Ludwig insists that self-preservation is not a moral principle in this particular case because it is inherent to the role of PM that obligations *qua* PM should be taken to override any personal obligations, and that in any clash where public and private moralities make conflicting demands, the private morality must be sacrificed, assuming that the occupant of the role of PM has entered voluntarily into that role (Ludwig 2007: 420-21). I disagree that such sacrifices are inherent; I think it depends entirely on the collective in question, on its rules and on what it accepts as justifying reasons for action, so Ludwig's insistence amounts to a solution by stipulation. Although non-moral considerations may come into play, we do not need to assume that they must; there is room to say that the PM morally should release the prisoner. So, I do not see why we should disregard self-preservation as a *pro tanto* obligation, and on a plausible assumption that it is morally permissible in some circumstances to act on principles of less moral weight and that to expect more is to expect unreasonably that we are dealing with a moral saint, we should not hold the PM acted immorally in saving himself, although he arguably should be praised if he refused to re-

lease the prisoner.⁴ This point seems to be acknowledged by Copp and Ludwig (Ludwig 2007: 419).

I think that the weak collective moral autonomy thesis is supported by this example. Generally, we can conclude that the individual members can have justifying reasons that, because these are also private reasons, the collective itself does not have. These reasons can defeat the *pro tanto* obligation of the individual but not of the collective. Hence, the individual has a defeated *pro tanto* obligation while the collective has an undefeated *pro tanto* obligation, which is to say, an all things considered obligation, and the collective is consequently all things considered responsible for the action.

ii) *The Prison Board*

Bob, Carol, Ted, and Alice are members of a prison board called to vote on resolutions aimed at improving security. Carol votes in favor of the resolution but the other members do not, so the resolution fails. Shortly afterwards, due to poor security, there is a prison riot resulting in bloodshed. The prison board is held to blame for the bloodshed since it would not have happened if the resolution had passed. Are the members of the board responsible? Carol is not, since she voted in favor. Bob, Ted, and Alice seem to be the guilty parties. But suppose that each had good reasons⁵ and that these are private reasons (Copp 2007: 377-79). What then?

Miller (2007: 402) argues that a participant in a decision-making process is institutionally responsible for any decisions made even if they did not personally vote in favor of the decision unless, perhaps, they take steps afterwards to dispute the decision or to withdraw from the process. If the decision is morally significant, then those who are institutionally responsible are also morally responsible. So for Miller the entire board is morally responsible, including Carol who voted for the resolution. However, Copp could happily concede this

⁴ This introduces a complication, since praise is obviously a reactive attitude. Should an agent be responsible for not-A if, although a reactive attitude does not seem appropriate for not-A, it is appropriate for A? According to Copp's definition, she is not responsible, but an argument could be made for changing the definition. I will ignore this here.

⁵ Objections based on whether the actual reasons Copp names are good enough to absolve responsibility seem to me to be beside the point.

since it shows only that the board members are *pro tanto* responsible. It does not touch the weak thesis.

Ludwig asks us to consider the case where only Ted is on the board but everything else is the same, including Ted's excuse. He (2007: 417) says:

Surely Ted is to blame in this case, if we say the Board was to blame in the previous case. There is *no one else to blame*. But it seems that Ted has just as good exculpatory reasons in this case as in the other. If we blame him here, then we can blame him in the other case. If we excuse him in the first case, we must excuse him in this case. But in neither case can we blame the Board and let Ted off.

But this is just what Copp denies and illustrates in the PM example. The point that Ludwig fails to appreciate is that Ted can have exculpatory reasons that, because they are agent-relative reasons, do not apply to the collective.

iii) The Tenure Committee

This example of Copp's makes use of discursive dilemmas. A three-person committee is deliberating about whether to grant tenure to a candidate, for which purpose it has a procedure where each committee member is to judge whether the candidate excels in three areas, and an additional vote about whether to grant tenure. Each member of the committee believes that the candidate excels in two out of the three areas but they disagree over which area he did not excel in, such that a majority believes of each area that the candidate did excel in it, which result is reported to the candidate. When it comes to the final vote, though, the candidate is rejected for tenure because he did not, in the opinion of any member of the committee, excel in all three areas (Copp 2007: 379-80).

Copp claims that this rejection is unfair and that the university is all things considered blameworthy for it. Since the procedure is faulty, individuals might be held blameworthy for those procedures, but to eliminate this possibility he stipulates that the designers of the procedure could not reasonably have been expected to anticipate such problems (Copp 2007: 380). The result is that no individual is blameworthy. This seems to support the strong collective moral autonomy thesis.

This is the line of thought that Ludwig takes: perhaps the committee members are not blameworthy but somebody in the university, namely the person responsible for the design of the procedure, is. The example has been complicated by the fact that we are dealing with two collectives here, the tenure committee and the university. Taking the university as the collective in question, this example does not show a case where no member of the collective is blameworthy. He rejects Copp's maneuver to evade this response: 'We may imagine that no individual anticipated the problem, but it is not in general an adequate response to a charge that someone is responsible for something in a way that warrants blame for a bad outcome that he or she did not anticipate a kind of problem that arose as the result of something he or she did' (Ludwig 2007: 414). Of course, it is not always an adequate response, but the question is whether it is adequate in this particular case. It seems to me that it is. Furthermore, I would not even say that it is 'in general' inadequate, unless it involved culpable ignorance, which is expressly ruled out by Copp in this case. Otherwise we would be committed to saying that agents are responsible 'in general' for outcomes that are both unintended and unforeseen.

Ludwig (2007: 415) tempers this rather harsh view with the following:

That is not to say that there may not be untoward consequences that no-one could have reasonably anticipated and so for which they should not be blamed. . . . However, if an individual were to be absolved on that score, a group agent would be due the same consideration, for a group cannot be expected reasonably to anticipate something its members cannot. So one cannot have it both ways . . . if that lets individuals off the hook, it must let the institution off the hook too.

But this is just the same question-begging assumption over again. However, he continues: 'The university may still have a responsibility to offer redress to the candidate even if it is not blamable.' Suppose, Ludwig seems to be saying, that we concede to Copp that no members of the university are blamable. This does not mean that the university as such must be blamable, but we are tempted to think this because we agree that the candidate was treated unfairly and deserves compensation. We have misidentified our reactive attitude because

our intuitions have been confused by the admitted fact that the person deserves compensation.

Miller argues in a similar way to Ludwig that the moral obligations of the university and its members must mirror each other. If the university has a *pro tanto* obligation to adopt a better procedure and grant tenure to the candidate, then so does the committee, although the committee may also have a *pro tanto* obligation to follow established procedure. The obligation that wins, on Miller's view, is the obligation to grant tenure, since the procedure that led to refusing it has been rejected as unsound. Hence, both the university and its members have the same all things considered obligation.

I agree with Miller that the committee members have a *pro tanto* obligation to grant tenure despite the initial plausibility otherwise and consequently I do not think that this example works as a demonstration of the strong thesis. But I do not see why we should assume that this is also their all things considered obligation. One's main obligation as a member of the committee is to follow the procedures laid out by the committee; if one contravenes those procedures, then one's action always runs the risk of not representing the committee, just as the PM's action does not represent the government if the rules that give the PM the authority she has are breached when she is held at gunpoint. The validity of the outcome, what makes the outcomes of the individual deliberations qualify as the decision of the committee, is following those procedures, even if from a wider point of view, another procedure would be fairer.

So, both the university and its members are *pro tanto* responsible for failing to grant tenure. The committee members are not all things considered responsible for failing to grant tenure because the obligation to follow procedure could override it and does in the scenario Copp describes. Other members of the university, e.g., those who designed the procedures, are also *pro tanto* but not all things considered responsible, because there cannot be a coherent all things considered obligation to consider every possible consequence. The question now is: what is the university's all things considered responsibility? We can agree that the candidate has been treated unfairly and deserves compensation. But is the university all things considered blamable?

I am not sure. I am inclined to think that, as in the kidnapping case, whilst it would have been praiseworthy for the university to revise its procedures, it is also supererogatory. Therefore, I do not

think that this example works, even if reinterpreted as an illustration of the weaker thesis. One of the reasons that it doesn't work is that, unlike the first example, all the reasons considered are group reasons – they all belong to public morality. Since the tenure committee is a part of the same institution as the university, it does not seem plausible to me that the tenure committee, or its members *qua* committee members, have excuses that the university does not. Therefore, I reject the strong thesis. If a collective has an obligation to A or is responsible

for A-ing, then at least one of its members must also have an obligation or be responsible by virtue of which the collective is responsible, although, since I am asserting the weak thesis, this obligation or responsibility is not necessarily for A itself. The moral properties of a collective supervene on the moral properties of its members and do not require a non-distributive analysis.

2. The Collective Intentional Autonomy Thesis

The same kind of supervenience operates also on intentional properties, and discursive dilemmas are usually phrased not in terms of responsibility for decisions but in terms of attitudes towards propositions. Just as the collective might be responsible for A whilst its members might not, a collective might accept the proposition *p* whilst its members might not. However, if methodological individualism applies here too, the collective's acceptance depends on some attitudes or other of its members. Discursive dilemmas show how attitudes towards some particular proposition at the group-level can be radically discontinuous from those at the base-level.

Pettit argues that this 'discontinuity' entails a kind of psychological autonomy on the part of the group such that the group is functionally equivalent to an agent. He uses the term *group agents* to denote groups that 'mimic the more or less rational way in which individual agents act' (List and Pettit 2006: 85). This includes having representational and goal-seeking states that direct the group's actions and that must be moving towards reflective equilibrium. List and Pettit (2006: 87) identify deductive closure as the critical condition that a group must satisfy in order to be a group agent, arguing that if a group is seen to be behaving irrationally, it is not likely to attract new members or be able to keep its old ones. In this way the group agent, like the indi-

vidual, sees itself and its agency as something persisting over time. I will be arguing that this is a sufficient condition for a weak form of intentional autonomy, but not the strong one that they seem to want. There are different ways in which beliefs of the group may supervene on the beliefs of its members. They call these ways *proposition-wise* and *set-wise* supervenience.

Let us reconsider the tenure committee from the point of view of the members' beliefs about the areas in which the candidate excels. Individual M_1 believes that candidate D excelled in research and teaching, but not in service. Since M_1 believes that D did not excel in all three areas, for these beliefs to be deductively closed and hence rational, M_1 must also believe, and does, that D should not be granted tenure. Individual M_2 believes that D excelled in research and service but not in teaching, and so agrees with M_1 that D should not be granted tenure. Similarly, individual M_3 believes that D excelled in teaching and service but not research, making it unanimous that D should not be granted tenure. These beliefs are all rational on the individual level.

However, the majority believed p that D excelled in teaching (individuals M_1 & M_3), a majority believed q that D excelled in research (individuals M_1 & M_2), and a majority believed r that D excelled in service (individuals M_2 & M_3). They also all believe *not-g* that D should not be granted tenure. The problem is that they also believe the biconditional (g if and only if (p and q and r)), i.e., that tenure should be granted if these propositions are true; furthermore, the majority do believe of these conjuncts that they are true, so if they believe the biconditional, then they should also believe D should be granted tenure in order for the group to be rational. Otherwise, the beliefs of the group will not be deductively closed. Proposition-wise supervenience, where the majority belief about some proposition (in the example, g) is taken directly as a function (in this example a majority, but other functions are possible but lead to the same results) of the members' beliefs about that same proposition, leads to refusal of tenure and irrationality at the group-level (List and Pettit 2006: 92-93).

The solution is *set-wise* supervenience. Instead of supervening on g , the group's belief about g is based on the members' sets of beliefs about p , q , r , and (g if and only if (p and q and r)). If they accept this last as a valid rule, then the procedure to ensure group rationality is to take some function of p , q , and r , e.g., a majority vote, and to

apply the rule (List and Pettit 2006: 95-99). Admittedly this can lead to some curious results, as Pettit (2004: 171) shows in a scenario where a group believes that g even when no members of the group at all believe that g . Such a scenario might be said to illustrate a *weak collective intentional autonomy thesis*. The board's belief that g supervenes on its members beliefs that the premises p , q , r , and (g if and only if (p and q and r)) are true, and the ascription of the belief that g to the tenure board can be distributively analyzed as the ascriptions of belief in the premises to the members. Pettit calls this a *premise-based* procedure, because only the members' judgments of the premises are taken into account, and the conclusion calculated by a rule of inference they accept. It should be noted that in using this procedure the attitudes of the group still depend on attitudes of its members; we have not satisfied a strong collective intentional autonomy thesis.

All groups will eventually reach such dilemmas where they are forced to choose between ensuring coherence with its own past by collectivizing reason, or risking social fragmentation by individualizing reason. Pettit says that 'you can have individual responsiveness or collective rationality but you cannot have both' (Pettit 2004: 174) and goes on to claim that groups that choose collective rationality 'deserve ontological recognition as intentional and personal subjects' (Pettit 2004: 175) because its attitudes are deductively closed.⁶ He says (2004: 182):

[T]he integrated collectivity, as characterised, is going to display all the functional marks of an intentional subject . . . there is no reason to discount those marks as mere appearances. Within relevant domains it will generally act in a manner that is rationalized by independently discernible representations and goals; and within relevant domains it will generally form and unform those representations in a manner that is rationalized by the evidence that we take to be at its disposal. In particular, it will manifest this sort of functional organisation, not just at a time, but over time; it will display the degree of constancy as well as the degree of coherence that we expect in any intentional subject.

List and Pettit want to argue that if we can provide a functional equivalent at the group level for whatever satisfies the criteria that

⁶ I would assume, although Pettit does not mention it, that the group-beliefs must be motivationally effective.

operate at the individual level for ascriptions of intentional properties, then the group is an autonomous agent in a stronger sense than that given by the weak collective intentional autonomy thesis.⁷ But

⁷ As we have seen, the criterion they concentrate on is deductive closure. List and Pettit argue that the group enforces a decision-making procedure such as the premise-based procedure to guarantee group-level rationality. This type of collectivized reason is a functional equivalent of individual reason, and satisfies deductive closure at the group level as the latter satisfies it at the individual level. It also provides a functional equivalent of the feature of agency as something persisting over time, since they argue that a group will not last long if it is seen to be inconsistent, and collectivizing reason will lead to a society that is strongly based on precedent and its previous decisions.

I am not convinced by this latter alleged equivalence, at least without further assumptions. In the first place, it is not clear to me that a sufficient distinction has been drawn between appearing inconsistent and being inconsistent. In the example, if the individual's judgments of the premises p and q are not common knowledge, then proposition-wise supervenience will appear rational although it is not and set-wise supervenience will not appear rational although it is. Secondly, I do not think that a group will last long if its decisions over some g consistently fail to match with the members' judgments over that g . We have seen that it is possible that the group's judgment does not match the judgment of any of its members. If this happens frequently, I think that such a group would not survive. Furthermore, I think that this is true even if it is proved, analogously to the way I have done above, that the decision reached is rational and based on procedures that all have agreed to. Emotional investments are invested in the conclusion, e.g., in whether a criminal ought to be punished, and not in factual, non-normative questions. If these emotions are continually frustrated, I do not think that appeals to rationality are likely to get very far. Helm (2008) gives an account of group agents that takes account of the emotional aspect.

Of course, there is any number of ways in which the survival of a group agent may be secured, and any number of ways in which a group agent might dissolve into a mere mass of people. List and Pettit emphasize the role of rationality and a commitment to rational unification, whereas I would emphasize the role of affect and a commitment to emotional unification. Conceivably, both might be considered as functional marks of an intentional subject. A society may even be able to survive under *modus tollens* and reject one of its premises; Orwell's "1984", with its departments for rewriting history and its double-think ensure that no-one, under fear of room 101, points out its contradictoriness and fallibility, describes such a society. It is not necessary that the rule, by following which reason is collectivized, is itself reasonable; a stupid rule would do almost as well. If a society can survive even without deductive closure, then the identification of the latter as a critical mark of intentional subjects is undercut. Perhaps List and Pettit would respond that it is enough that survival can be achieved this way, closure being a sufficient if not necessary condition and a functional equivalent in some cases even if it is not in all cases.

this is the wrong criterion anyway for determining whether something is an appropriate bearer of intentional or other agential properties. The proper criterion is whether the ascriptions of such properties require a non-distributive analysis, that is to say, whether the realizers of such functional states require anything outside of the supervenience base of their members and the agential properties of their members. To put it another way, conscious beings are the basic units of methodological individualism because no parts of such beings are themselves conscious or have intentional properties. This is itself a functional mark of intentional subjects, and is not satisfied by group agents, whose parts (the group members) do have intentional properties. I do not see anything in the examples we have considered to suppose that this criterion has been, or even can be, satisfied.

Conclusion

I summarize the argument in this paper as follows:

1. If an object's possessing a property P supervenes entirely on properties of its mereological parts then the ascription of M to the object can be distributively analyzed as the ascriptions of those supervened on properties to those mereological parts.⁸

⁸ Perhaps an objector might dispute the close relation that I have made between supervenience, distributive analysis, and autonomy. They might object that I have assumed that only the most fundamental objects in our ontology are genuine entities and genuine bearers of properties, or perhaps that demonstrating a supervenience base of an object's properties is enough to exclude that object from our ontology. They might reply that there is a difference between saying that a property being ascribed to an object supervenes on others and in saying that it is in a metaphorical, rather than genuine, sense that the object bears the property – linguistic reduction is not necessarily ontological reduction. They might agree that a collective's intention to A supervenes on its members' attitudes, and yet still maintain that the collective exemplifies this intentional property. For them, the collective could still be autonomous.

It is true that I make this assumption. For me, collectives are mereological sums. Whether or not the object is identical to its mereological sum (and I am not even sure whether this question makes sense) should not affect my thesis unduly because the argument given above does not actually rely on an assumption this strong. We can say without inconsistency that some properties, like agential properties, are emergent at the level of conscious beings, and that their supervenience base does not include agential properties, but only physical properties. However, any putative

2. If the ascription of agential property M to a collective can be distributively analyzed (as described in (1)) in terms of the agential properties of its members, then we have methodological individualism.
3. If we have methodological individualism, then the strong thesis has not been demonstrated.
4. There are no cases of agential properties possessed by collectives that do not satisfy the antecedent in (2). In particular, even the most problematic cases produced so far have been shown to satisfy the antecedent in (2), which is to say that they demonstrate only the weak thesis.
5. Therefore, the strong thesis has not been demonstrated.

I would go further than saying that the strong thesis of collective moral autonomy has not been demonstrated and reject it on the theoretical grounds already alluded to. If a collective has a responsibility or obligation, then some of its members must also have at least a *pro tanto* responsibility or obligation. However, private (and, I suggest, only private) considerations can override an individual's *pro tanto* obligations without overriding the collective's *pro tanto* obligations, or to put it another way, some of the individual's *pro tanto* obligations are also the collective's *pro tanto* obligations, and some are not, so the collective's all things considered obligations do not consider these latter. It follows from this that moral properties ascribed to collectives derive in some manner – however complex – from moral properties (not necessarily the same properties) ascribed to its members. We can still analyze such ascriptions distributively, giving a conception of autonomy 'thinner' than required for the claim that collectives can be considered on a par with individual human subjects.

The same kind of reasoning applies for intentional properties as for moral properties, so I also reject the strong thesis of collective intentional autonomy. The propositional attitude towards some proposi-

exemplifications of this property at a higher level than, and that supervenes entirely on, the level at which the property emerged, are not genuinely exemplified and statements that they are exemplified must be read metaphorically. My arguments refer to agential properties, and do not exclude the possibility that the ascription of other kinds of property require a non-distributive analysis. Maybe some properties are genuinely emergent at the collective level, but I have yet to see a convincing example of one.

tion ascribed to the collective is always derived from the propositional attitudes of its members to some set of premises from which the proposition in question can be inferred through some agreed-upon rule of inference, as Pettit himself demonstrates. Taking this *premise-based* procedure as some kind of functional equivalent to an individual's conception of his agency enduring through time, Pettit argues that this amounts to the collective's being an individual agent. I reject this as a faulty criterion. As long as we have a distributive analysis – which is what a *premise-based* procedure basically is – we have no reason to grant autonomy to collectives in the sense that Pettit (2004) would have us do.

That Copp is not arguing for autonomy in this strong sense is shown when he says: 'It might seem to follow that the properties of a collective must be accounted for by the 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*. . . . collectives are not 'independent agents'' (Copp 2007: 369 [my italics]). Thus, although he argues for the strong moral autonomy thesis, this condition is satisfied even by the weak thesis; significantly, one that still allows for a commitment to distributive analyses and, by extension, methodological individualism.

David Botting
 Faculdade de Ciências Sociais e Humanas
 Universidade Nova de Lisboa
 1069-061
 Lisboa, Portugal
 davidbotting33@yahoo.co.uk

References

- Copp, David. 2007. The collective moral autonomy thesis. *Journal of Social Philosophy* 38 (3):369-388
- Helm, Bennett W. 2008. Plural Agents. *Noûs* 42 (1): 17-49
- List, Christian and Pettit, Philip. 2006. Group Agency and Supervenience. *Southern Journal of Philosophy* 44 (S1):85-105.
- Ludwig, Kirk. 2007. The Argument from Normative Autonomy for Collective Agents. *Journal of Social Philosophy* 38 (3):410-427
- Miller, Seumas. 2007. Against the Collective Moral Autonomy Thesis. *Journal of Social Philosophy* 38 (3):389-409

- Pettit, Philip. 2004. Groups with minds of their own. In *Socializing Metaphysics* ed., Frederick Schmitt. New York: Rowman and Littlefield, pp. 167-93.
- Velleman, J. David. 1997. How to Share an Intention. *Philosophy and Phenomenological Research* 57 (1):29-50.