Erkenn (2014) 79:1663–1684
DOI 10.1007/s10670-014-9634-9

ORIGINAL ARTICLE

Group Agency and Individualism

Carol Rovane

Received: 11 April 2014 / Accepted: 11 April 2014 / Published online: 30 April 2014

Abstract Pettit and List argue for realism about group agency, while at the same time try to retain a form of metaphysical and normative individualism on which human beings qualify as *natural persons*. This is an unstable and untenable combination of views. A corrective is offered here, on which realism about group agency leads us to the following related conclusions: in cases of group agency, the sort of rational unity that defines individual rational unity is realized at the level of a whole group; rational unity is never a metaphysical given but always a product of effort and will; just as it can be realized within groups of human beings it can also be realized within parts of human beings, as well as within whole human lives; in cases of group agency, the rational unity that is achieved at the level of the group typically precludes rational unity at the level of its human constituents within their whole lives, though it can be realized within parts of those human lives. Along the way, a contrast is drawn between cases of genuine group agency and the cases of political agency envisaged by Rousseau and Rawls (and by Pettit and List) which leave individual human beings intact as separate agents in their own rights.

When philosophers are open minded about the possibility of group agency, they accept that there can be a significant similarity, or parity, between how a group functions and how a normal adult human being functions. More specifically, they accept that a group of human beings can realize, or at least approximate, at the level of the whole group the same kind of rationality that is characteristically realized by normal adult human beings. When this happens, the group itself functions as an individual agent in its own right. The question then arises, what becomes of the human constituents of the group agent? Do they remain individual agents in their

C. Rovane (✉)
Philosophy Department, Columbia University, New York, NY 10027, USA
e-mail: cr260@columbia.edu

own rights, even as they constitute a group agent? Although most philosophers would be inclined to answer affirmatively, I shall argue that this is not the case.

Let me clarify first of all, that if the theorist of group agency wishes to retain a commitment to some form of individualism, it cannot be the doctrine of *methodological* individualism that is sometimes espoused in the social sciences. The defining commitment of that doctrine is to reduction—the reduction of social facts to facts about the individual psychologies of human beings. Whereas, those who argue for group agency typically argue for an irreducibility claim, to the effect that the facts about group agents are *not* reducible to facts about the individual psychologies of their human members.

But what other form of individualism might be affirmed by the theorist of group agency who, unlike the methodological individualist, is committed to the irreducibility of the group phenomenon? It is best described as a form of *metaphysical* individualism, which I will preliminarily gloss as follows: whenever individual human beings bring about or participate in group agency, they nevertheless remain individual agents in their own rights.

In their recent work on group agency,[1] Philip Pettit and Christian List have advocated precisely this combination of views. On the one hand, they affirm the reality of group agents in a sense that makes the facts about them irreducible to facts about the individual psychologies of individual human beings, while on the other hand, they affirm that individual human beings are *natural persons* who necessarily remain the individual agents they naturally are, no matter what their involvements in group agency may be. They supplement their metaphysical individualism with a corollary *normative* individualism, which recommends that group agents be expressly organized so that they do not dominate their human members, and moreover, so that they safeguard the individual rights and interests of their human members.

Pettit and List also offer a somewhat perplexing, and I think ultimately unstable, account of how group agents might aim to achieve, at the level of the whole group, the sort of rationality by virtue of which they might count as rational agents in their own rights. They begin their account by reviewing various methods of group decision making that do *not* yield rationality at the level of the whole group, because these methods *aggregate* the views of individual members and mere aggregation is not itself a rational process. This was first demonstrated by Arrow with respect to voting on pairs of preferences,[2] and List has contributed additional formal results along similar lines.[3] But rather than infer that it is not possible to achieve rationality within a group, Pettit and List propose various constraints on aggregative methods of group decision making, which are designed to help group agents avoid the failures of rationality that such methods can be demonstrated to yield.

If these proposed constraints sufficed to ensure that a group agent could employ *the very same sorts of rational procedures* that normal adult human beings characteristically employ when they aim to be rational, that would completely

[1] Pettit and List (2011).

[2] Arrow (1951).

[3] Many of these are offered in Pettit and List (2011).

vindicate the sort of realism about group agency that Pettit and List wish to defend. I myself have argued elsewhere that groups of human beings *can* employ such rational procedures, and so they *can* qualify as individual agents in their own rights. However, my argument does *not* portray a group's will as being in any sense a *function* of the individual views of its members; and more importantly for my purposes in this paper, it shows that realism about group agency is in serious tension with the residual commitment to metaphysical individualism that Pettit and List wish to retain.

Here is a very brief statement of the reason why: When human beings achieve rationality together at the level of a whole group, parts of their intentional lives must (together) constitute the thoughts and actions of the group agent, and these thoughts and actions must necessarily proceed from the group's point of view; as a result, insofar as the human constituents of a group agent also retain 'separate' points of view, these points of view are bound to be somewhat smaller than what we would ordinarily think of as a whole human being's point of view, for they do not include the thoughts and actions that constitute the group agent's point of view. To put my point in another way: when human beings achieve rational unity together at the level of a whole group, this tends to produce a certain kind of *rational fragmentation* in their lives, because not all of the thoughts and actions associated with their brains and bodies proceed from the same point of view—some proceed from the group agent's point of view, while others proceed from a point of view that is somewhat smaller than human size.[4]

When Pettit and List claim that human beings are natural persons, part of what they are claiming is that such rational fragmentation within the human being does not follow upon involvements in group agency. It is unclear to me whether they might be making a stronger claim as well. Perhaps they are claiming that rational fragmentation *cannot* follow upon such involvements, or perhaps they are claiming that even if it can follow, it nevertheless *ought not to be permitted* to happen. In any case, when I argue that such rational fragmentation does follow upon involvements in group agency, part of what I am claiming is that Pettit and List are mistaken when they claim that human beings are natural persons. The more general lesson of my arguments is that there are no natural persons at all, as I shall explain in the final section of the paper.

Traditionally, the philosophers who have taken the most interest in group agency have approached it from the perspective of political philosophy. This seems to make good sense, since it would seem that if there is such a thing as political agency, it must consist in, or at least involve, a form of group agency. It would also seem that this particular form of group agency ought to be compatible with a commitment to

[4] These claims about the different sizes that rational points of view can have are not meant metaphorically. On the account I am offering, a rational point of view consists of the thoughts that together comprise the proper basis of deliberations that proceed from that point of view. Since the thoughts that comprise a group agent's rational point of view are scattered across different human lives, it literally takes up a larger region of space–time than a rational point of view that consists solely of thoughts within a given human life. I'll say more as I go on about the different sizes that rational points of view can have, and why human beings who constitute group agents typically also house rational points of view that are—again, literally—smaller than human size, because some of the thoughts that occur within those human beings' brains belong to the group agent's point of view.

metaphysical individualism, since the primary goal of political agency is often conceived to be the protection of the rights and interests of individual human beings. This is especially clear in the political arrangements envisaged by Rousseau and Rawls, in which a common will is formed at the level of a whole political group, which is rationally directed at safeguarding the rights and interests of individual human beings. Yet we shall see that the specific methods by which Rousseau and Rawls aim to ensure that the common will of a political group is rationally directed at safeguarding the rights and interests of individual human beings have special features in the light of which they are a very poor model for group decision making in general, and for this very reason the forms of political agency envisaged by Rousseau and Rawls barely qualify as cases of group agency at all, and are at best a limiting case of it.[5] Nevertheless, it will be helpful to approach the general topic of group agency by starting with their picture of political agency. Not only is this in keeping with a long tradition in which philosophical interest in group agency has been tied to preoccupations in political philosophy (and in fact, Pettit's own interest in group agency has grown out of his work in political philosophy); but also, their picture provides an instructive contrast with *genuine* group agency, in which the formation of a group will *does* affect the status of its human constituents as individual agents in their own rights.

In Sect. 1 I will contrast the cases of political agency envisaged by Rousseau and Rawls and genuine group agency, and then I will go on to draw a further contrast between these cases and the many cases of group agency that Pettit and List discuss—in which 'aggregative' methods of group decision making are 'improved' with the aim of achieving rationality at the level of the whole group.

In Sect. 2, I will argue that groups of human beings can achieve the very same kind of rationality that is characteristic of an individual agent, and furthermore, when they do they qualify as individual agents in their own rights. I will also explain why the achievement of rationality at the level of the whole group tends to produce rational fragmentation within its human constituents.

I should mention that I will not be offering my account of group agency as a realistic model for most political decision making, or for decision making within any very large group of human beings. But the account is still illuminating, because it provides a criterion or standard by which to determine whether a group agent counts as *real*. Thus, the greater the distance between how groups actually operate and the account of group agency that I will be offering here, the less clear grounds we have for supposing that these groups really do qualify as agents in their own rights.

[5] It was suggested by a referee for this paper that when Rawls envisaged the original position from which principles of justice are to be selected, he wasn't really trying to characterize a form of group agency. I won't try to argue that he or she is mistaken about this. I'll merely point out that the principles of justice are interesting only insofar as they can direct political action, and moreover—with the exception of one-person dictatorships—political action is characteristically a *group* activity. So it doesn't really matter for my purposes whether Rawls himself intended that the choices made from his original position be thought of as constituting the choices of a political group; what matters is that political theorists who are interested in political agency might naturally think of it in this way.

In Sect. 3 I will clarify what I take to be the general lesson of my account of group agency, which is that the existence of *any* agent is always a product of effort and will, and is never a metaphysical given—which is just to say, *there are no natural persons*. The reason why is that the boundaries between agents can always be re-drawn so as to configure agents of different sizes—not only agents of human size, but also, group agents composed of many human beings and multiple agents within a human being, and hybrid cases of the sort I described above, in which the existence of a group agent goes together with rational fragmentation within the lives of its human constituents.

1 Political Agency as Group Agency

In *The Republic*, Plato made a direct analogy between individuals and the state, suggesting that justice amounts to the same thing in both—namely, a proper harmony of parts directed by reason. But when Plato portrayed the just state as functioning in exactly the same way that a just individual would, he didn't display much concern about the points of view of individual citizens. He reasoned that since most citizens lack the insight of philosopher kings, it would be best if they were directed by others, without regard for their own beliefs about how the state should be ordered, and indeed without regard for their own beliefs about how their respective individual lives should be ordered.

Unlike Plato, Rousseau aimed to incorporate some genuinely democratic commitments into his account of political agency.[6] Thus, while on the one hand he believed that a political group can achieve the sort of rational unity that is characteristic of the individual agent, by forging what he called a *general will*, on the other hand, he also portrayed the general will as constituted by the individual wills of each and every citizen. More specifically, he held that each citizen must be autonomously identified with the general will, with the result that when citizens are directed by the general will they are *self*-directing as well. Rawls's account of political reasoning from the 'original position' seems to fulfill the very same desideratum, by ensuring that the most fundamental principles of justice by which the political community is to live are separately chosen by each individual when she reasons from the perspective of the original position.[7]

There are two crucial and striking features of the political reasoning envisaged by Rousseau and Rawls: first, the common will of the political group is to be achieved through identical reasoning that can be carried out separately by each individual citizen; second, the goal for the sake of which political agency is exercised is to safeguard the rights and interests of the individual citizens who comprise the political group. The first feature is really the mechanism through which the second feature is to be secured. Thus, if we ask, how is it that the will of a group agent manages to stay focused on the goal of securing the individual rights and interests of its human members, the answer given by Rousseau and Rawls is that the group's

[6] Rousseau (1997).

[7] Rawls (1971).

will is arrived at in a very special way, through individual reasoning that can be carried out separately by each member of the group.

This is somewhat more explicit in Rawls than in Rousseau, since Rawls explicitly framed the original position as one in which individual human beings would all take up the very same questions, about what fundamental principles of justice should govern their political community, and then, they would all reason in the very same way towards the very same conclusion. Yet although Rousseau was somewhat less explicit about this, he did say that when citizens are well informed and have good judgment, it is best that they arrive at an 'original' social contract by reasoning on their own, because when they reason together they may divide into factions and this will deflect their attention away from the proper goal of their reasoning, which concerns the truly *common* interests of *the whole*. It is a matter of some debate whether these common interests were conceived by Rousseau as identical *individual* interests, as opposed to *communal* interests. Again, the matter is much more explicit in Rawls, who instructed individuals in the original position—who are behind a veil of ignorance that shields them from knowledge of their particular identities—to reason in a self-interested way, so as to secure the best outcome for themselves no matter who they turn out to be when the veil is lifted. But I take Rousseau's many references to individual rights and freedoms to show that he did conceive the general will as serving individuals. That, at any rate, is how I shall read him for the purposes of this paper. Thus, for my purposes here, the significance of Rousseau and Rawls lies in the prospect they hold out, of forging rationality at the level of the whole political group, while at the same time safeguarding the individual rights and interests of its human members, the citizens.

So let me return to the point I made above, about how the group's will is to be harnessed to this latter goal according to Rawls—and I think according to Rousseau as well. As I've already indicated, this is to be accomplished through a particular way of forming the group will, as emerging through independent reasoning on the part of individual citizens: when each citizen reasons about what fundamental principles should govern the political group to which she belongs, she can recognize that, and also why, every other citizen will reason in the way same that she does; and so she can regard herself as an exemplar for the whole group, with the result that the group's will takes on the character of *her* individual will, and with the further result that this group will is perfectly aligned with every other human point of view within it.

A similar exemplariness of the individual reasoner can be found in Kant's account of moral reasoning, but with the following differences: when I reason my way to the categorical imperative, I regard myself as an exemplar of *all rational agents*, as opposed to a group of rational agents who might together constitute a particular political group; and the moral principle that Kantian moral reasoning leads me to embrace is a principle that is to govern my exercise of my *own individual agency*, and this marks a contrast with the sorts of principles that Rousseau and Rawls had in mind, which would govern a *group's* exercise of its political agency.

Let us suppose for the sake of argument that Rousseau and Rawls managed to give a convincing account of how the wills of individual citizens might thus

perfectly align, so as to form a group will about the most fundamental principles by which they are to live together in political community. It seems pretty obvious that once citizens go beyond such basic matters, in order to take up the many more specific and concrete political issues that are bound to come before them, they will require a different procedure for political decision making—that is, they can no longer expect that if they reason *separately* they will arrive at a consensus. For this reason, most theorists in the social contract tradition hold that once the fundamental political commitments of the group are determined through such a consensus (this is what Rousseau referred to as the original contract), other matters can be handled by voting with majority rule—the thought being that a commitment to majority rule about these other matters is itself something about which the political group might achieve consensus.

In their account of group agency, Pettit and List are very clear that majority rule is an inadequate basis for group decision making, in part because it does not serve their commitment to normative individualism (it does not protect individual human beings from being dominated by the groups to which they belong), but also because it does not secure the form or level of rationality that is required for their realism about group agency. The real problem, however, isn't just majority rule. It is a wider problem that concerns the nature of voting, and all other methods of group decision making in which a group's decisions are a function of the individual views of the group's human members—the methods that Pettit and List call *aggregative*.

Arrow drew attention to this problem in his ground-breaking work on voting, in which he demonstrated that voting can lead to violations of the requirement of transitive preference ranking at the level of a whole group, even if each member of the group individually satisfies that requirement.[8] Pettit and List see this as just one instance of a more general problem. They find another instance of it in the Discursive Dilemma that arises in legal theory, which arises in the following way: a majority of a panel of judges vote for A, and a majority of the judges vote for B; and yet, although A and B entail C, a majority of the judges do not vote for C. As in Arrow's case, we get a violation of rationality at the level of the group even though its members are all individually rational. List has contributed several more formal results, to show that various other methods by which we might hope to generate group decisions by aggregating individual attitudes can yield violations of some basic requirements of rationality.

Here is the root of the difficulty, as I intuitively see it: very often, individuals who vote the same way on a particular occasion do so for different reasons, and these different reasons will naturally lead them to vote differently from one another on other occasions; thus, because there is no common set of rational considerations to which the members of the group are all rationally responding in the same way when they vote (*unlike* what Rousseau and Rawls envisaged in connection with the choice of fundamental political principles), the overall patterns of their voting on a range of issues are not likely to exhibit the sort of rationality that is characteristic of an individual agent's deliberations.

[8] Arrow (1951).

Although Pettit and List emphasize that voting and similar aggregative methods of group decision making may lead to violations of basic requirements of rationality, I noted in my introductory remarks that they are nevertheless optimistic about the prospects that group agents might be able to achieve rationality at the level of the whole group. Their optimism seems in part to be empirically driven—they are struck by the extent to which many groups seem to function in rational ways, in the sense of controlling their behavior for the sake of achieving specific goals, and even entering into dialogue with other agents about the nature of their goals and the reasons for them. Their optimism also seems to be driven in part by sympathy for functionalism, a view of the mind that is particularly hospitable to group agency owing to the fact that it is completely liberal about the material conditions in which mental processes might be realized when they are construed as functional processes. Since functionalists are committed to being open minded about the potential mindedness of robots, for example, why not groups as well? But Pettit and List also invoke normative considerations in favor of their realism about group agents. In spite of all they have to say about why aggregating individual attitudes in order to arrive at group decisions can fall short of delivering rationality at the level of the group, they maintain that groups can take many measures to avoid the pitfalls of these aggregative methods—for example, by using straw polls in order to track how patterns of voting are emerging within the group, or by refining the issues to be voted on in order to make their logical relations explicit, or by thinking carefully about the order in which issues are to be voted on, so that the group has a chance to allow their votes on premises to determine their conclusions by logical implication (thereby avoiding the Discursive Dilemma).[9]

From a philosophical point of view, there is something peculiar about Pettit's and List's strategy of trying to show that group agents can achieve rationality at the level of the whole, by first arguing that various aggregative methods of group decision making are bound to fail on this score, and then seeking to refine those very methods so that they don't fail. We would do far better to start with the following questions: What, *in general*, or *in the abstract,* does rationality require of individual agents? How do they normally achieve rationality when they do—and if they can't fully achieve it, how do they still strive to achieve it? Is there any obstacle to supposing that a group of human beings might follow these very same rational methods? These questions set the topic of my next section.

2 Agency and Rational Unity: In General and in Groups

As I explore the conditions of rational agency, I shall confine my attention to cases of *reflective* rational agency—by which I mean, agency that is exercised *from* a particular point of view, which is a site of deliberation and choice, and a site from which actions proceed. I will assume that it makes no sense to speak of particular

[9] Maskin (2001) has argued that when rankings are restricted in various ways, including *ideological* rankings, majority rule is a good (where this connotes *rationally acceptable*) procedure for group decision making.

thoughts and actions without attributing them to a particular agent, so that the thoughts must always be portrayed as figuring in that agent's point of view, and the actions must always be portrayed as proceeding from that agent's point of view. This assumption flags something of importance for the social dimensions of reflective rational agency: insofar as an agent has its own point of view from which to deliberate, choose and act, it can enter into distinctively interpersonal relations in which others *engage* its point of view, such as conversation, argument, criticism and promising; furthermore, such an agent can be held responsible for what it does in the fullest sense, which involves engaging its capacity for *self*-criticism.

Occasions for deliberation arise when an agent takes itself to face multiple options from which to choose and then seeks grounds for its choices. This involves evaluating its options with a view to determining which among them it would be better to choose and which among them it would be worse to choose. This comparative evaluation of options presupposes a certain *ideal* of rationality, which is to determine what it would be *best* to do—not best in a sense that would imply omniscience, but rather, *best from the agent's own point of view*. An agent who is committed to this ideal is committed to arriving at and acting upon *all-things-considered-judgments*, which are judgments about what it would be best for the agent to do in the light of *all that it thinks*. Thus, to deliberate *from* a point of view is to take into account the various thoughts that *constitute* that point of view, which are the agent's *own* thoughts—the thoughts the agent would regard in a first person way, as *mine*.

The point of introducing this ideal of rationality is not to suggest that agents typically live up to it in their deliberations, or for that matter, that they typically deliberate before acting. Agents often act without deliberating, and even when they take time to deliberate they usually fall short of the ideal of arriving at an all-things-considered judgment. All the same, *if* an agent is committed to being rational, then it is committed to realizing the ideal—for it believes that what it *ought* to do is what would emerge as *best* if it took due account of all that it thinks. I will not try to offer a complete list of all of the specific rational activities that an agent would have to carry out in order to take due account of all that it thinks. But here are three that any such list would have to include: the first is to resolve conflicts among one's attitudes, by getting rid of contradictions among one's beliefs and by sorting out priorities among one's evaluative commitments (by weighing, ranking and aggregating them); the second is to identify the options before one, given what one believes about one's general abilities and one's specific circumstances; the third is to work out implications of one's various thoughts *taken together*, in order to arrive at an answer to the deliberative question, which among one's options would it be best to choose—and this, of course, is what an all-things-considered judgment articulates.

When agents are committed to this rational ideal, it is appropriate to hold them responsible, because it is meaningful to ask them for an account of why they did what they did. I want to emphasize that this can be meaningful even in cases when an agent has fallen far short of the ideal, either because it didn't deliberate at all, or because it acted against its own, deliberated, all-things-considered judgment (as in cases of weakness of will). In such cases, the agent's own commitment to being

rational provides a normative standard in the light of which it can assess whether the actual reasons for which it acted are in accord with its own best reasons. This ensures that it can be, and moreover, that it should be, *self*-critical for its rational failings; it also ensures that it can and should be receptive to corresponding criticisms from others—thereby *taking* responsibility for what it does. (Obviously, an agent may disagree with others' criticisms, and this will generally be an occasion for moral and political argument. But this does not speak against my point connecting responsibility and a capacity for self-criticism, since agents who are willing to engage in moral and political argument are arguing, in part, about what grounds there are for criticism of both oneself and others, and in these arguments they are all taking responsibility for themselves as well as holding one another responsible).

I'm now ready to explore how a group of human beings might come to be committed to this ideal of rationality, and to have a point of view from which to deliberate, choose and act, and to engage with others in distinctively interpersonal ways, and to take responsibility for itself. I want to explore this by considering the case of a philosophy department that is faced with the task of designing and implementing a Ph.D. program in philosophy.

Let us begin by supposing that the faculty members who comprise the department all have personal views about what the requirements for the Ph.D. should be. And let us suppose that the department's options are determined by compiling all of the various possible requirements for the Ph.D. that might be proposed by one faculty member or another. Given Arrow's famous result on voting, along with List's companion results, we can predict that if the department were to vote separately on each possible requirement, the results might well amount to an incoherent set of Ph.D. requirements. This will happen if faculty members who vote for the same requirements do so for different reasons, with the result that their votes on different requirements fail to line up in any rational way. It might seem that the department could avoid this problem and still employ a voting procedure, by placing appropriate constraints on it in the spirit that Pettit and List recommend. They might, for example, restrict the options to be voted to whole slates of requirements, each of which is internally coherent. But even this procedure would leave the department without a unified sense of *why* the slate that the majority voted for is *best*, all things considered, since even with respect to these options, faculty who vote the same way might do so for different reasons. Consequently, if the Director of Graduate Studies were faced with having to explain to students *why* the department imposes the Ph.D. requirements that it does, she would be left having to report that the majority of the faculty had voted for them, albeit for various different reasons.

It seems obvious to me that the only way in which the department overall could have reportable *reasons* for its Ph.D. requirements would be if it had arrived at them through a department wide deliberation, whose aim it was to work out which set of requirements really is best by the lights of the department's point of view. There might still be an impression that this could in principle be accomplished through voting, so long as each slate of requirements to be voted on was accompanied by reportable reasons—a statement of its overall rationale. Perhaps so. But if the stated

reasons are to be binding, then they should also bear on future decisions of the department when it has occasion to revisit and re-evaluate its Ph.D. program. This would require that the department employs a method of decision that is systematically responsive to the normative force of those reasons across a range of conditions, and such a method would, by its very nature, be a *deliberative* method—that is, it would no longer in any sense be an aggregative method.

So let us inquire more directly into *how* a whole department might deliberate rather than vote. It might seem that a department could not do this unless its members already agreed on a very great deal—in particular, on the *basis* from which their deliberations should proceed. In the case at hand, this would mean that the members of the department would have to agree, antecedently to their deliberations, about all of the various matters of fact and value that are relevant to designing and implementing a Ph.D. program. Now, if the members of a department did agree to this extent, then it would seem that they could employ the method by which Rousseau and Rawls proposed to achieve a group will, which doesn't require a *group* deliberation at all, but rather, separate deliberations by the group's members—deliberations that are presumed to be identical because they all proceed from exactly the same basis. But I've already indicated that it isn't realistic to suppose that the members of a group do, in general, agree to this extent.

I can report that when the Columbia Philosophy Department set out to re-examine its requirements for the Ph.D., the faculty brought to the table strikingly divergent views about what those requirements should be. Some thought there should be no requirements at all besides the dissertation; some thought there should be a two-language foreign language requirement; some thought there should be a one-language foreign language requirement; some thought there should be a 'serious' logic requirement; some thought there should be a history requirement to study philosophy prior to the twentieth century; some thought there should be a history requirement that could be fulfilled by studying early analytic philosophy and twentieth century phenomenology; some thought there should be comprehensive examinations; some thought there should be topical examinations in the field of a student's dissertation; some were against all examinations besides the defense of the dissertation. Faced with this diversity of viewpoints, and wishing to have a rationale for our Ph.D. program, we recognized that we needed to have a genuinely *group* deliberation.

This brings me to a point about Rousseau and Rawls that I have not yet sufficiently emphasized, which is that their methods do not, strictly speaking, involve a *group* deliberation per se—indeed, they can do without it precisely because they envisage a perfect unanimity within the political group, both about the goals for the sake of which the group's agency is to be exercised, and about the basis from which deliberations about how those goals should be realized are to proceed. As I put it earlier, each member of the political group is supposed to be an exemplar for every other, with the result that each member's reasoning can stand in for the reasoning of any other. This is really a *limiting case* of a group point of view, because there is no distinction between the political group's point of view and the points of view of its individual citizens. But in cases where the members of a group do not enjoy such antecedent unanimity, no member's individual reasoning can

stand in for the reasoning of any other, or for the group as a whole. So if the group is to pursue its goals in a rational manner, it must undertake group deliberations which are carried out by the *group itself*.

It is somewhat easier to imagine what such a genuine group deliberation would involve by turning to a case of scientific inquiry.[10] When Robert Oppenheimer was asked to direct the Manhattan Project, General Groves—the army officer in charge—wished to keep the various teams of scientists who were working on different parts of the project completely isolated from one another for security reasons. But Oppenheimer insisted that this would impede scientific progress—because the different teams of scientists needed to communicate, in order to work out the significance of their various results together. In this same spirit, he also instructed the scientists themselves that they must desist from coming to their joint meetings armed with prior conclusions, with the aim of convincing their colleagues. If they did this, their meetings would descend into *argument*, whereas what scientific progress required was that the scientists work out the joint significance of all of their respective findings—their *all-things-considered* significance. In effect, then, Oppenheimer imposed a commitment to the ideal of rationality that I described above, in the light of which the members of the Manhattan Project could deliberate together as a single group agent with a single group point of view. And this did *not* involve each scientist carrying out this deliberation by herself, and serving as an exemplar for identical deliberations on the part of every other member of the group, in the way that political theorists so often envisage following Rousseau and Rawls; it involved the scientists deliberating *together*. Yet I don't mean to imply that it would be impossible for a single scientist to think all of the thoughts that constituted the group's deliberations. This *is* possible. In fact, it is in the nature of the case, for it is a general feature of thought and agency that anything that *one* agent can think can also be thought by *another* agent. So if what I'm calling the group agent that carried out the Manhattan Project can think all of the thoughts in question, so might an individual scientist. The point remains, however, that the individual scientists involved in the Manhattan Project were not in a position to do all of this thinking on their own. In their particular circumstances, what was required for progress on their project was that they should *pool* their various findings and hypotheses, even though none of them would have been prepared to deliberate from that large set of pooled considerations if they had been working on their own—partly because some were simply ignorant of what the others' findings and hypotheses were, and partly because some of them disagreed about some of them. While such a lack of convergence and agreement might, in many circumstances, lead a group that is faced with having to make a decision to settle it through a vote, that would have been an outlandish thing for the members of the Manhattan Project to do. They were faced with questions that required *reasoned* answers, which is what they aimed to give under Oppenheimer's direction, through genuinely group deliberations, which were carried out from a group point of view that was clearly distinct from the points of view of its members—since it was constituted by everything that each of them brought forward as relevant to the deliberative questions before them.

[10] I've discussed the example to follow in Rovane (2005).

Let us try to imagine what it would be for the members of a philosophy department to go to a department meeting, and to address the question of how best to design and implement a Ph.D. program in the spirit that Oppenheimer was urging on the scientists who participated in the Manhattan Project. Rather than each department member deliberating beforehand from her own point of view, and arriving at a private conclusion about what the department should do with the aim of persuading her colleagues, the aim would rather be this: the members of the department should bring relevant considerations forward, so that the department can pool them together for the purpose of working out their all-things-considered significance. This is not correctly conceived as an antecedent *agreement* about what the basis for our deliberations should be, since we could not even *know* what would be brought to the table until we began to discuss the matter. The *pooling* of considerations is a quite different act from arriving at a unanimous agreement among the members of a group; it signals a commitment on the part of the group as a whole to deliberating from those considerations together, regardless of how the various members of the group might privately view them.

If a philosophy department is to proceed in the way I'm suggesting, its members must come to meetings prepared to use their rational capacities together, so as to work out the all-things-considered significance of all that is put forward as a relevant consideration. As they do this, they will have to bracket what we might naturally refer to as their 'individual' perspectives on the matters before the department. So for example, in the case of the Columbia department's actual deliberations, I happened to be one of the faculty members who was against imposing any standardized requirements on Ph.D. students aside from the production and defense of a dissertation that makes an original scholarly contribution in the field of philosophy. This is not to say that I didn't think students should take lots of courses and learn lots of things. It is only to say that there is no one course requirement or certification that I thought every student should satisfy regardless of the topic of her dissertation project—such as a logic requirement, a foreign language requirement, a history requirement, or a distribution requirement, or any comprehensive or topical exams. However, I recognized that I was not dictator of the Columbia Philosophy Department. And once my colleagues put on the table all of the possible requirements for the Ph.D. that the department should consider, it was clear that the ensuing deliberations would have to weigh the merits of a program with no requirement beyond the dissertation against all of the other possible requirements that might be imposed, and indeed all of the other possible combinations of those requirements; and furthermore, the merits of these different requirements would have to be weighed in the light of a much broader set of considerations than I myself was bringing to the table, which included all of the considerations that had been brought to the table by all of the members of the department.

One particularly thorny issue for the department concerned the foreign language requirement. Here are five conflicting considerations that arose in connection with this issue: first, it seems desirable that students should not be confined in their research, or indeed in their direct philosophical exchanges at conferences, to what is available in English; but second, it is not clear how much good is to come of being

certified in a foreign language unless a very high level of proficiency is required for certification; third, a strong requirement would take significant time away from more directly philosophical studies; fourth, the higher the level of proficiency the program requires, the fewer languages can be required; fifth, if the requirement is reduced to just one foreign language—so that as high a level of proficiency can be required with as little compromise as possible to other work for the Ph.D.—this contradicts the spirit of the first reason for the language requirement, which is to ensure that students not be confined in their research and interactions by their inability to communicate in foreign languages. I'm going through these details in order to clarify that my *department* might have to address all of these matters even though if I were dictator I would just set them all aside. But more than that, I also want to clarify that when *I*, Carol Rovane, go to a department meeting, a portion of the intentional life associated with Carol Rovane's brain and body can help to constitute deliberations about all of these matters, where these deliberations do not proceed from Carol Rovane's personal point of view, but from a much larger point of view that comprehends all that the members of the department have put on the table as relevant to the department's deliberations about what requirements to impose on the Ph.D. In these larger deliberations of the whole department, the department will need to strive for *consistency* in the face of the wide diversity of considerations it is taking up. So while it may already be difficult to sort out what the foreign language requirement should be, the department will also want to ensure that whatever *reasons* finally justify it in imposing whatever foreign language requirement it decides upon, those reasons should be consistent with the reasons for its decisions about the logic requirement, the history requirement, the distribution requirement, the exams, etc. (For example, if it is decided that a foreign language requirement is useless unless a high level of proficiency is required, does the same reasoning apply to the case of logic and why or why not?)

If a department proceeds in the ways that I'm recommending, it will obviously not be subject to the Discursive Dilemma that Pettit and List discuss, nor any of its variants. Furthermore, there will be no difficulty about the group losing sight of various logical relations that hold among the considerations and options before it, since the whole purpose of deliberation is to attend to those relations, in order to work out the all-things-considered significance of all that it thinks. Finally, there will be no problem of having to determine a specific *order* in which the options before the department are to be decided—which is one of the ways in which Pettit and List suggest a group might avoid illogicality, by voting on premises first and then letting logic, rather than another vote, determine conclusions. The reason why there is no problem of order is that the ideal of arriving at all-things-considered judgments is an ideal that directs us to seek an overall *balance* of considerations through a process that is *holistic* in character. Not only is it the case, then, that all of a department's options need to be weighed together, but also, when various considerations are adduced for and against those options, their relative priority and importance is always up for re-assessment, and the process does not end until the *joint* significance of *all* such considerations has been worked out—which, as I've said, is what an all-things-considered judgment articulates.

To repeat, I am claiming that when Carol Rovane goes into a meeting of the Columbia Philosophy Department, a portion of the intentional life associated with Carol Rovane's brain and body can help to constitute such a departmental deliberation. I think it should be intuitively clear that the human *capacity* to do this is bound up with some related social capacities. One thing that we 'rational agents' are all able to do is to imaginatively project ourselves into points of view not our own, in order to work out what others have reason to do by their own lights. This is something we do whenever we try to engage in *internal* as opposed to *external* criticisms of others' points of view. I want to suggest that when human beings join efforts in order to engage in group deliberations, they do something similar—only it is not a use of projective *imagination*, because it is the carrying out of an *actual deliberation*. In order to carry out a deliberation from the point of view of the group, the rational capacities that human beings possess, by which they can *in general* work out what follows from what, and what the joint significance of a certain set of considerations is, are exercised in such a way that these things are worked out at the level of a whole group. When human beings do this, they literally *forge* a group point of view, which is not identical with any of their 'individual' points of view—so long as none of them is a dictator of the group. To thus *forge* a group point of view, as opposed to merely *imagine* one, is to constitute a process of deliberation that eventually issues in choice and action—the choices and actions of the group itself.

Would I be rationally incoherent if *I* were to *let go of* a portion of what would otherwise be *my own* intentional life to deliberations carried out from the department's point of view—given that I believe that the department will not choose as I would have chosen were I dictator of the department? *No*. I may believe that it is better that the department exist than not; and I may believe that it is better that the department proceed rationally rather than vote; and I may believe that it is better that the department proceed rationally from its *own* group point of view rather than be dictated to by any of its members, including myself. The upshot is that I must recognize that an *interpersonal* relation obtains between myself and the department; and insofar as some of the intentional activities of 'my' body and brain are carried out from the department's point of view, they are literally *not mine*. This is just the sort of rational fragmentation to which I was referring at the start.

When we think of the identities of reflective rational agents as metaphysically given, it may seem that these last claims can't be true. What seems true is that *I* have an entirely separate domain of intentional control that is set by the biological conditions of my existence as an individual animal, and everything that goes on within this domain of my intentional control belongs to *me*, and is something for which I bear personal responsibility. It seems to follow from this metaphysical individualist view that I cannot possibly bear an interpersonal relation to anything that goes on within 'my' body and brain. Furthermore, it also seems that if I could somehow 'give over' a portion of the intentional activities associated with 'my' body and brain to the life of the Columbia Philosophy Department, *I* would still be the one who is responsible for those activities. This seems to be confirmed by the following reflection: if I ever came to disapprove of what the Columbia Philosophy Department is doing, I would always have the option of *withdrawing* my

contributions to its intentional life; and on the assumption that I am a *natural person*, I would bear personal responsibility for *my own* contributions to the group's life, and so I *ought* to withhold them when they cannot be justified from my own—as opposed to the department's—point of view. It might seem that additional support for this metaphysical individualist view can be drawn my own example of the Manhattan Project. Although all of the scientists who signed on were clearly committed to scientific research in atomic physics, after a point many of them began to have political reservations about participating in the creation of an atomic bomb, and some of them felt obliged to withdraw from the project out of personal conscience. This seems to be the spirit in which Pettit and List advocate their combination of metaphysical and normative individualism. They suggest that it can help to keep group agents in moral line if their members are generally vigilant about taking personal responsibility for their own involvements in group agency, rather than just 'going along' with the group—especially in cases where the group's actions do not appear justified from their own, 'personal' points of view.

In spite of the apparent attractions of metaphysical individualism, I hope it is clear, by this point, why I regard it as metaphysically mistaken. In the next section, I'll complete my argument against it by articulating its more general lesson, which is that the existence of *any* rational agent, no matter what its size, is always a product of effort and will. But before completing my argument, I want to note a respect in which I find the suggestion of the last paragraph morally inadequate. However appealing the suggestion might seem to be—that human beings are always ultimately responsible for their 'individual' contributions to group agency—the policy of withdrawing from group agents who are not acting well leaves the following glaring problems completely unaddressed: first, when human beings are determined to keep their own hands and conscience clean by withdrawing from group agents when the latter do not act well, this often leaves the group agents intact and able to continue acting badly; second, in cases where a group really does function as an agent in its own right, what the group agent does is not an expression of the point of view of any of its human constituents, and so by focusing on the individual responsibility of its human constituents for their individual contributions to group agency, we lose site of the real *locus* of responsibility which is the group agent itself; third, just as a group's actions do not proceed from the individual points of view of any of its human constituents, likewise, those human constituents do not have intentional control over what the group agent does. Surely, what *realism* about group agency invites us to seek is a way of holding a group agent *itself* responsible for what lies within *its* intentional control. We make no headway on this matter by retaining residual commitments to forms of metaphysical and normative individualism which, I think we can now see, are in serious tension with realism about group agency.

3 Why the Existence of an Agent is Always a Product of Effort and Will

At the start of the last section, I claimed that all thoughts and actions must proceed from a particular point of view, which I described as a *deliberative* point of view.

The question I want to take up in this section is, how do such deliberative points of view come to be?

If we think of human beings as *natural persons* as Pettit and List do, we will think of them as having points of view that are somehow *given* to them—perhaps with their very first experiences, or perhaps more gradually in the course of their cognitive development. In a sense these thoughts are correct—only we need to bear in mind that there are three quite different conceptions of what a point of view is, and this bears on the question how a point of view comes to be.

One conception is presupposed by the very idea of consciousness: whenever there are conscious episodes, there must be a *phenomenological* point of view from which those episodes are apprehended in, or through, consciousness. I suppose it is arguable that a human being must have such a phenomenological point of view as soon as it has its very first experiences—though I have no particular stake in denying what philosophers in the Kantian and phenomenological traditions hold, which is that consciousness is not possible unless a very rich set of cognitive conditions are met, which are surely not met by a newborn baby when it first experiences hunger, satisfaction, etc. It does not affect my arguments about agency how this particular issue gets sorted out.

A second conception of a point of view is the *bodily* point of view from which an animal perceives and moves. Such a bodily point of view requires a specific form of cognitive organization, in which the contents of an animal's current perceptions reflect its current bodily movements, and in which its current bodily movements are guided by its current perceptions. This cognitive organization also requires a special way of representing one's own body which is, in a low-grade sense, first personal—it singles out one's own body as providing the spatial location *from which* one perceives other things (as in front of me or behind, as to my left or right, as near to me or far, as moving with respect to me or stationary, etc.). It seems pretty obvious that a human infant is not born with a ready-made bodily point of view of this kind, but comes to have one as it learns to control and move its body; yet it also seems overwhelmingly plausible that coming to have such a bodily point of view is a natural and unavoidable outcome of a human being's normal biological and cognitive development, and so it is not a product of effort and will in the sense that I shall be arguing a deliberative point of view is. I think the same can be said of a phenomenological point of view, regardless of the position we take on how much cognitive development is required for having one. For regardless of whether we do or do not allow that a newborn infant counts as having one just by virtue of being sentient, the point remains that a human being comes to have one as a part of its normal cognitive development.

The philosophical debate about personal identity that Locke inaugurated raises the following question: is a given phenomenological point of view necessarily rooted in a given human being's life, in such a way that it necessarily coincides with that human being's bodily point of view? When Locke distinguished personal and animal identity he argued that this is not so, on the ground that the same personal consciousness can persist in a new and different animal body, and perhaps without any body at all. In making this argument he was further developing and qualifying the position for which Descartes had argued in the second *Meditation* when he

claimed that he was not a man but a purely thinking thing. On both sides of this debate—the Lockean side and the opposed 'animalist' side that equates the person with the human being—it has generally been assumed that what I am calling a *deliberative* point of view must coincide with a phenomenological point of view; and it goes together with this assumption that deliberation is a *conscious* process, in which the sort of rational unity that an individual agent achieves when it deliberates is always to be achieved within a single, unified consciousness. But, of course, the argument for group agency calls this directly into question, by pointing to a case in which such rational unity can be achieved *without* phenomenological unity. Furthermore, if I am right that group agency tends to produce rational fragmentation, then just as the phenomenological unity of consciousness isn't necessary for the rational unity of an agent (in the group case), likewise, it isn't sufficient either. This invites the following, admittedly controversial, conclusion: although normal cognitive development may ensure that each human being comes to have its own phenomenological point of view and its own bodily point of view, it does not ensure that each human being comes to have its own deliberative point of view.[11]

But if the possession of a single, unified deliberative point of view within the human being is not an inevitable outcome of a process of normal cognitive development, then how does such a point of view come to be? I answer that *any* deliberative point of view, no matter what its size, comes to be *through* the process of deliberation itself. What human beings are born with is a *capacity* to deliberate, which cannot be exercised at all until fairly late in its cognitive development; and insofar as human beings are born with a capacity to deliberate, they are also born with a *potential* to form a deliberative point of view by coming to recognize different deliberative considerations as things to be taken into account together. The acts by which such considerations are recognized as things to be taken into account together are the first steps in the deliberative process, and it is really through them that these various recognized considerations—which are really just *thoughts* about various matters of fact and value—come to constitute a single point of view from which deliberations proceed, the aim of which is to work out their all-things-considered significance. This is exactly how I characterized the case of a philosophy department's group deliberations. When a department pools various considerations from which its deliberations are to proceed, it thereby begins to have its own deliberative point of view; and once it follows through on deliberating from that deliberative point of view it passes the social test for personhood, as I explained in the last section—it can be engaged in distinctively interpersonal ways about why it does what it does. My claim, then, is that this is so not just in the case of group agents, but in all cases of fully reflective rational agency.

On the account of agency that I'm offering, human beings begin as wantons who don't deliberate at all, and subsequently, they come to be sites of deliberation, through which particular deliberative points of view begin to emerge. But why should human beings ever come to exercise the deliberative capacities through

[11] For a sustained investigation into the first person, including the psychological pre-cursors to the first personal thought of fully reflective agents, see Burge (2011)—though his ultimate conclusion is precisely the view against which I am arguing in this paper.

which this happens? A partial answer to this question is, *simply because they can*. Thus if we ask, why do human beings ever begin to walk and talk, part of the answer is surely that it belongs to their nature to have and exercise their native capacities for walking and talking, and to give this answer is tantamount to saying that they do these things simply because they can. But human beings also have specific goals for the sake of which they walk and talk, which lead them to walk here rather than there, and to say one thing rather than another. The same holds for deliberation. If we ask, why should human beings ever exercise the deliberative capacities through which different deliberative considerations can be taken into account together, a partial answer is, simply because they can. But a more complete and informative answer will point to the fact that deliberation makes it possible to embrace and pursue certain sorts of goals—goals that can be realized only through *coordinated* thought and action. As human beings mature, their first steps towards forming a deliberative point of view typically involve the coordination of their thoughts and actions over time. I'm not referring to prudence—doing things now to ensure that future desires will be satisfied. I have in mind carrying out coordinated long-term activities such as building things, or engaging in certain forms of training (in athletics, dance, music, etc.), or earning and saving money in order to purchase something. Commitments to these sorts of projects are not compatible with wantonness—always acting on one's current impulse—because they require one to consider whether acting on a current impulse is compatible with carrying out the project. In addition to this task of assessing the worth of acting on current impulses in the light of one's other, longer-term goals, there is another way in which an agent who is engaged in long-term projects must take things into account together—namely, to work out what steps of the project need to be taken now, bearing in mind what one has already done, and bearing in mind as well how one's past and present efforts will require to be completed by future efforts. So my point is, longer-term projects provide reasons to deliberate, and once the process of deliberation begins, there arises within the human being a temporally extended deliberative point of view that encompasses more than 'current' thoughts. However, I also want to emphasize that the deliberative points of view that initially emerge within human lives are not very long-lived, and could not properly be characterized as points of view from which a whole human life is lived—that would come only with a recognition of the possibility and feasibility and worth of certain *life projects*, such as a career and family.

The capacity to envisage such life-projects, which would provide reasons why a human being should deliberate as one across the whole span of its biological life, stands alongside the capacity to envisage group projects that would give groups of human beings reasons to deliberate as one within the whole group. In these cases, considerations are taken into account together in order to afford coordination of thought and action across bodies as well as across time, for the sake of realizing goals that require the coordinated efforts of a whole group. When this happens, a kind of rational unity is achieved—or at least strived for—which signals the presence of a single deliberative point of view, that can be addressed and engaged in interpersonal ways.

In spite of all that I have said, the claim that *all* agents come to be through effort and will may seem paradoxical. It might be thought that there must be *original* agents through whose effort and will agents of *other* sizes can come to be. It might also be thought that if there are such original agents then it would be a violation of their nature to alter or destroy them; in particular, it might be thought that if human beings are such original agents then it would be a violation of their nature to alter or destroy them through the more fragmented forms of existence that I have claimed follow upon involvements in group agency. These thoughts are of a piece with Pettit's and List's claim that human beings are natural persons; and these thoughts also underlie various recommendations that Pettit and List make about group agency, which are aimed at preserving various aspects of human individuality—so that human beings remain loci of individual rights and responsibilities even as they exercise their agency to form group agents. But as human beings come to exercise their deliberative capacities, in order to coordinate thoughts and actions for the sake of realizing various goals, there is no metaphysical or natural necessary that dictates that they must forge rational unity within the biological boundaries set by their existence as separate organisms; and so there is no metaphysical or natural necessity that they must forge deliberative points of view that fall one-to-one with human bodies and human centers of consciousness. In a way, this is a claim to which Pettit and List are already *implicitly* committed by virtue of their realism about group agency. They simply haven't understood that it is a perfectly general claim, to the effect that it lies within human capability to form agents of many different sizes, smaller as well as larger than human size.

Of course, if a given human being were to fail to achieve rational unity within her life overall, at least in some significant degree, then certain characteristically human projects would be foreclosed, such as certain sorts of careers and personal relationships. Yet there may be other worthwhile goals that could be accomplished by foregoing precisely these projects. Some of these goals may involve participations in various forms of group agency; and in some cases they may demand more of their human constituents than philosophy departments typically do—think of military service for example, or certain political movements, not to mention the Manhattan Project. Other goals may require less than even one single human life, and in such cases we may find that a human being is not the site of one, sustained rational point of view. Consider, for example, a human being who lives in a somewhat fragmented social condition—who lives in a very traditional immigrant home in New York City, and who must also navigate the very different social space of an elite New York City public high school, and comes to pursue the sorts of projects supplied by that academic environment. It may seem like hyperbole to suggest that such a human being may not have one single deliberative point of view from which she acts, but it is not. Such a human being may literally be deliberating from quite distinct points of view when she is at home and at school. It makes no odds that certain basic beliefs and desires may be shared by both, since it is typical that a great many beliefs and desires are shared by distinct agents. What makes such agents distinct is supplied by the social test of personhood—of having an engageable point of view. The point of view presented and available for engagement

in the traditional home may be quite distinct from the point of view presented and available for engagement in the public high school.

This last example puts in mind a somewhat different claim for which Elizabeth Anderson has argued, namely, that reasoning requires a prior identity, which is given by membership in a group that supplies us with norms from which to reason and act.[12] As it happens, Anderson portrays the imparting of norms to group members as a form of collective agency through which a group is able to solve certain problems. But this is obviously a very different account of group agency from the one that I have developed in this paper, since there is no suggestion that when groups exhibit this form of collective agency they qualify as individual agents in their own rights—on the contrary, the point of calling it *collective* agency is to emphasize that it is really being exercised at the level of individual human agents through their embrace of common norms.[13] All the same, Anderson's picture of group agency leads her to make a further claim about multiplicity that is very close to the one I just made in connection with the example above, for she holds that when a human being belongs to more than one group, she becomes the site of multiple identities—and that is precisely how she would view the child of traditional immigrant parents who attends an elite public high school in New York City.[14]

Here is Pettit and List's response to Anderson's claim about multiple identities: "It is little short of comic to suggest that we are each an arena in which such different identities have autonomous voices." (p. 197) There is a ragbag of terms that philosophers use against views they don't like when they are short on argument. "Absurd" and "insane" are perhaps the most common, but "comic" surely counts among them. Insofar as Pettit and List really do find Anderson's suggestion comic, any laughter on their part ought, I think, to be dismissed as *nervous* laughter in the face of a conclusion to which *their* arguments should have led them, if only they had had the courage to follow their logic all the way to the general lesson that I have drawn—namely that their *own* individual identities are not metaphysically given but are, rather, contestable and re-negotiable.

Once we understand this point, we must take care how we express ourselves. What does it really mean to deny that "we" are "each" an arena in which multiple agents have autonomous voices? If Pettit and List had meant to be referring to individual *agents* then they would have been right, for *by definition* an individual agent possesses a kind of rational unity that precludes multiplicity. Since this definitional point holds for agents of every size, it follows that a group agent is no

[12] See Anderson (2001).

[13] Although I would be less quick to attribute any form of *agency* to identity-conferring groups, I am sympathetic to Anderson's claim that many of the norms by which we reason are given to us with membership in groups—though I would put the point in terms of being located in a moral context. See Chapter Four of Rovane (2013) for further discussion.

[14] If I read her correctly, her primary concern is not to emphasize that human beings may be the site of rational fragmentation, even though this does follow from her view; her primary concern is to criticize the economists' picture of the rational individual. Her concluding point is that if a human being's primary identity is given by being a woman, then she might not emerge as someone who reasons as an individual in the economists' sense unless she first possesses *other* identities, and then in the process of navigating these other identities she gains some distance on the norms imparted to her through her identity as a woman.

more an arena in which different identities have autonomous voices than an agent of human size is. In fact, this definitional point goes to the heart of why it is so hard to reconcile the phenomenon of group agency with sort of metaphysical individualism that Pettit and List wish to retain. When human beings achieve rational unity together, they forge a unified group point of view *within* which there can be no separateness of viewpoint. (Though it is also true that if the group endeavor does not wholly absorb its human members, then each of those human beings may be the site of another point of view that is smaller than human size, and also separate from the group's point of view—and of course, the definitional point would hold for those smaller-than-human points of view as well).

In the passage I quoted above, Pettit and List are not trying to underscore this definitional point about how *any* agent is bound to exhibit rational unity, no matter what its size. They are trying to make a point about *individual human beings*—that *they* are not arenas in which multiple identities have autonomous voices. But this is simply not correct. The correct thing to say about human beings is much more complicated: each human being *can be* the site of a single agent, but it *need not* be, because it *can also be* a site of multiple agents *and* a site of intentional activity belonging to a group agent. Thus, what I have argued *in all seriousness* is that the possibility of multiple agency is one among many human possibilities, and it is of a piece with the possibility of group agency which is the topic of Pettit and List's recent book. If a group agent has its own autonomous voice that is not reducible to the voices of its human constituents—as they seem to want to allow—then it cannot be that each human being is a separate arena of intentional activity that is wholly governed by just one autonomous voice.

References

Anderson, L. (2001). Unstrapping the straightjacket of 'preference': A comment on Amartya Sen's contributions to philosophy and economics. *Economics and Philosophy, 17*, 21–38.

Arrow, K. (1951). *Social choice and individual values*. New York: Wiley.

Burge, T. (2011). The Dewey lectures 2007: Self and self-understanding. *Journal of Philosophy, CVIII*(6/7), 287–315.

Maskin, E. (2001). Is majority rule the best election method? *Occasional Papers in the School of Social Science*, *11*, 1–8.

Pettit, P., & List, C. (2011). *Group agency: The possibility, design and status of corporate agents*. Oxford: Oxford University Press.

Rawls, J. (1971). *A theory of justice*. Cambridge: Harvard University Press.

Rousseau, J. (1997). *The social contract and other later political writings of Rousseau*. Cambridge: Cambridge University Press.

Rovane, C. (2005). What is an agent? *Synthese, 40*, 181–198.

Rovane, C. (2013). *The metaphysics and ethics of relativism*. Cambridge: Harvard University Press.