

The Intentions of a Group

Michael E. Bratman 10-30-15 forthcoming in Eric Orts and Craig Smith, eds., The Moral Responsibility of Firms (Oxford University Press)

Three questions

In my previous book on *Shared Agency*, I develop a plan-theoretic model of shared intentions of a group, for example the shared intention of a group to paint the house together, or to rob a bank together, or together to help curb an outbreak of a disease (Bratman 2014). The model is broadly individualistic in the sense that it sees the proper functioning of a certain structure of the plan-like attitudes of the participants - attitudes with certain contents, certain inter-relations, and in a certain context -- as sufficient for forms of social functioning characteristic of shared intention. In the cases on which this model focuses, what makes it true that the group intends, say, to rob the bank is that there is an appropriate inter-connected structure of relevant plan-like states of each of the participants, plan-like states that concern, inter alia, the group's bank robbery. In such cases, if we want to consider the targets of judgments of moral responsibility for robbing the bank, it can seem natural to suppose that those targets are simply the individual participants, though what they are each responsible for might well be affected by the facts underlying the sharing. (One can, for example, be responsible for participating in a conspiracy.) In such cases it can seem natural to understand a judgment that the group is morally responsible for the bank robbery as a summation of judgments about the responsibility of each of the individual participants. After all, insofar as the group intends to rob the bank, what is true is just that each of the participants has relevant intentions -- ones that concern, inter alia, the bank robbery -- and these intentions are, in the context, inter-connected in relevant ways.

Matters would seem potentially different, however, if there were cases in which the group has intentions that are not themselves shared intentions. If the group intends, say, to curb the outbreak of a disease in a certain town, and if this is not itself a shared intention of the group, then perhaps there is room for the idea that the group itself is appropriately subject to relevant moral praise (or, in a different case, blame), and in that sense is appropriately held morally responsible for the cited activity.

In saying this I am taking for granted, without argument, the idea that in central cases S's moral responsibility for X is grounded, at least in part, in the role of associated intentions of S concerning X, in S's relevant "quality of will". (Strawson 2003). A second idea that I will take for granted without argument, and that will figure in my discussion below, is that S's intentions provide particularly significant support for such attributions of responsibility when the relevant guidance by those intentions constitutes S's governing her own actions, when the relevant guidance constitutes S's self-governance.¹

I do not say that the functioning of such self-governance-related intentions suffices for relevant moral responsibility. My assumption is that such intentions are basic to moral responsibility, not that they are the only condition of such responsibility. My focus in this Chapter will be on the possibility of group intentions that satisfy this basic condition for group responsibility. I leave open the question whether there are further necessary conditions for such responsibility.

¹ See T.M. Scanlon's appeal, in his discussion of Strawson's approach to responsibility, to "the capacity for critically reflective, rational self-governance ... [a capacity that] is not specifically moral." (Scanlon 2003, 368-9)

This leads me to the three questions I aim to discuss in this Chapter. First, does the plan-theoretic approach to shared intention provide resources for modeling robust group intentions that are not themselves shared intentions? Second, if there are such intentions of a group will their normal functioning at least prima facie help constitute the group's relevant self-governance? Third, what is the implication of our answers to these questions for the possibility of group responsibility?

Concerning my first question, I do not say that a negative answer would establish that there are no group intentions that are not shared intentions. Perhaps there are other routes to a coherent model of such group intentions. Perhaps we can look directly at the overall functional organization of a group. And perhaps when we look directly at that group-level functional organization we can sometimes justifiably ascribe intentions to that group as a part of an overall theory of the group's diachronic functioning that has significant explanatory power.² We thereby try to see group intentions as functional (or functionally specified) states within a group-level functioning system. And to do that we need not identify these group intentions with corresponding shared intentions among members.

In contrast, the strategy of *Shared Agency* was to try explicitly to construct relevant forms of sociality and group functioning out of the materials of inter-connected individual planning agency. In that book I tried to show the extent to which we can build characteristic forms of group functioning out of characteristic forms of individual functioning, given distinctive aspects of that individualistic functioning. And my first

² This is, broadly speaking, the approach of (List and Pettit 2011). As they say, they draw on “a broadly functionalist theory of agency”.(75) (Though in this book List and Pettit focus primarily on beliefs and desires of a group, rather than intentions.) And see also (Pettit 2003).

question in this Chapter is whether, in articulating those materials of inter-connected individual planning agency, we thereby have the resources to construct the infrastructure of robust group intentions that are not shared intentions. My answer will be that we do, but that in acknowledging these group intentions we need to give up the idea that such intentions must be embedded in a robust holistic structure of attitudes of that very group.³ It will sometimes suffice, instead, that the intentions of a group are embedded in a larger structure of attitudes that to a significant extent involves attitudes of the individual participants. And such group intentions can indeed play a central role in the self-governance of the group.

With these results in hand I will turn briefly at the end to their bearing on the possibility of group responsibility.

The planning theory of shared intentionality

In the background of my discussion will be the plan-theoretic model of shared intention and shared intentional action that I have developed in *Shared Agency*. So let me briefly sketch relevant aspects of that model here. I begin with the thought that an important feature of individual human agency is that it normally involves complex forms of plan-infused cross-temporal organization. Given limits in our knowledge and our mental resources, these plans will typically be partial, and will need to be filled in as time goes by. These forms of practical thinking will be guided by an (implicit)

³ List and Pettit write: “Let a collection of individuals form and act on a single, robustly rational body of attitudes... and it will be an agent. ...[I]t is possible for collections of individuals to coordinate their individual contributions so as to achieve this level of functioning. Hence [on the assumption that this possibility is realized] group agents exist.” (List and Pettit 2011, 75) I take it that the “single, robustly rational body of attitudes” to which they appeal is a body of attitudes of the group. So the group’s intentions, on this model, would be embedded in a robust holistic structure of attitudes of that very group.

acceptance of norms of plan rationality – norms of consistency, agglomerativity, means-end coherence, and stability. In this sense there is a distinctive rational dynamics of planning agency. And intentions are states in this planning system.

The next step is to describe inter-related planning structures of participating individuals such that the rational functioning of those planning structures – functioning in accordance with the rational dynamics of individual planning agency -- would constitute the rational, social functioning that is characteristic of shared intentional activity. This would be to show that our planning capacities are a common core that lies behind both our capacity for cross-temporal organization in individual intentional agency and our capacity for social organization in our shared intentional activities. (This is a version of the thought that there is a significant parallel between the temporal structure of an individual's life and aspects of inter-personal sociality.) I aim in this way to articulate sufficient conditions for robust forms of shared intentionality in terms of these inter-related, individualistic planning structures; but I leave open the question whether these conditions are strictly necessary for shared intentionality. This is the strategy of sufficiency. And it supports one important aspect of the general thesis of the fecundity of planning agency, the thesis that our capacity for planning agency lies at the bottom of a range of basic practical capacities – in this case, our capacity for shared intentional activity.

In developing this model I suppose that, at least normally, shared intentional activity is explained by a corresponding shared intention so to act. I then describe inter-connected structures of individual planning agency whose functioning in accordance with the rational dynamics of individual planning agency would ensure the social-

psychological functioning that is characteristic of shared intention. To specify this social-psychological functioning, I draw on analogies with the roles of individual intentions in organizing individual thought and action. The characteristic social-psychological functioning of a shared intention to X will involve inter-personal coordination of action and planning in pursuit of X and in ways that accommodate the rational agency of each of the participants. It will also involve framing shared deliberation or bargaining concerning how the group is to X. And in the background will be analogous social norms – norms of social agglomeration and consistency, social means-end coherence, and social stability -- whose violation would normally undermine these social roles of shared intention.

My proposal is to articulate this structure by appeal, in part, to conditions along the lines of:

1. Each intends that we X.
2. Each intends that we X by way of each of our intentions that we X (in this sense these intentions inter-lock with each other) and by way of X-tracking mutual responsiveness in sub-intention and action, and so by way of sub-plans of each that mesh with each other. (Plans mesh when they are co-realizable. Plans may mesh even though they do not match.)
3. Each believes correctly that there is persistence interdependence between the intentions in 1.
4. All this is out in the open.

Putting aside some details, what I argue in *Shared Agency* is that when such a public structure of intentions of each functions in accordance with the rational dynamics of

individual planning agency, it thereby realizes the roles characteristic of shared intention in part by way of conformity to the associated social rationality norms. Further, when this structure of intentions of each leads to the joint activity by way of relevant mutual responsiveness there is shared intentional activity. I thereby highlight conceptual, metaphysical and normative continuities with individual planning agency, while also providing the resources to articulate central ways in which shared intentionality goes beyond simple individual planning agency.

A key is the appeal to the *intention* of each in favor of X by way of the other's intention, mutual responsiveness, and meshing sub-plans. It is not just that each intends his part in X and merely expects the other to play her part. So the rational pressure on each to make her plans coherent and consistent ensures rational pressure on each to support the success of the joint activity and the meshing role of the other in that activity. Each is thereby under rational pressure to coordinate with the other – perhaps by way of helping actions – and to avoid ways of acting that are incompatible with the joint activity. These rational pressures on each, pressures grounded in the rational dynamics of individual planning agency, induce, given these distinctive contents and inter-relations, pressures in the direction of social coherence and consistency and associated social coordination and effectiveness.

There can be such shared intentions even though each participates for different reasons. Further, since each intends the joint activity in part by way of the relevant intention of the other, and by way of mutual responsiveness and so meshing sub-plans, each is under rational pressure to seek to ensure that the sub-plans of each, agglomerated together, both are adequate to the shared task and do indeed mesh

interpersonally. So, in the absence so far of adequate, meshing sub-plans, a shared intention will tend to structure bargaining or shared deliberation (or the like) in the pursuit of such mesh.

This gives us a plan-theoretic model of those intentions of a group that are shared intentions. A group intends, say, to rob the bank, when its members share an intention to rob the bank together. And the plan-theoretic model of such a shared intention consists of the cited structure of inter-related intentions of the participants, in an appropriate context.

Shared policies of weights

The next step is to see how this plan-theoretic model can be extended to an important kind of shared policy, namely a shared policy to give certain weights to certain considerations within the group's shared deliberation.

Shared deliberation of the sort of interest here is a shared intentional activity in which the participants bring to bear common standards in their effort to settle questions about how to proceed within a shared intentional activity in which they are engaged. In a central case this will be a matter of bringing to bear shared commitments to weights in the shared deliberation. We might, for example, share a commitment within our shared painting of the house to our giving substantial weight to environmental concerns as we deliberate about which paints to use, and how to dispose of various materials. Or perhaps an academic department has a shared commitment to its giving weight to collegiality, or to issues of sub-field, in its shared deliberations about faculty appointments. Or perhaps a business group has a shared commitment to giving weight

to increasing profits, or to satisfying certain environmental concerns, in its shared deliberations.

I argue in *Shared Agency* that while such shared commitments to weights will normally be grounded to some extent in relevant evaluative judgments of the individual participants, sameness of value judgment, even in a context of common knowledge, is neither sufficient nor necessary for a corresponding shared commitment to weights in shared deliberation. And my proposal is that our shared commitments to weights are better modeled as shared intentions that favor our giving certain weights to certain considerations in relevant shared deliberation. Since such shared intentions will normally have a characteristic generality, they will be shared policies of weights – where policies are intentions with suitably general contents. We then extend the plan-theoretic constructivist architecture of shared intention to these shared commitments to weights: a shared commitment to give weight to R is a shared policy to give weight to R in relevant shared deliberation; and this shared policy consists in public, interlocking, and interdependent general intentions, on the part of each, in favor of our giving weight to R in relevant shared deliberation.⁴

Shared policies of weights are policies about how the members of the group are to think together in their shared deliberation. The next step is to consider other kinds of shared policies concerning how those members are to think together.

Shared policies of procedure

Begin by noting that shared deliberation can itself be embedded in yet a further structure of shared policies of procedure. A group might have a shared policy to turn to

⁴ For a more extensive discussion of these matters see (Bratman 2014, chap. 7). In (Bratman 2007d) I classify such shared policies of weights as a kind of shared valuing

shared deliberation – in contrast, say, with various forms of bargaining -- in sorting out issues concerning how it is to carry out a shared project. And such a shared procedural policy might help explain why the group is now engaging in such shared deliberation.

Shared policies of procedure might favor various group decision procedures.⁵ A shared policy might favor, say, a majority vote procedure. Or it might favor a consensus procedure characteristic of a Quaker meeting. And so on.

One kind of procedure that might be embraced within a shared policy is a procedure that authorizes a certain sub-group to settle certain issues for the group in its execution of certain shared activities. (Note, though, that such a policy of authorization need not insure legitimate authority.) A Philosophy department might have a shared policy of procedure for graduate admissions, one that authorizes a sub-committee to make the decision. And once such procedural policies of authorization are available they can iterate: a sub-group that is authorized by the larger group may have its own policies of authorization of a sub-sub-group, or even just a different group, to make certain decisions that it has been authorized to make.

One can extend a plan-theoretic construction to such shared policies of procedure, including shared procedural policies of authorization. As a first approximation, such a shared procedural policy can be modeled as a public, interlocking and interdependent structure of relevant policies of each that favor the

⁵ A point highlighted by Thomas Hurka in conversation. Appeal to such shared procedural policies is in the spirit of Peter French's work on "Corporate Internal Decision Structures." See (French 1984) esp. chaps. 3-4. Related ideas can also be found in (Tuomela, 1995), esp. 176-180; (Pettit 2003); (List and Pettit, 2011) esp. chap. 3; (Hindriks, 2008); (Hindriks, 2014); and (Ludwig, 2014).

operation of a given procedure (e.g., some voting procedure), where that includes follow-through with the output of that procedure.

Two potential elements in such follow-through can be distinguished. Follow-through may involve further shared intentional activity of the group, shared intentional activity that is guided by the outcome of the procedure. This will involve relevant interlocking, mutual responsiveness and intended mesh across the group. In contrast, follow-through may be distributed among the activities of certain individuals or sub-groups. Such distributed follow-through need not involve interlocking and mutual responsiveness across the overall group, and so need not involve shared intentional activity of the overall group itself.

A shared policy of procedure will normally favor follow-through that to some extent involves further shared activities of the group, shared activities that are guided by the outcome of that procedure. However, the shared policy may well also favor follow-through that consists in activities distributed across individuals and/or sub-groups. I will return to this difference below.

Group self-governance

So the plan-theoretic model can include shared policies of weights and of procedures, including shared policies of authorization. Intended follow-through with the outcomes of procedures favored by such shared policies may be to some extent a matter of further shared activities of the group itself, and to some extent a matter of distributed activities of individuals or sub-groups. This plan-theoretic architecture includes shared policies of weights and/or procedures as structures of relevant, interconnected policies of each of the participants. Such shared policies of weights and/or

procedures are compatible with divergence in the underlying evaluative judgments of the participants and with differences in the reasons for which each participates. Nevertheless, given their central role in the functioning of the group, it is plausible to see such shared policies as speaking for the group on relevant matters. In this way we extend to the group an idea that is important in our understanding of individual agency, namely: that certain attitudes are not merely wiggles in the agent's psychic stew but speak for the agent in the sense that their guidance is, *prima facie*, the agent's self-governance.⁶ And my proposal is that such shared policies of weights and/or procedures will normally be such that when they guide relevant thought and action the group is, in that domain, governing itself.⁷

In a particular case, if we share a policy in favor of weights or a procedure and you knowingly violate that policy, then you are in violation of norms of consistency and coherence on your own intentions and plans. I can appeal to that in criticizing you for your violation. I may also be able to criticize your violation as tending to undermine our group self-governance. Insofar as you yourself value group self-governance, this criticism will engage your own concerns. In this way the connection with group self-governance can help stabilize such shared policies. And this stability of such shared policies contributes to the claim that those policies do indeed speak for the group. So the contribution of such policies to the group's self-governance is self-reinforcing in the sense that this contribution supports the stability of those shared policies and that

⁶ This idea is rooted in work of Harry Frankfurt. See his reflections on "where (if anywhere) the person himself stands" in (Frankfurt, 1988) at 166. For my development of this idea see (Bratman, 2007b) and (Bratman 2007c).

⁷ And this supports the idea that participation in such shared policies ensures membership in the group.

stability in turn supports the status of those policies as attitudes whose guidance is, prima facie, the group's self-governance. Or at least this is true for groups whose participants value the group's self-governance.

So the planning model provides for structures of shared policies of weights and procedures whose guidance is, prima facie, the group's relevant self-governance and whose functioning is thereby potentially stabilized in a way that supports the role of those shared policies in the group's self-governance. I now want to argue that such shared policies of procedure can induce group intentions that are not shared intentions.

Procedure-based intentions of a group

Suppose a group is engaged in a shared intentional activity of providing medical aid in crises around the world. In the face of many claims on its resources it has a shared policy that specifies procedures for deciding whether to provide aid in a specific case. Let's say that this policy authorizes a sub-group to make a decision by way of its own, shared deliberation, followed by a majority vote. Suppose that in a particular case this procedure issues, by way of a split vote on the part of the authorized sub-group, in a decision to provide aid in a certain context C. Given that this is the output of procedures favored by policies that are shared by the participants, this output will be poised to guide relevant follow-through in the sense that given relevant knowledge and rationality on the part of relevant participants, the output would guide such follow-through. This follow-through may simply involve distributed activities of individuals or sub-groups, distributed activities carried out in accordance with the outcome of the procedure. But this follow-through may also involve further shared intentional activities of the group itself, shared intentional activities that are guided by the output of the

procedure. In both cases the outcome of the procedure guides thought and action (given knowledge and rationality), but in the second case it guides, in particular, shared thought and action of the group itself.

Now my proposal is that when the output of such a procedure is set to guide (inter alia) further shared intentional activities of that very group, the fact that the procedure has issued in that output ensures that the group itself intends to act in accordance with that output. In the present case, it is likely that the outcome of the procedure is set to guide, for example, further shared deliberations of the group concerning how to coordinate this aid with other aid the group plans to provide (though breakdowns in the guidance remain possible). If this is true then the group itself intends to provide the aid in context C: this is a procedure-based group intention.

Will this procedure-based group intention be a shared intention? Well, there may be members of the group who judge, at least prior to the outcome of the procedure, that it would be better for the group not to do this. But that does not show that the group's intention to provide aid in C is not a shared intention: a group can share an intention despite such differences in evaluative judgment. What stands in the way of identifying this group intention with a shared intention lies not in such evaluative disagreement but, rather, in the possibility of limited knowledge or rationality. Some members might simply be unaware of the outcome of the procedure. (This will be increasingly common as the group grows in size and complexity.) Given this ignorance, though the group intends to provide medical aid in context C there may not be a shared intention to provide this aid. Again, suppose some members are aware of the outcome but -- because they think that it is a serious mistake -- resist intending that the group provide aid in context C even

though they continue to participate in the general shared policies of procedure. This would be a breakdown on their part in plan rationality; but such breakdowns are all too possible. And given such a breakdown there will not in fact be a shared intention that corresponds with the group intention to provide the aid.

So if there is relevant ignorance, or relevant breakdown in individual plan rationality, there may not be a shared intention to provide aid in context C. But such ignorance or rational breakdown may not stand in the way of its being true that the group intends to provide this aid. The outcome of the procedure establishes an intention of the group so long as the outcome would guide relevant shared thought and action under conditions of relevant knowledge and plan-theoretic rationality. So the plan-theoretic architecture can make room, in this way, for procedure-based group intentions that are not shared intentions.

Granted, so long as there is no relevant ignorance or rational breakdown, the procedure-based group intention will correspond to a shared intention.⁸ Nevertheless, the procedure-based group intention will not always correspond to an actual shared intention; and the existence of the procedure-based group intention does not depend on the existence of a corresponding shared intention. In this sense procedure-based group intention and shared intention are separable. They are separable even though the infrastructure of procedure-based group intention systematically involves shared intention. The underlying structures and processes that make it true that there is the procedure-based group intention in favor of the outcome of the procedure need not ensure a corresponding shared intention.

⁸ This point was emphasized in conversation with Olle Blomberg, Brian Epstein, and Arto Laitinen.

Suppose that, in this procedure-based way, the group intends to provide medical aid in context C. What makes it true that the group so intends is a complex structure of shared intentions and policies, actual social procedures and their output, and resulting tendencies toward relevant guidance of (inter alia) relevant shared intentional activities and forms of shared thinking. Appeal to this plan-theoretic structure remains in the spirit of the general thesis of the fecundity of planning agency. But this structure does not itself ensure a corresponding shared intention to provide aid in context C.

Group intentions and dispersed beliefs and plans

But now there is a problem. The idea that the group itself intends to provide aid is in tension with a common idea about the holism of the mental. The idea of the holism of the mental is the idea that a mental attitude of S -- for example, an intention of S -- must be embedded in a substantial, holistically coherent web of related attitudes of S (Davidson 2001) (Rovane 1998). In particular, if we ascribe to S an intention to X we are committed to seeing that intention as part of a holistic web of attitudes of S, attitudes that include, for example, beliefs about the nature of X and how it might be done, intentions not to interfere with X and to take needed steps toward X, beliefs about the nature of such needed steps, beliefs needed to have the more immediately implicated beliefs, other pro-attitudes toward X and/or its expected upshots, and so on. But while there may be special cases in which there are sufficiently rich plan-theoretic structures to provide the infrastructure of such a robust holistic web of attitudes of the group, this does not seem to be ensured by the procedure-based infrastructure of the group intention that has been described.

But if the group itself does not have specific beliefs about how to execute its intention to provide aid, and plans about how to do that, how will that intention be effective? The answer is that the efficacy of this procedure-based group intention can involve the operation of a range of related beliefs and plans dispersed among the individual participants. For example, certain participants may know how to prepare and distribute a needed vaccine and have specific plans for doing this in their particular circumstances. It is not necessary to suppose that the group itself has this knowledge or this specific plan for preparing and distributing the vaccine (though it may). Although the group itself has the procedure-based intention to provide the aid, the operation of that intention can draw on relevant beliefs and plans dispersed amongst relevant participants, beliefs and plans that are responsive to their specific circumstances.⁹ We need not suppose that these dispersed beliefs and plans must be beliefs and plans of the group.

Granted, the specific plans of certain participants for preparing and distributing the vaccine are downstream outcomes of the procedures that are supported by relevant shared policies. And I have claimed that the basic outcome of those procedures – an outcome in favor of providing the aid – may be an intention of the group to provide the aid. So why not say that these specific plans are, as well, an element in a procedure-based intention of the group to prepare and distribute the vaccine in this specific way? The answer is that to induce a procedure-based intention of the group an outcome of a procedure needs to be poised to guide shared intentional action on the part of that

⁹ This is broadly in the spirit of F.A. Hayek's emphasis on "knowledge of the particular circumstances of time and place," knowledge "that is dispersed among many people". See (Hayek 1945) at 521 and 530. For the use of this idea in understanding the distinctive strengths of ancient Athenian democracy, see (Ober, 2008).

group. But the cited specific plans for preparing and distributing the vaccine are only poised to guide the thought and action of a sub-set of individuals.

I conclude that the functioning of procedure-based group intentions can draw substantially from a web of dispersed beliefs and plans of individuals or sub-groups that need not be elements in corresponding intentions of that overall group.

This is an extension of an idea implicit in our plan-theoretic treatment of shared intention. Suppose that you and I share an intention to paint the house together. For this shared intention to be effective each will need to have relevant sub-plans, and those sub-plans will need together sufficiently to specify preliminary steps, means, and the like so that in following through with those sub-plans we do indeed paint the house together. But the theory does not insist that all the sub-intentions associated with these sub-plans must be elements in corresponding shared intentions. In intending that our relevant sub-plans mesh we need not intend or expect that our relevant sub-plans match. And aspects of my sub-plans may not themselves be out in the open, or may not involve intentions of mine that interlock with corresponding intentions of yours. For example, my sub-plan for shopping for the paint at Ace Hardware need not be an element in a shared intention in favor of shopping at Ace Hardware. It will many times suffice that you know that some sub-plan or other of mine will solve the problem of how we are to get the paint. You do not need to know what that sub-plan is, nor do you need yourself to have an intention that matches and interlocks with my intention concerning Ace Hardware.

The further web of intentions that supports the efficacy of our shared intention need not be solely a matter of other shared intentions; it can, and normally will, involve

as well relevant sub-intentions dispersed among the participants, sub-intentions that are not themselves elements in corresponding shared intentions. And the situation with a procedure-based intention of a group is analogous in the sense that the further web of intentions that supports the efficacy of that intention of the group need not be solely a matter of other intentions of the group; it can, and typically will, involve relevant, distributed intentions of the participants, intentions that themselves need not be elements in corresponding intentions of the group. There will be rational pressure for all these various intentions to mesh within a consistent and effective overall web; but they can mesh without themselves being elements in corresponding intentions of the group.

We therefore should not insist that for the group to intend to provide the aid its intention must be embedded in a substantial, holistic web of attitudes of the group. It may suffice that it is embedded in a social context in which relevant attitudes of the participants do a lot of the work.

Between shared intention and a robust holism

This may seem to be an unstable result. We do normally expect the intentions of an individual to be embedded in a robust web of relevant attitudes of that individual. But our plan-theoretic construction of procedure-based group intentions allows for group intentions that are not embedded in a holistic web of attitudes of the group. So our plan-theoretic construction of procedure-based group intentions is in tension with a plausible view about the holism of the mental.

In “How to Share an Intention,” J. David Velleman addresses an analogous issue. Velleman begins with a Searle-inspired idea that an intention is “a mental representation that causes behavior by representing itself as causing it.” (Velleman,

2000 at 207) (Searle, 1983 chap. 3.) Now, Velleman has a distinctive view about why certain representations tend to cause their own fulfillment, a view that appeals to a purported concern with self-knowledge. But I want to put that view aside here and focus on an independent issue, one to which, as Velleman shows, we are led when we see intention as “a mental representation that causes behavior by representing itself as causing it.”

Suppose that you and I arrive at a shared intention to walk together, and we arrive at this shared intention by way of a normal conversational exchange in which you say “I will if you will” and I reply “I will”. About this pair of interdependent statements Velleman says:

Our statements ... combine to form a joint statement saying, in effect, that they will jointly prompt us to take a walk; and they jointly prompt us to take a walk, as they jointly say. They consequently add up to a single representation that causes our actions by representing itself as causing them -- a single token intention that is literally shared between us.” (Velleman 2000, at 217-8)

But how could a joint statement literally be an intention? Aren't intentions states of mind? If Velleman is right that this joint statement satisfies his Searle-inspired condition for intention, isn't that an argument against the sufficiency of that condition for intention?

Velleman's response is concessive: “if a commitment's being oral or written entails that it isn't mental, our commitment to taking a walk may not be a mental act or state at all.” (Velleman 2000, at 219) And the important point for my purposes here is

that this opens up the theoretical possibility of intentions that do not satisfy some further condition for being states of mind.

It is this general theoretical possibility – abstracted away from the details of Velleman’s theory -- that I want to highlight. I have described a plan-theoretic infrastructure of group intentions, one that is set to play standard roles of intention in guiding downstream thought and action of the group even though these group intentions need not ensure corresponding shared intentions. These group intentions may not be embedded in the kind of robust holistic web of attitudes of the group that we may plausibly suppose to be characteristic of mental states of that group. Nevertheless, these group intentions may still be poised to be effective in part because of the way in which they are set to interact with relevant attitudes of the participants. So we should recognize the theoretical possibility of procedure-based group intentions that are neither shared intentions nor embedded in the kind of holistic web of attitudes of the group that we might plausibly suppose is a condition of being a mental state of that group.

We can put the point in terms of the idea of a subject of an attitude. In one straightforward sense, if it is true that the group intends to provide the medical aid then the group is the subject of that intention. But there is also a stronger idea of a subject as a unitary locus of a robust holistic web of inter-related attitudes. And in this sense of a subject of attitudes, the group may fail to be a subject of its intention to provide the aid.

In this respect the situation parallels the situation Velleman aimed to describe in which (a) there is literally a single intention that is shared, namely the joint statement; (b) this joint statement plays the defining roles of intention, and so (c) we do indeed

intend to walk together; but (d) there is no single subject of that intention. (As Velleman says, the “joint making up of minds” that is involved in the joint statement “is not the making up of a joint mind.” (2000, at 219)) Similarly, the procedure-based group intention whose plan-theoretic infrastructure I have tried to describe need not be an intention of a group mind, in a sense of group mind that involves a robust holistic web of attitudes of the group. Nor need it be an intention whose subject (in the strong sense) is the group. So it is a mistake to suppose that a group intention is either a shared intention or an intention of the group that is embedded in a robust holistic structure of inter-related attitudes of that group. Certain procedure-based group intentions occupy the space between these two other possibilities. And the recognition of this intermediate space gives us important theoretical resources.

Even if these procedure-based group intentions are not embedded in such a holistic web of attitudes of the group, and even if the group is not, in a strict sense, the subject of those intentions, these group intentions may still be poised to be effective in part because of the way in which they are set to interact with relevant attitudes and activities of the participants. And, given that the background shared procedural policies normally speak for the group, the functioning of such group intentions will normally be a form of group self-governance.

So we arrive at answers to our first two questions. First, the plan-theoretic approach can model robust group intentions that are not themselves shared intentions; and such group intentions need not be embedded in a holistic web of attitudes of the group. Second, the normal functioning of such group intentions will normally help constitute the group’s relevant self-governance.

Group responsibility

This takes us to our third question, concerning group responsibility. And here the discussion of our first two questions suggests that procedure-based group intentions are candidates for the kind of intention-based grounding needed for attributions of responsibility to the group. This is because such procedure-based group intentions are set to provide relevant downstream guidance both of the thought and action of the group and of relevant distributed activities of individuals and/or sub-groups, and their playing these roles is, *prima facie*, the group's self-governance. And I have argued that all this can be true even if these group intentions are neither shared intentions nor embedded in a robust holistic web of attitudes of the group.

But what should we say about those cases in which the group intention does correspond to a shared intention? Well, our main reason for thinking that group intentions can provide the intention-based ground for group responsibility was that such intentions guide thought and action both of the group and of relevant individuals and/or sub-groups, and that their playing these roles can constitute the group's self-governance. Suppose then that a group has a procedure-based intention to X, that this intention plays these roles in its self-governance, but that, because of relevant ignorance on the part of some of the members of the group, there is not a corresponding shared intention. We have in this case the basic intention-based ground for group responsibility. Suppose now that the relevant information newly becomes available to all the members and as a result there newly comes to be a shared intention that corresponds to the procedure-based intention of the group. It does not seem plausible that this change would newly block the group's responsibility, though it might

affect the responsibility of relevant individuals. If the procedure-based group intention can provide the intention-based ground for group responsibility in the absence of a corresponding shared intention, then it seems plausible that it can continue to do so even when there is a corresponding shared intention.

So one fundamental condition for group responsibility – a condition of group intention that is poised to be involved in the group’s self-governance – can be realized in a procedure-based way. Since there can be such a procedure-based group intention in the absence of a corresponding shared intention, we can resist certain kinds of pressure to reduce relevant judgments of responsibility to judgments solely about the responsibility of each of the participants. And since there can be such a procedure-based group intention that is not embedded within a robustly holistic group mind, we can resist the idea that this basic condition of group responsibility requires such a holistic group mind.¹⁰

REFERENCES

Bratman, Michael E. 2007a. Structures of Agency: Essays (Oxford University Press)

Bratman, Michael E. 2007b. “Reflection, Planning, and Temporally Extended Agency” as reprinted in (Bratman 2007a).

Bratman, Michael E. 2007c. “Three Theories of Self-Governance,” as reprinted in (Bratman 2007a)

¹⁰ Some of the ideas in this essay were presented at the Insead-Wharton December 2013 Conference on “The Moral Responsibility of Firms: For or Against?”. A later version was presented at the University of Copenhagen October 2014 Conference on “Thinking (about) groups.” I benefited from comments from participants at both of these conferences. I have also benefited from written comments from Olle Blomberg, Frank Hindriks, and Philip Pettit, and related discussion with Christian List.

Bratman, Michael E. 2007d. "Shared Valuing and Frameworks for Practical Reasoning" as reprinted in (Bratman 2007a)

Bratman, Michael E. 2014. Shared Agency: A Planning Theory of Acting Together (Oxford University Press).

Davidson, Donald. 2001. "Mental Events," as reprinted in Davidson, Essays on Actions and Events (second edition) (Oxford University Press)

Frankfurt, Harry. 1988. "Identification and Wholeheartedness," as reprinted in Frankfurt, The Importance of What We Care About (Cambridge University Press)

French, Peter. 1984. Collective and Corporate Responsibility (Columbia University Press)

Hayek, F.A. 1945. "The Use of Knowledge in Society." American Economic Review 35: 519-30

Hindriks, Frank. 2008. "The Status Account of Corporate Agents," in Concepts of Sharedness; Essays on Collective Intentionality edited by H.B. Schmid, K. Schulte-Ostermann, and N. Psarros (Ontos Verlag)

Hindriks, Frank. 2014. "How Autonomous Are Collective Agents? Corporate Rights and Normative Individualism," Erkenntnis 79 1565-1585.

List, Christian and Pettit, Philip. 2011. Group Agency: The Possibility, Design, and Status of Corporate Agents (Oxford University Press).

Ludwig, Kirk. 2014. "Proxy Agency in Collective Action," Nous 48:1 75-105

Ober, Josiah. 2008. Democracy and Knowledge (Princeton University Press)

Pettit, Philip. 2003. "Groups with Minds of Their Own," in Socializing Metaphysics: The Nature of Social Reality edited by Frederick F. Schmitt. (Rowman & Littlefield)

Rovane, Carol. 1998. The Bounds of Agency: An Essay in Revisionary Metaphysics (Princeton University Press).

Scanlon, T.M. 2003 "The Significance of Choice," as reprinted in Free Will (second edition) edited by Gary Watson (Oxford University Press)

Searle, John R. 1983. Intentionality (Cambridge University Press).

Strawson, Peter. 2003. "Freedom and Resentment," as reprinted in Free Will (second edition) edited by Gary Watson (Oxford University Press)

Tuomela, Raimo. 1995. The Importance of Us (Stanford University Press)

Velleman, J. David. 2000. "How to Share an Intention," as reprinted in Velleman, The Possibility of Practical Reason (Oxford University Press)