

TCRS



Teoria e Critica della Regolazione Sociale

Quaderno 2011

Scott J. Shapiro

MASSIVELY SHARED AGENCY

Centro Studi TCRS

Via Crociferi, 81 - 95024 Catania - Tel. +39 095 230478 - tcrs@lex.unict.it

Scott J. Shapiro
Yale University
scott.shapiro@yale.edu

In:
Intenzionalità collettiva e figure della responsabilità
Quaderno 2011

ISSN: 1970-5476
Centro Studi
"Teoria e Critica della Regolazione sociale"
Via Crociferi, 81 - 95024 Catania
Tel. +39 095 230478 - Fax +39 095 230462
tcrs@lex.unict.it
www.lex.unict.it/tcrs

Scott J. Shapiro

MASSIVELY SHARED AGENCY

«[T]he cooperation of the wage-laborers is entirely brought about by the capital that employs them. Their unification into one single productive body, and the establishment of a connection between their individual functions, lies outside their competence. These things are not their own act, but the act of the capital that brings them together and maintains them in that situation. Hence the interconnection between their various labors confronts them, in the realm of ideas, as a plan drawn up by the capitalist, and, in practice, as his authority, as the powerful will of a being outside them, who subjects their activity to his purpose».¹

(Karl Marx)

In his classic work *The World We Have Lost*,² Peter Laslett describes the life experienced by the English of the Seventeenth Century, the period roughly between the Middle Ages and the Industrial Revolution, as one defined by its smallness and intimacy. According to Laslett, virtually every individual lived with a family, either their own or another. There were neither firms nor factories; businesses were located within the home and invariably disbanded with the death of the *paterfamilias*. Few of the 2,000 schools in England had more than one teacher. Gatherings among commoners were tiny by contemporary standards, the largest typically being the church congregation on Sundays. Laslett notes that the biggest crowd assembled in the Seventeenth Century, the Parliamentary Army at Marston Moor, “would have gone three, four or even five times into the sporting stadium of today.”³

The world we have gained, by contrast, is characterized by the enormous scale of social life. Business corporations, consumer cooperatives, trade unions, research universities, philanthropic organizations, professional associations, standing armies, political parties, organized religions, governments and legal systems, not to mention the collaborative ventures made possible by the internet,

¹ *Das Kapital*, 449-50 (trans. Ben Fowkes). Thanks to Matthew Smith for the quotation.

² London: Methuen, 1971(2nd edition).

³ *Ibid.* at 10.

such as Wikipedia, MMOG's (massively multi-member on-line games), open-access software and the World Wide Web itself, all harness the agency of multitudes in order to fulfill certain objectives. The modern world, we might say, is one defined by "massively shared agency."

And yet, philosophy has no viable theory for analyzing these ubiquitous activities. Although the theory of action has seen a recent turn from a more or less exclusive concern with individual agency to concerns with pervasive forms of shared activity, as when you and I sing a duet together or paint a house together, the accounts of shared agency produced are unable to account for the existence of massively shared agency.

The reason for this is two-fold. First, action theorists have largely eschewed giving analyses of activities involving authority structures. They have confined themselves to egalitarian activities such as the aforementioned duet-singing and house-painting, where neither participant has normative power over the other.⁴ Unfortunately, many forms of massively shared agency are unlikely to succeed in the absence of authority. Without some centralized control over behavior, the odds that many people will organize themselves towards the same objective and resolve their conflicts in a peaceful and efficient manner is apt to be low. To give an account of the kinds of massively shared activity that we normally encounter one must also provide an account of authority.

Secondly, philosophers of action have largely concentrated on analyzing shared activities among highly committed participants. The working assumption has been that those who sing duets or paint houses together are all committed to the success of the activity. It is unclear whether action theorists have intentionally limited the ambitions of their theory or have been operating under the notion that shared activity requires these forms of commitments. This restriction, however, has rendered these theories inapplicable to instances of massively shared agency. For it is highly unlikely that large-scale ventures such as industrial production or fighting a war under conscription can be staffed with individuals who are all

⁴ See, e.g., Michael Bratman, *Faces of Intention* (Cambridge: Cambridge UP, 1998), p. 94 and p. 110. In her most recent book, *A Theory of Political Obligation: Membership, Commitment, and the Bonds of Society* (Oxford: Oxford UP, 2006), Margaret Gilbert offers a theory of shared agency involving authority. It should be noted, however, that Gilbert is less interested in giving an analysis of the authority relation than in specifying the conditions under which political authority legitimately imposes obligations. Thus, she neither provides an analysis of institutional authority nor show how authority can arise in the face of alienation. One of the aims of this paper is to provide an analysis of institutional authority involving alienated participants (see, e.g., Section V below).

committed to the same goals, especially in pluralistic societies such as our own. Alienation and massively shared agency usually go hand in hand.

Since the modern world depends on massively shared agency at every turn, it must be possible for individuals to work together despite the existence of authority and alienation. And to the extent that models of shared agency developed by philosophers rule this out, these models must be supplemented or revised accordingly.

My strategy in this paper, therefore, will be as follows. In Section I, I will begin by setting out what I take to be the most interesting and plausible theory of shared agency that currently exists, namely, the one developed by Michael Bratman. In Sections II and III, I will point out two major limitations of Bratman's theory: first, that it applies only to ventures characterized by a rough equality of power and second, that it applies only to small-scale projects among similarly committed individuals. I will then attempt in Sections IV and V to develop the important insights of Bratman's theory into a new account of shared agency, one that will be applicable to small egalitarian ventures as well as large scale institutional practices involving authority structures.

I. Shared Intention and Shared Agency

I cooked dinner last night. My neighbor also cooked dinner last night. But we did not cook dinner together. Rather, my wife and I cooked dinner together. We might ask: what makes it the case that my wife and I cooked together, but my neighbor and I did not?

A plausible response is to say that my wife and I cooked dinner together because we intended to cook dinner together, whereas my neighbor and I did not. My neighbor and I *each* intended to cook, but we did not intend to cook with *each other*. Shared agency, it is natural to say, is distinguished from individual agency by virtue of the intentions of the agents. Even if my neighbor used my kitchen to cook and cooked it at the same time as my wife and I cooked, and even if we cooked the same food, our cooking was distinct from his cooking because we did not intend to cook with him and he did not intend to cook with us.

The pressing question, then, is to explain the nature of the intentions that lead to shared agency. What kind of an intention is an intention to act together? Because it would be beyond the scope of this paper to survey and evaluate every attempt to answer such a question, I will confine myself to the theory developed by Michael Bratman. Although the focus will mainly be on Bratman's model, I will have occasion to mention other accounts of shared agency and compare them in certain relevant respects to his account.

A. *The Functional Roles of Shared Intentions*

Bratman's strategy for distinguishing individual from shared intentions is broadly-speaking functionalist in nature. On his account, shared intentions are characterized by a set of dispositions and associated norms.⁵ A shared intention is any complex of states that plays, and ought to play, three roles in shared activity. First, a shared intention coordinates the actions of each participant towards the realization of their goal and has associated with it norms that require such coordination. Our shared intention to cook together, for example, will tend to lead one of us to cook part of the meal and the other to cook the other part. Moreover, the same shared intention exerts rational pressure on both of us to coordinate our actions so that we each prepare part of the dinner. Second, the shared intention coordinates (and ought to coordinate) the planning of each participant so that they can achieve the intended goal. Our shared intention to cook leads (and should lead) me, say, to plan to buy the vegetables if I know that you aren't planning to do so. Third, shared intentions specify a background for bargaining in the case of conflict. If we disagree, say, on which food to cook for dinner, our shared intention will tend to (and should) limit bargaining over the kind of food to cook and not what we should eat for, say, breakfast the next day.⁶

⁵ *Ibid.* at 112.

⁶ A great virtue of this functionalist strategy is that it takes a unified approach to intention and agency. As Bratman has forcefully argued in the case of individual agency, intentions are mental states that tend to organize behavior towards the achievement of some objective and are governed by norms of rationality that guide agents towards the same end. See his *Intention, Plans and Practical Reason* (Cambridge: Harvard UP, 1987), especially chs. 1 and 2. Individual intentions organize behavior by disposing agents not only to act on their intentions, but also to fill in their plans as to the appropriate means for the intended ends as well as ruling out as inadmissible plans that conflict with other intentions and beliefs of the agent. Associated with these dispositions are norms that require means-ends coherence, consistency with other intentions and beliefs and stability of intentions over time. Bratman's theory of shared agency requires that shared intentions play the same organizing role and be governed by the same norms as individual intentions. The difference between shared and individual intentions,

Having specified the functional roles that shared intentions are supposed to play, Bratman next considers which complex of mental states play such roles. Two basic options present themselves. A shared intention might either be an intention in the mind of the group. Or it might be constituted by the individual intentions of the group members.

Bratman rejects the first option, given the apparent metaphysical extravagance that comes with the postulation of group minds.⁷ He opts instead for a reductionist account which seeks to explain shared intentions as a complex of individual intentions. This raises the next question: what is the content of these constituent mental states?

Again, two options present themselves. The intentions of each group member may refer either to their own actions in the shared activity (call these "singular" intentions) or to the group's activity (call these "plural" intentions).⁸ According to the first option, our shared intention to cook dinner consists in my singular intention to do *my* part in cooking dinner and my wife's singular intention to do *her* part in cooking dinner.⁹ On the second alternative, our shared intention consists in my plural intention that *we* cook dinner and my wife's plural intention that *we* cook dinner.

Bratman opts for the second alternative: shared intentions are reducible, at least in part, to the plural intentions of the group members. Although Bratman does not provide an explicit argument for this choice, it is clear that he thinks that a complex of singular intentions cannot play the necessary functional roles. Suppose I intend to do my part in cooking dinner by making the side dish. My wife, however, intends to do her part by baking a cake for dessert. Since I am only committed to making the side dish and my wife is only committed to baking dessert, there is no rational pressure on either of us to cook the main part of the meal. This complex of singular intentions, therefore, will not lead us to cook

rather, concerns the different demands that organizing takes in these varying contexts. Organizing the behavior of a group requires different mental operations than organizing the behavior of an individual towards an individual objective. The three roles associated with shared intention – the coordination of action and intention between participants and provision of a basis for bargaining – reflect this change in context.

⁷ For a provocative attempt to defuse metaphysical objections to postulating group persons, see Carol Rovane, *Bounds of Agency* (Princeton: Princeton UP, 1998), especially ch. 4.

⁸ Bratman does not use the term "plural intention." I borrow it from J. David Velleman, "Review of *Faces of Intention*," *Philosophical Quarterly* LI (2001): 119-21, p. 119.

⁹ Tuomela and Miller opt for this approach. See Raimo Tuomela and Kaarlo Miller, "We Intentions." *Philosophical Studies* LIII (1988): 367-89.

dinner. By contrast, Bratman believes that plural intentions can play the requisite roles. It is to this demonstration that we turn.

B. *Interlocking Intentions and Meshing Subplans*

Bratman's first suggestion is that shared intentions are fully reducible to plural intentions. Thus, we share an intention to cook dinner if and only if I intend that we cook dinner and you intend that we cook dinner.¹⁰ The idea behind the proposal is that a simple combination of plural intentions is guaranteed to play the necessary functional roles. My plural intention that we cook dinner will set up our cooking as an objective to which I am rationally committed. Since satisfying this objective requires that I contribute to our cooking and you contribute to our cooking, I will be rationally committed to coordinating my planning and action with your planning and action. Absent this coordination, we won't end up cooking dinner and hence satisfying our mutual commitments. Furthermore, since my intention that we cook rationally commits me to ruling out options that are inconsistent with our cooking and likewise for you, we will be led in our bargaining to consider how we will cook together. Our plural intentions, in other words, provide a background from which rational bargaining can take place. It would seem, therefore, that a simple combination of plural intentions can play the three roles characteristic of shared intention.

Unfortunately, this argument does not quite go through. As Bratman points out, one way in which I can satisfy my intention that we cook dinner is for me to take your hands and chop the vegetables and then use your hands to stir them in

¹⁰ It might be wondered how this proposal escapes circularity: how can we explain shared agency via plural intentions when plural intentions (intending that *we J*) seem to use the concept of shared agency? Bratman's solution is to notice that the descriptions of the joint activity in the content of the plural intentions can be read in two different ways. The content-phrase "that we cook dinner" is ambiguous between a "cooperatively-loaded" and "cooperatively-neutral" sense. A cooperatively-loaded description is one that describes the activity in question as a shared activity. Thus when I intend that we cook dinner under a cooperatively-loaded description, I intend that we intentionally cook dinner together as a cooperative venture. A cooperatively-neutral description describes an activity in a manner that is compatible with its being a cooperative venture or one that is non-cooperative. Thus, if referring to my neighbor and I, I said that we cooked dinner yesterday, you would not know whether we intentionally cooked dinner together or whether I cooked my own dinner and he cooked his own dinner. Bratman proposes to understand the content of the constituent plural intentions in the cooperatively neutral sense. In this way, Bratman is able to explain shared intentions, and hence shared agency, by reference to mental states whose contents do not characterize the action in question as shared activities. Thus, our intention that we cook together, understood in a cooperatively-loaded sense, can be analyzed in terms of intentions that we cook, understood in a cooperatively-neutral sense. See generally, *Faces of Intention*, pp. 96-7.

the wok. In this instance, my intention that we cook dinner will lead us to cook dinner but not by virtue of your agency. My plural intention, in other words, will not lead us to intentionally cook dinner together.

The problem with the first proposal is that plural intentions do not ensure that agents who possess them will coordinate their *actions* with one another. My intention that we cook can be satisfied even if you don't intentionally cook. What is needed is a complex of intentions that ensures that each participant to the shared activity coordinates their actions with each other.

Bratman's next suggestion is that we share an intention to J not only when we each intend that we J but also when we intend that we J by way of the other's intention that we J. Thus, my wife and I intend that we cook just in case I intend that we cook because of her intention that we cook and she intends that we cook because of my intention that we cook. Bratman calls such intentions "interlocking intentions."

Notice that it is not possible to satisfy interlocking intentions by taking over the intentional agency of one of the other participants. If I take your hands and chop the vegetables I will not have acted in accordance with my intentions because my intention is that we cook dinner by way of your intention that we cook. A complex of interlocking intentions, thus, results in the coordination of the actions of the participants.

As Bratman points out, once we require plural intentions to interlock, it would seem odd if the only intentions to which these intentions refer are the intentions of the other participants. Why would I care about the successful execution of my partner's intentions and not my own? It seems more natural to say, for example, that I intend that we cook because of my wife's intention that we cook as well as my intention that we cook.

Requiring interlocking plural intentions, however, is still not strong enough. Suppose you think that we should deglaze the pan with water, but I think we should deglaze it with wine. While you are not looking, I pour out the water from the measuring cup and fill it with wine. You then unwittingly deglaze the pan with wine. While it is true that we cooked dinner by virtue of your actions, we did not do it in the way that you intended. Although you helped cook dinner, you did not intend to help by pouring *wine* into the pan. Although our interlocking intentions

led us to coordinate our actions, they did not lead us to coordinate our planning. Our activity came about in a manner that one of us did not plan.

In order to ensure that shared intentions lead to the coordination of intentions as well as actions, Bratman builds in a further commitment that the plural intentions be interpersonally consistent. That is, not only must each participant intend that they engage in the joint activity because of each other's plans and sub-plans but they must also intend to do so in accordance with subplans that "mesh." When I pour out your water and substitute wine, our cooking dinner is not in accordance with sub-plans that mesh because your sub-plan of deglazing with water is inconsistent with my sub-plan of deglazing with wine. We intend to cook together only if we are committed to resolving this conflict. Should conflict arise, the participants must be committed to deliberate, negotiate and/or bargain in order to overcome the impasse. This is not say that this commitment is an absolute one. It may be that each party cannot accept the subplans of the other, in which case the attempt at cooperation will ultimately fail.

Bratman imposes one more requirement on shared intentions. In order for us to share an intention, it must be common knowledge that we have these interlocking plural intentions. We wouldn't say that we intended to cook together if you didn't know that I intended that we cook together.

In sum, Bratman's proposal for shared intentions is as follows:

We share an intention to J if and only if

- (1)(a)(i) I intend that we J.
- (1)(a)(ii) I intend that we J in accordance with and because of meshing subplan of (1)(a)(i) and (1)(b)(i).
- (1)(b)(i) You intend that we J.
- (1)(b)(ii) You intend that we J in accordance with and because of meshing subplan of (1)(a)(i) and (1)(b)(i).
- (2) It is common knowledge between us that (1)¹¹

Having characterized shared intentions, Bratman proceeds to define shared intentional activity [or "SIA"]. In order for two people to act together intentionally, it is not enough for each to intend that they engage in the activity and to successfully perform the activity. The intention and action must be appropriately connected.

¹¹ See *ibid.* 100–102.

Bratman specifies two ways in which the attitudes in question must be related. First, there must be mutual responsiveness of intention. During the planning of the joint activity, each participant must be attuned to the subplans of the other participants.¹² Each must be committed to adjusting their intentions in light of the intentions of the other members of the group, or at least committed to convincing the others to adjust their intentions so as to achieve a mesh. Secondly, there must be mutual responsiveness in action.¹³ During the execution of the joint activity, each participant must be attuned to the behavior of the others and adjust their behavior accordingly so as to achieve the intended results. When we sing a duet together, I guide my behavior by listening carefully to which parts of the song you are currently singing. Bratman contrasts mutual responsiveness in action to “prepackaged cooperation.”¹⁴ In prepackaged cooperation, each participant acts according to, and only according to, the previously set plan. In a prepackaged duet, I might gauge my parts by looking at my watch, rather than my partner, and coming in at the pre-designated time. In such a case, the planning was a shared intentional activity, but the singing is not.

In conclusion, a shared intentional activity results whenever we J, we share the appropriate attitudes and these attitudes are connected to our J-ing.

Our J-ing is a SIA if and only if

- (A) We J.
- (B) We have the attitudes specified in (1)(a) and (1)(b) and (2) and
- (C) (B) leads to (A) by way of mutual responsiveness (in the pursuit of J-ing) in intention and action.¹⁵

II. Shared Intentional Activity with Authority

As Bratman himself notes, he developed his theory by focusing on joint activities that do not involve authority structures.¹⁶ In this section, I would like to

¹² *Ibid.* 106.

¹³ *Ibid.*

¹⁴ *Ibid.*

¹⁵ *Ibid.*

¹⁶ *Ibid.* 94. See also *Ibid.* 110.

explore whether his account of shared agency can be extended to cover these sorts of cases and, hence, is capable of providing an appropriate foundation for modeling massively shared agency.

A. Can Bratman's Model be Extended?

As a general matter, when adults cook dinner together, sing a duet together, or fly to another city together, neither has authority over the other. If while preparing dinner I order you to "fire" the chicken, our cooking together would probably end then and there. There are many cases, however, where one member of the group has authority over others and yet it is entirely appropriate to consider their joint activity a cooperative venture.

Imagine that a group of friends decide to sail from New York to Nova Scotia. Having sailed together once before, they remember the problems that arose the last time they set out. They recall that during storms, despite their best efforts, they were unable to coordinate their behavior effectively and, as a result, everyone's safety was seriously put at risk. In addition, the friends were generally indifferent between stopping in Maine and proceeding directly to Nova Scotia, and this collective indifference made the decision difficult to make, causing substantial delays. In order to prevent the same problems from arising again, they resolve to appoint a group member as the captain for the trip. They figure that the trip will be faster, safer and more pleasant if one is designated as having authority over the others.

Let us assume that their first trip was a shared intentional activity. I think it is reasonable to assume that the mere appointment of an authority would not preclude their second trip from being one as well. That the new captain has authority over her crew is not inconsistent with their cooperating with one another in order to achieve a common goal. Indeed, they set up the authority structure precisely because they thought that they could not achieve their goal, or achieve it as well, without hierarchy.

At first glance, however, it would seem that this second trip could not be a shared intentional activity. Consider the conditions that Bratman imposes on shared intentions. As we saw, participants share an intention only when each intends to act in accordance with subplans that mesh. However, this condition

appears to fail in joint activities involving authority. The captain of the boat does not intend to mesh her subplans with those of her crew – rather, they are supposed to mesh their plans with her. Similarly, mutual responsiveness does not seem to be present in the second voyage. While the crew must be responsive to the captain, the captain does not have to be responsive to the crew. Such is the nature of hierarchy: responsiveness to intention and action need not be mutual. Those below must respond to those above, but not vice versa.

These incompatibilities, though, are only minor or non-existent. As I will argue, the conditions that Bratman imposes on shared intentional activity are either consistent with the existence of shared intentional activities involving authority [hereinafter “SIAA”], or must be slightly modified in order to achieve the requisite fit. In addition to showing the consistency of Bratman’s general model, I will set out further conditions that will distinguish an ordinary SIA from one where authority relations are also present.

B. *Orders and Meshing Subplans*

Let us begin with the meshing subplans condition. Upon reflection, it is apparent that the captain of the boat is committed to acting in accordance with meshing subplans. Given that she wants everyone’s behavior, including her own, to be organized in a certain way, she must ensure that everyone’s subplans, including her own, mesh. If the captain believes it preferable to rig the sail herself, she must find out whether anyone else intends to rig the sail as well. More importantly, if the captain has decided that the boat will dock in Maine, she must guarantee that this decision meshes with everyone else’s plans. The captain courts disaster if she plans for the crew to dock the boat in Maine but they plan to dock it in Nova Scotia.

That the captain is committed to meshing subplans does not mean, however, that she is committed to revising *her* plans in case of conflict. To be sure, the captain may adjust her plans to mesh with those of her crew; she may also negotiate and bargain with them in order to persuade them to revise their subplans, rather than vice versa. What is special about the captain’s position, though, is that she possesses a tool for achieving a mesh that her crew does not have: she may *order* her crew to act as she intends. If the captain exercises her

authority, her crew will be committed to revising their subplans so that they mesh with the captain's. Once orders are given, it will then be inappropriate for the crew to negotiate and bargain with the captain – to do so would be a slight to the captain's authority.

In a SIAA, therefore, every participant is committed to acting in accordance with meshing subplans. Authorities and subjects are distinguished, however, by what those commitments require them to do given their differing roles. Whereas authorities can achieve interpersonal consistency either by revising their subplans or issuing orders, their subjects do not have that luxury. If they are ordered to act by someone having authority over them, they must revise their subplans just in case their intentions do not yet mesh.

As this discussion suggests, orders are devices that an authority uses in order to enforce its intentions. When someone in authority issues an order, that person intends that the subject (a) adopt the content of the order as her sub-plan and (b) revise her other subplans so that they mesh with the order. If the captain orders the crew to head towards Maine, the captain intends that every crew member adopt her plan as their plan, as well as revise their other plans so that they are consistent with this newly adopted plan.

SIAA's can now be characterized as follows:

Our J-ing is a SIAA if and only if

- (1) J is a SIA.
- (2) If one of us has J-authority over the other, then
 - (a) The authority intends that the subject adopt the content of her orders as subplans as well as revise the subject's subplans so that they mesh with the orders.
 - (b) The subject intends to adopt the content of the authority's orders as subplans as well as to revise his subplans so that they mesh with the orders.
 - (c) (a) and (b) are common knowledge
- (3) Either I have J-authority over you or you have J-authority over me.

One limitation of this formulation, however, is that it takes the J-authority relation as given (why the relation is called "J-authority," rather than "authority," will be explained shortly). It does not show us how to determine whether someone has J-authority over another, only that certain inferences can be generated from

such a determination. In principle, though, it should be possible to construct a definitive test for establishing whether someone has authority in a shared intentional activity. After all, it is plausible to suppose that authority relations in a SIAA are created in part by that very activity. The captain became the captain of the boat because the friends decided, during the planning phase of their trip, that one will have the power to regulate certain activities on the boat. The test for determining who has J-authority over whom should track the process by which authority relations are generated.

How, then, are authority relations generated from SIA? Here's a hypothesis: Someone comes to have J-authority over another just in case he intends that others adopt and revise certain subplans and these others intend to respond to these directives in the manner intended. We will say that authority arises in shared intentional activities just in case the intentions of the participants "vertically" interlock. This idea can be captured by converting condition (2) above into an explicative definition of "J-authority."

(V) A has J-authority over S if and only if:

- (a) A intends that the other adopt the content of her directives as subplans as well as revise S's subplans so that they mesh with the directives.
- (b) S intends to adopt the content of A's directives as subplans as well as to revise his subplans so that they mesh with the directives.
- (c) (a) and (b) are common knowledge

The first question that naturally arises with respect to the relation of J-authority as set out in (V) is whether vertical interlocking creates reasonable authority relations. Do the participants of J now have objective reasons to listen to their new "superiors"? My own view is that mere vertical interlocking of intentions is insufficient to create reasonable authority. This is so because many shared activities are morally noxious and there can be no obligation to participate in morally noxious enterprises. Bank robberies, terrorist plots and unfair business practices may all involve structures of J-authority but there is no reason to act so

as to further their cause. For this reason, I have termed the relation "J-authority" rather than "authority" *simpliciter*.¹⁷

While the vertical interlocking of intentions is insufficient for reasonable authority, it is enough to create *rational* authority, that is, to render it rationally required for subjects to heed the directives of those who have J-authority. The reason is that, according to (V), submission to authority involves *intending* to take the content of another's directives as one's subplans. As a result, the norms of rationality that govern intentions come into play.¹⁸ If someone submits to the authority of another, and yet ignores an order directed to him, then he will be acting in a manner inconsistent with his intentions. His disobedience will be in direct conflict with his intention to defer. Moreover, this "subject" is not rationally entitled to give up his intention to defer to authority, unless he has a sufficiently good reason to reconsider. As Bratman has argued, intentions are subject to norms of reasonable stability. Since the function of intentions is to guide our conduct over time so that we need not deliberate about what to do next at every turn, it would be self-defeating if agents were rationally permitted to reconsider their intentions absent some good reason to do so.¹⁹

While submission to authority creates rational authority relations, such that disobedience becomes irrational upon acquiescence, agents might nonetheless be irrational for submitting to authority in the first place. The authority may be a nitwit. Nevertheless, once someone has formed an intention to treat the authority's directives as trumps to their own planning, they have changed their normative situation and are rationally committed to follow through unless good reasons appear that force them to reconsider.

Whether submission to authority is rational depends on whether it is rational to delegate one's own planning authority to another. In general, such delegation is irrational because we are normally the best and cheapest judges of what we should do. But when this is no longer the case, when others know more than we do about what we should do and can be trusted to point us in the right direction, or when we

¹⁷ For stylistic reasons, I will sometimes use "authority" in what follows instead of the more awkward "J-authority."

¹⁸ See *Intention, Plans and Practical Reason*, chs. 1 and 2.

¹⁹ In order for vertically-interlocking intentions to generate the J-authority relation, it is necessary that the intentions be fairly stable. As Margaret Gilbert pointed out to me, individuals whose demands can be ignored virtually at will are best not considered authorities of any sort.

can conserve precious cognitive resources by deferring to others without risking too much error, we should plan for others to plan for us.

In shared activities, there are additional reasons to delegate planning authority. Insofar as shared activities usually require a certain degree of organization in order to succeed, and because it is usually costly for participants to organize themselves, it will often be rational for participants to rely on an authority to guide their conduct. In submitting themselves to an authority, they obviate the need to deliberate and bargain with one another in order to establish a mesh in their subplans.

Authorities are useful in shared activities, we might say, because they are “mesh-creating” mechanisms. When disputes between participants break out with respect to the proper way to proceed, authorities can create a mesh between the subplans of the participants by demanding that both sides accept a certain solution. By accepting this authoritative settlement, participants satisfy their commitment to act in accordance with meshing subplans. And in situations where superiors and subordinates are in conflict, superiors can satisfy this commitment precisely by issuing an order, thereby imposing meshing subplans on their underlings.

C. Mutual Responsiveness

It should be clear by now that authorities are as responsive to the intentions and actions of their subjects as their subjects are responsive to their intentions and actions. The function of authorities in SIAA is to ensure that the participants’ actions are organized so that goals of the activity can be achieved. This will require that authorities be attuned to the intentions and actions of the participants, modifying their own subplans in some cases and ordering the revision of their subjects’ subplans in others. Moreover, they must be sensitive to the success or failure of their subjects’ actions so that compensatory action may be taken if judged appropriate.

There is an important sense, however, in which the participants in SIAA are not mutually responsive to one another. When an authority exists within a group, subjects may look towards the authority in order to coordinate their behavior rather than taking their cues from each other. Moreover, when orders are issued, the directives may remove discretion from the subjects on how to proceed. Depending

on the scope of authoritative regulation, aspects of a SIAA may start to resemble prepackaged cooperation, insofar as the subjects will be forced to respond to the rules laid down rather than to the each other's intentions and actions.

D. Summing Up

As we have seen, Bratman's account of shared agency is compatible with authority relations, provided that some slight modifications are made. In particular, the requirement of mutual responsiveness in action must be dropped. The major conceptual generalization involves broadening the class of mesh-creating mechanisms to include authority. In shared activities involving authority, superiors can satisfy their commitment to meshing subplans not only by deliberating, negotiating and bargaining, but by issuing orders and thus imposing a mesh on their subordinates.

III. Massively Shared Agency

Having argued that, with a bit of tweaking, Bratman's model of shared agency is able to account for joint activities involving authority, I would now like to examine whether they can be extended to large-scale ventures. I will argue that they cannot. The problems center on Bratman's demand that all participants in shared activities must share a plural intention in favor of the activity. This requirement, I will argue, excludes activities that employ large numbers of participants, given the doubtfulness that participants in these activities will, or can, all share the necessary commitments.

A. Alienated Participants

It is one of the merits of Bratman's theory of shared agency that it does not require each participant to have the same motivations for engaging in the shared activity. You may want to paint the house because the paint is chipping; I may want to paint the house because I am sick of its color. Nevertheless, we can paint the house together if our respective reasons support the same intention, namely, the intention that we paint the house together.

Despite this liberality, I would like to argue that Bratman's account is still too restrictive. For it is often the case that participants engage in a shared activity even though some, perhaps all, are not committed to the joint activity. People can work together, in other words, despite the absence of a shared intention that the group J.

To see why the requirement of shared plural intention is too strong, consider the case of the "alienated" painters. Suppose Abel wants Baker and Charlie to paint his house. Abel offers \$1000 dollars to Baker if Baker does what he tells him to do. Abel offers Charlie the same terms. Baker and Charlie both agree. Abel then tells Baker to scrapes off all the old paint and Charlie to paint a new coat on the scraped surface. Charlie waits until Baker scraps the old paint from the front of the house and then proceeds to paint a fresh coat on it. While Charlie paints the front, Baker scrapes the paint off the back of the house. When Baker finishes scraping, Charlie paints the rest of the house. It would seem that both Baker and Charlie have intentionally painted the house together.

It follows on Bratman's theory that Baker and Charlie share a plural intention that they paint the house. But this does need not be so. Suppose that halfway through his painting the fresh coat, Charlie announces that he quits. Baker replies that it doesn't matter to him – after all, he will get his money regardless. If we assume that Baker is rational, sincere and hasn't changed his mind once he accepted the job, we can infer from this exchange that Baker never formed a plural intention. This is so because to intend that they paint the house entails a rational commitment on the part of Baker to the joint activity of their painting the house. This commitment, if present in a rational participant, must express itself in some form of action designed to result in their painting of the house. Baker might pick up the brush and paint the house himself, or notify Abel of Charlie's departure, or try to convince Charlie that he shouldn't leave. The fact that Baker does nothing indicates that he is not so committed. He has the singular intention to do as Abel says and, hence, to scrape the paint off the house, but not the plural intention that *they* paint the house. Yet, despite the fact that Baker is not so committed to the joint activity, it will be true that they intentionally painted the house together if Charlie changes his mind and stays to finish the job.

Bratman's model of shared agency, therefore, appears to exclude cases of shared agency among alienated participants. Baker and Charlie are both alienated

from the project of painting the house. They don't care a wit about painting the house, only in getting their money. Indeed, they may hate Abel and not want him to have a nicely painted house. Yet, they can still paint the house together, and do so intentionally, even though neither of them intends that they paint the house together. They paint together, in other words, despite the fact that neither of them possesses a plural intention, let alone share that plural intention.

I do not wish to claim that Bratman's account cannot apply to any activities involving alienated participants. For it is often the case that participants can be committed to a joint activity to which they are apathetic or even hostile. Suppose that Baker and Charlie are told that the only way they will get their money is if they manage to get the house completely painted. In that case, Baker and Charlie, if rational, will be committed to their painting the house and hence will share an intention in favor of painting the house together. Yet, there are situations in which alienation does not result in shared intentions, such as the first scenario mentioned, and it is these cases that Bratman's model unjustifiably rules out.

If Bratman's model is inapplicable to the case of the alienated painters, then it will be unsuitable for virtually *any* instance of massively shared agency. For in any large scale activity, there are bound to be participants that intentionally contribute to the group effort but are not committed to the success of the group venture. Dilbert might be hired by Microsoft to work on a new version of their operating system. Because he gets paid only if he programs, he intends to program. However, he may not care at all if the software group is successful, given that he is paid regardless of whether the group is successful. He does not form a plural intention in favor of the group venture because he does not share its goals. If the group is successful, and is successful in part because of Dilbert's efforts, then the group has intentionally worked together, regardless of Dilbert's lack of commitment to their success.

B. *Failure of Interdependence*

According to Bratman's accounts of shared agency, it is not enough that participants have as their goal that the group engages in the joint activity. The participants must each *intend* that the group engage in the activity. Shared agency requires shared plural intentions, not just shared plural goals.

David Velleman has objected that Bratman's requirement of a shared plural intention is too strong.²⁰ Intention, he pointed out, is an attitude that settles matters in favor of the intended action. This settling function of intention makes it difficult to see how a plural intention can be shared. How, Velleman wondered, can my intention that we paint the house settle the matter for us – isn't the question of whether we will paint the house up to you as well? To be sure, if I have authority over you, I can settle the matter for you. But if so, Velleman continued, how can you intend that we paint the house, given that this would require that you have authority over me as well?

How can I frame the intention that "we" are going to act, if I simultaneously regard the matter as being partly up to you? And how can I continue to regard the matter as partly up to you, if I have already decided that we really are going to act? The model seems to require the exercise of more discretion than there is to go around.²¹

To this objection, Bratman responded by noting that my settling the matter whether we paint the house does not entail that I have authority or control over your actions.²² For I might confidently predict that my intending that we paint the house will lead you to similarly intend. I settled the matter for us by intending, but I did so via your free agency, namely, by my knowing that you will take my intention as a decisive reason to adopt the same intention. Likewise, your intention that we paint the house settles the matter for us because you are confident that whether we paint the house together is completely dependent on whether you so intend. Given my intention, you know that we will paint the house if you intend that we paint the house and we won't paint the house if you don't intend that we paint the house. Thus, both of us can intend that we paint the house because both of us can settle the matter whether we will paint the house.

Bratman's strategy here is to show that participants can share a plural intention whenever the plural intentions are *interdependent*.²³ I can settle that we paint the house because I know that your intention depends on my intention. Likewise, you can settle the same matter because you know that my intention

²⁰ See J. David Velleman, "How to Share an Intention," *Philosophy and Phenomenological Research* LVII (1997): 29-50.

²¹ *Ibid.* at 35.

²² See Bratman, "I Intend that We J" in *Faces of Intention*, p. 155.

²³ *Ibid.* at 153.

depends on your intention. There is no incoherence in each of us settling what the other will do because each of us knows that the persistence of their intention is the ground for the persistence of the other one's intention.

Now it is plausible to suppose that intentions will be interdependent in many activities involving a small number of participants. This is so because in such instances each participant is likely to be pivotal to the success of the activity and hence each participant will ground their intentions in the persistence of the others' intentions. But it is doubtful whether this interdependence can endure in the face of expanding numbers. For in larger-scale shared activities, the marginal contribution of most contributors falls and, correspondingly, the likelihood that those participants will be deemed pivotal also decreases. Assume that Abel, Baker and Charlie decide to throw a party together. Each is happy to throw the party as long as at least one other person contributes. From Abel's perspective, it is not necessary that both Baker and Charlie help – either will do. Baker and Charlie have similar preferences. Suppose further that all of this is common knowledge. It would seem that neither Abel, nor Baker, nor Charlie can coherently intend that they throw the party together because each knows that they are not pivotal to the project. Abel knows that the persistence of Baker and Charlie's intentions that they throw the party does not depend on his intention. As long as Baker and Charlie intend, Abel is not pivotal. It would appear, therefore, that interdependence fails.

Bratman might respond by arguing that the intentions of Baker and Charlie do in fact depend on Abel's intentions. For if they learn that Abel is backing out, Baker and Charlie cannot have the intention that they (meaning Abel, Baker and Charlie) throw a party together. They will have to share a new plural intention, namely, an intention that Baker and Charlie throw a party together. Moreover, the shared activity will be transformed, because it will no longer involve Abel's agency. Abel can be confident, therefore, that he has settled the matter for Baker and Charlie because the persistence of Baker and Charlie's intentions do indeed depend on the persistence of his intention.

While this response might solve the problem of marginal participants in comparatively small-scale ventures, I doubt that it will be adequate in cases of massively shared agency. When multitudes work together on a project, it is unlikely that any participant will know the identity of all the other participants. Most of the employees at Microsoft, for example, are unknown to each other.

Because these workers don't know the identity of many of the other workers, the intentions of the latter cannot affect the intentions of the former. If Abel does not know that Baker works at Microsoft, Baker's plural intention cannot affect Abel's plural intention. Baker cannot, in other words, intend that the group of employees work together because he cannot settle the matter for them.

In cases of massively shared agency, therefore, interdependence, even of the weak kind we are now considering, must fail. The participants in such ventures cannot coherently possess, let alone share, the same plural intentions because each cannot settle the matter for the rest.

C. Functionalism about Shared Agency

In the last two subsections I argued that Bratman's shared plural intention requirement fails in cases of massively shared agency. In any large-scale venture, it is likely that some participants will be alienated, marginal or virtually invisible. When this is so, not all participants in the group activity will, or can, intend that the group engages in the joint activity. Since Bratman's account assumes that all instances of shared agency are explainable by shared plural intentions, it follows that his account rules out the possibility of massively shared agency.

Despite the fact that Bratman's shared plural intention requirement is misconceived, I think that it is nonetheless embedded within a highly appealing picture of shared agency. Indeed, reflecting on the compelling features of Bratman's model will show why the shared plural intention requirement is needlessly restrictive and ought to be rejected for reasons quite apart from its incompatibility with instances of massively shared agency.

Let's start by recalling that Bratman's model is founded on two basic premises. The first assumption, accepted by all theorists, is that an event is a shared action only if it is explainable by a shared intention. Call this the "Shared Agency-Shared Intention" Principle.²⁴

Bratman's second premise is that a shared intention is any complex of intentions that play three roles in shared activity. It must coordinate the actions of

²⁴ Notice that Bratman accepts a stronger principle, namely, than that an event is a shared action of J-ing only if it is explainable by a shared intention to J. This principle might be termed the "Simple View for Shared Agency." Bratman famously rejects the Simple View in the case of individual agency. See, e.g., his "Two Faces of Intention," *The Philosophical Review*, XCIII (1984): 375-405.

each participant towards the realization of the goal in question; it must coordinate the planning of each participant so as to achieve the intended goal; finally, it should specify a background for bargaining in the case of conflict. Call this premise "Functionalism about Shared Intention."

We have seen, however, that both premises are problematic. As for the Shared Agency-Shared Intention Principle, the alienated house painters case demonstrates that shared agency is possible without any shared intention. Baker and Charlie intentionally paint the house together even though they don't share an intention to paint the house together or do anything else together.

Two problems arise in connection with Functionalism about Shared Intention. First, in light of our earlier discussion of authority, it seems too strong to require that shared intentions provide a basis for bargaining. In many situations, bargaining is very costly. A complex of attitudes that precludes bargaining, but does so in favor of a cheaper conflict resolution mechanism, such as authority, should also count as a shared intention.

Second, and more importantly, it seems wrong to say that any complex of mental states that satisfies the three Bratmanian roles is a shared intention. Consider again the case of the alienated painters. Baker and Charlie do not share an intention to paint the house together. Yet, the fact that each separately accepted Abel's offer enables them to coordinate their actions, planning and resolve their conflicts. Because Abel has planned for them to paint the house, their actions will be coordinated towards the goal of painting the house. Moreover, their planning will be coordinated because Abel has given them separate tasks to fulfill and has thus ensured that the subplans that Baker adopts will not interfere with those that Charlie adopts. Finally, if their planning clashes, Abel is there to resolve the conflict. It appears, then, that a complex of attitudes other than a shared intention – in this case Baker and Charlie's singular intentions to listen to Abel and Abel's plural intention that Baker and Charlie paint the house – plays the three Bratmanian roles as well.

While the two premises on which Bratman's model rests prove to be mistaken, I believe that they nevertheless are on the right track. The first premise is correct to state that shared agency is action explainable by a complex of mental states of a distinctive kind. The second premise is correct to the extent that it claims that this distinctive kind is functional in nature, where the roles in question

involve the coordination of action, planning and the resolution of conflict. Putting the two ideas together, we get an attractive picture that might be called "Functionalism about Shared Agency," namely, that shared agency is action explainable by any complex of mental states that coordinate action and planning and resolve conflict.

It is important to see that Functionalism about Shared Agency does not entail a shared plural intention requirement. For while Bratman is right that a web of public and interlocking plural intentions satisfies the three roles he identifies, it does not follow that *only* this web of intentions can satisfy these roles. In other words, the roles that are characteristic of shared agency may be multiply realizable and hence be instantiated by a complex of mental states not construable as a shared intention.

Not only is the shared plural intention requirement not entailed by Functionalism about Shared Agency, it is actually inconsistent with it. For as we saw with the alienated painters, the three Bratmanian roles are realized by a complex of attitudes not wholly constituted by plural intentions (in fact, the participants in question only possess singular intentions). And it is plausible to suppose that many other combinations of propositional attitudes will be capable of fulfilling such functions as well.²⁵ Bratman's shared plural intention requirement, therefore, is unwarranted because it rules out the possibility that other complexes of attitudes might too serve these roles and be responsible for shared activity. A theory that took Functionalism about Shared Agency seriously, or at least wished to be consistent with it, should allow for the multiple realization of the roles characteristic of shared intention and hence credit events explainable by any such realizers as instances of shared agency.

D. *Hypercommitment*

Before I go on to suggest a more flexible theory of shared agency, I should note that Bratman is not alone in imposing overly restrictive conditions on the intentionality of participants in shared activities. Margaret Gilbert, for example, has

²⁵ Consider, for example, the following attitudes held by Abel, Baker and Charlie: each intends to contribute his fair share to their throwing a party and to do so in accordance with and because of the meshing subplan of each other's intentions to do so. Although no one holds a plural intention, each is able to coordinate their actions and intentions, and resolve their conflicts, for each intends to contribute their fair share to their throwing a party and to do so via each other's meshing subplans.

argued that participants in shared activities must be jointly committed to act together as a body.²⁶ You and I can cook together, according to Gilbert, only if we are jointly committed to doing so. Although there are many differences between Bratman's shared intentions and Gilbert's joint commitments, both accounts are alike in requiring too great a commitment on behalf of the participants in a shared activity.²⁷ Baker and Charlie are painting together even though neither is committed to painting together "as a body."

More than any other writer, Christopher Kutz has recognized that theories of shared agency must make allowances for the phenomenon of alienation.²⁸ Accordingly, Kutz drops the demand for shared plural intentions, requiring only that each participant possess what he calls a "participatory intention."²⁹ A participatory intention is an intention to *contribute* to the joint activity. The problem with Kutz's account, however, is that, like Bratman's and Gilbert's, it too is hyper-committed.³⁰ The case of the alienated painters shows that shared intentional activity can take place even without participatory intentions. Neither Baker nor Charlie intends to contribute to the painting of the house and yet they intentionally paint the house together.³¹

IV. Shared Activity and Shared Plans

As we have seen, shared agency is possible without shared plural intentions, and even without any shared intention. What, then, characterizes shared agency as *shared* agency?

²⁶ See, e.g., Margaret Gilbert, *On Social Facts* (Princeton: Princeton UP, 1986).

²⁷ On the some of the differences between Bratman's and Gilbert's accounts, see Michael Bratman, "Shared Intention and Mutual Obligation" in *Faces of Intention* and Margaret Gilbert, "What Is It for Us to Intend?" in *Sociality and Responsibility* (Lanham: Rowman and Littlefield, 2000).

²⁸ See Christopher Kutz, *Complicity: Ethics and Law for a Collective Age* (Cambridge: Cambridge UP, 2000), pp. 96-103. My thoughts on these matters have been influenced by Kutz's excellent discussion.

²⁹ See *ibid* at 81-82.

³⁰ I owe the term "hypercommitted" to Shelley Kagan.

³¹ It is true, of course, that even though they don't intend to contribute, they *intentionally* contribute to the painting of the house, insofar as they know that following the shared plan will lead to their contributing to the painting of the house.

A. *Sharing Plans*

My suggestion is that shared intentional activity is activity guided by a shared *plan*.³² A group G engages in a shared intentional activity to J, on this account, when five conditions are met:

- (1) There is a shared plan for G to J.
- (2) Each member of G intentionally follows her part of the shared plan;
- (3) Members of G resolve their conflicts about J-ing in a peaceful and open manner;
- (4) It is common knowledge that (1), (2) and (3);
- (5) J takes place in virtue of (1) and (2).

For example, both Baker and Charlie share a plan, namely, a plan to paint the house by having Baker scrape and Charlie paint. Each intentionally follows the parts of the shared plan that apply to them: Baker intentionally scrapes and Charlie intentionally paints. Since Baker and Charlie listen to Abel in case of conflict, they resolve their disagreements peacefully and openly.³³ Moreover, all this is common knowledge. Hence, when the house gets painted it will be true that Baker and Charlie intentionally painted the house together.

The primary failing of this account, of course, is that it doesn't tell us when a plan is shared. For the mere fact that members of a group follow the same plan does not mean that they share the plan. Men who take off their hat in church normally don't share any rule about respect in church and hence are not collectively showing respect in church (compare this with a chorus line where the dancers are taking off their hats together). What, then, makes a common plan a *shared* plan?

³² By a "plan," I don't mean the mental state of planning or "having a plan" but rather the object of that mental state, namely, an abstract entity that specifies actions to be taken depending on contingencies.

³³ Because Baker and Charlie resolve conflicts between their subplans by looking to Abel as an authority, they manage to paint the house in accordance with meshing subplans.

Return to the alienated painters. There, the reason that Abel's plan is shared by Baker and Charlie is clear: Abel intended for Baker and Charlie to paint his house and designed the plan so that they would be able to execute that intention. Abel assigned the task of scraping to Baker so that the old paint would be removed from the house. He assigned the task of applying a fresh coat to Charlie so that the house would be painted with a fresh coat. Finally, he instructed Baker to perform his part before Charlie's part so that only the old coat of paint would be scraped off and that the new paint would be applied to the bare wood. The plan was fashioned so that the parts added up to a completed activity, namely, to the painting of the house. This suggests the following preliminary characterization of a shared plan: a plan is shared by a group to J when (1) the plan was designed, at least in part, for the members of the group so that they may engage in the joint activity³⁴ J and (2) each member accepts the plan.

The requirement that shared plans be designed, at least in part, for members of the group to follow does not entail that *one* person designed the plan for the participants. The participants, for example, could have decided together how to divvy up the tasks. Indeed, each participant could have design their own part of the joint activity. Suppose Abel hired Baker and Charlie to paint his house but did not tell them how to go about it. Baker picks up the scraper and starts to scrape. Noticing this, Charlie says "Okay, then I'll paint." Baker replies, "That's great, because I can't stand the fumes." Despite the fact that these decisions were initially made unilaterally, they still share a plan to paint together, one that assigns scraping to Baker and painting to Charlie. They share a plan because the plan was developed for each member of the group so that they may successfully engage in the joint activity and each member of the group accepted that plan.

By requiring that each member of group accept the shared plan, I mean more than just that each member accept their part of the plan.³⁵ To accept a plan

³⁴ The phrase "joint activity" might be thought ambiguous as between a mere collection of actions and an integrated whole that has actions as its parts. My cooking dinner and your painting is a joint activity in the first sense because it is a collection of two individual actions, but not in the second sense because these actions bear no relation to one another. I am using "joint activity" in the second sense. Thus, our joint activity of painting together is an integrated whole because we don't want merely to, say, scrape old paint and apply fresh paint, but to paint a house that has scraping old paint and applying fresh paint as its parts.

³⁵ Someone accepts their part of a plan either when they (1) intend to do certain actions, knowing that these actions are required by their part of the plan or (2) intend to do what their part of the plan requires *under that description*. In the latter case, the participant may not know what their part requires: it is enough for acceptance of a part that someone intends to follow their part whatever it happens to be.

entails a commitment to let the other members do their parts as well. Thus, Charlie accepts the shared plan because, in addition to being committed to scrape, he is committed to allowing Baker to paint. The acceptance of a plan does not require that the participants actually know the full content of the shared plan; the commitment may simply be to allow others to do their parts *whatever they happen to be*. This suggests that a shared plan must be at least “publicly accessible,” namely, that the participants could discover the content of the plan that pertain to them and to others with whom they are likely to interact if they wished to find out.³⁶

In many situations, the full-fledged acceptance needed for sharing will be lacking at a certain level of detail. Charlie may allow Baker to paint, but may not have any commitment to let him use the ladder for the next hour. In this case, the shared plan will only specify who paints and who scrapes, but not who uses the ladder in the next hour. In fact, acceptance may be missing except for the commitment to engage in the joint activity, in which case the shared plan will be very rudimentary. “We paint the house” can be the sum total of the shared plan. Wherever the shared plan is silent, participants may be required to design their own subplans that will enable them to execute the plan and, unless these subplans come to be accepted by others, these parts will not be shared and may be contested some time in the future.³⁷

This suggests that Bratman’s model of shared intentional activity is actually a special case, namely, one that applies to ventures structured by highly rudimentary shared plans and where participants are consequently forced to devise for themselves which parts they will play. For whenever the participants are the principal designers of the activity, they must have plural intentions. The reason is simple: if the participants did not all have plural intentions, there would be nothing motivating them to finish planning the activity. Since Baker does not intend that they paint the house together, there is no rational pressure on him to figure out how to finish the project when Charlie threatens to leave. The only sorts of

³⁶ Requiring common knowledge of the plan’s content would be too onerous, for as the size of any shared activity grows, it becomes increasingly unlikely that most could know the entire content of the shared plan.

³⁷ Certain members of the group may devise and accept subplans that specify the parts that each are to play in certain joint subactivities, in which case these subplans will be shared by that subgroup, but not the group at large.

attitudes that will rationally propel such participants to completion of the project are plural intentions.

But if someone can design a plan for others to implement, then those who implement the plan need not have plural intentions. As long as Abel intends that they paint the house together, Baker and Charlie need only be committed to following what Abel tells them to do. These attitudes will play the coordinating and resolving roles characteristic of shared plural intentions. Hence, in such cases, shared plural intentions are superfluous.

These considerations suggest that Bratman's shared plural intention requirement should be substantially weakened and replaced with the following principle:

Commitment to a Shared Plan: The participants each have an appropriate commitment (though perhaps for different reasons) to a plan developed, at least in part, for them so that they may engage in the joint activity, and their engagement in the activity is in the pursuit of this commitment.

In contrast to Bratman's principle, the Commitment to a Shared Plan does not require that the participants intend that their group engage in the joint activity. Nor does it even demand that they intend to contribute to the joint activity. It merely requires that each be committed to acting on the shared plan and engage in the joint activity because of that commitment.

In addition, Bratman's requirement of Mutual Responsiveness in Intention must be dropped completely. This principle is appropriate only for situations in which the participants are the main designers of the activity. However, when shared activities are structured by "pre-packaged" plans, participants may be highly unresponsive to each other's intentions. They need only follow their parts of the shared plan.

It should be noted, however, that once we drop the requirement of mutual responsiveness we eliminate the possibility of representing all instances of shared agency as cases involving a group "working together." Consider a bookkeeper employed by AT&T in Omaha and a telephone line installer who also works for the same company but lives in Los Angeles. It would be odd to say that they "work together" at AT&T, given that they have never interacted with one another and know nothing of each other's existence. Or consider the masons who first laid the

foundation bricks for Saint John the Divine at the end of the 19th Century. Certainly we don't want to say that these masons worked together with those who are currently laying bricks for the cathedral.

Even though every member of a very large group can't, strictly speaking, work together, they still can be said to engage in shared agency, because each can intentionally play their part in a joint activity. If a verb is needed to describe these sorts of cases, we might choose "participate." The first and last masons each participated in the collective project of building a cathedral, even though they didn't participate at the same time and hence never worked together.

B. *Two Notions of "Planning for"*

Because the notion of "planning for" is crucial to the account of shared agency I have been developing, I should say a bit more about it and the related notion of "planning." Let us begin with "planning." "Planning" might either denote an activity or a mental state. Planning as an activity is the process of formulating a plan for subsequent adoption. It is the activity, in other words, of planning *how* to achieve a certain objective. Planning as a mental state is the state one is in after having adopted a plan. It is the state of planning *to* act in a certain way or planning *that* a certain state of affairs obtain. In this latter sense, planning is synonymous with intending. If I am planning to go to Mexico, I am intending to go to Mexico.

Because "planning" is ambiguous as between the activity and the state, "planning for" can refer either to the formulating of a plan for someone to follow or to the intention that someone engage in a certain activity or that something happen to her. When the travel agent develops a vacation plan for me, he is planning for me in the first sense, whereas when Able intended that Baker and Charlie paint the house together, he was planning for them in the second sense. It is possible, of course, for someone both to formulate a plan for another to follow and intend that they follow that plan, as the alienated house painters case shows.

In order for a plan to be shared, must it be planned for the group merely in the first sense, merely in the second sense or both? Let's say that I want my children to clean the mess they created in the family room and, as a result, ask them to pick up their toys. I make it clear to each of them that I am not

demanding that they do so, only requesting it. Suppose further that by some miracle they actually pick up their toys and do so only because each wants to please me. I think that we would say that they cleaned the family room together. This is so despite the fact that no one intended that my children clean the room. I did not intend that they clean the room because I did not think that my request would settle anything for them – I only hoped that it would. Nor did either of my children intend that they clean the room. Each child simply wanted to please me and could care less whether the other one did their job (in fact, they hoped that the other would ignore my request because it would make them look better in comparison). Yet, there is shared agency here because I developed a plan for them to follow and they followed that plan.

It turns out, therefore, that not only is shared agency possible in the absence of shared plural intentions, it can obtain even without any plural intention. What is necessary is that someone formulates a plan for the group to follow and communicate that plan to the group. If that plan is efficacious in coordinating planning, action and resolving conflict, that is, if it satisfies the functional role that shared intentions normally serve, it will issue in shared agency.

C. The Possibility of Massively Shared Agency

We can now see that massively shared agency is indeed possible in the face of alienation. In order for a group to act together, they need not intend the success of the joint enterprise. They need only share a plan. That plan, in turn, can be developed by someone whose does intend the success of the joint activity, or at least so hopes. As long as participants accept the plan, intentionally play their parts, resolve their disputes peacefully and openly, and all of this is common knowledge, they are acting together intentionally.

To be sure, the fact of alienation presents difficult logistical problems for planners. Because alienated participants aren't usually committed to the success of the joint activity, it is likely that they will have to be given detailed guidance on how to act. It may also be necessary to create authority structures so that conflicts can be resolved and performance monitored. The task of institutional design, in other words, is to create a practice that is *so thick with plans and mesh-creating mechanisms that alienated participants end up acting in the same way as non-*

alienated ones. That activities can often be structured so that participants intentionally achieve goals that are not *their* goals accounts for the pervasiveness of massively shared agency in the world around us.

The prevalence of alienated participants in massively shared agency does require that we modify our account of shared intentional agency in one further respect. The proposal offered above requires that *all* participants accept the plan in order for the plan to be shared. Yet, only shared intentional activities involving the smallest groups could pass a test of universal acceptance, for it is inevitable that some participants will either be apathetic, lazy, misguided, rebellious, heroic, evil or on the take, and will therefore not be committed to acting on their part of the plan or letting others do likewise. The only requirement that should be imposed is that most participants are so committed. The “most” is intentionally vague, as the concepts being explicated are vague in just this way.

A plan P to J is shared by a large group G if and only if:

- (1) P was designed for members of G so that they may J by following it; and
- (2) Most participants of G accept P.

Of course, we can't credit those who don't accept the plans as having acted together with those who do.³⁸ When we speak of a large group as participating in a joint project, strictly speaking we must be referring only to some large subset of the group in question.³⁹

V. Intentions and Institutions

In Section II, I argued that authority relations are generated in a shared activity on the basis of vertically interlocking intentions. As set out in (V), one participant has J-authority over another participant when the authority intends the subject to adopt his directives as sub-plans and to revise all other sub-plans so that

³⁸ Thanks to Shelley Kagan for pointing this out to me.

³⁹ As a result, the first condition for shared intentional activity set out on page 32 should be modified as follows:

G engages in a shared intentional activity to J if and only if:

- (1) There is a shared plan for G⁺ to J, G is a subset of G⁺ and most members of G⁺ are members of G;

they mesh with the directives and the subject intends to adopt the contents of the directives as his sub-plans and to revise other sub-plans so that they mesh with the directives.

Now, it will certainly be objected that while (V) will be suitable for many joint ventures involving authority, it cannot serve as a general account for all massively shared activities. First, it cannot account for authority relations involving alienated participants. Even if I don't intend to listen to my boss, that is, even if my intentions don't vertically interlock with his, he still is my boss and he has J-authority over me (where the J is the employment activity).

Second, the property of having J-authority set out in (V) is an extremely simple normative relation, one that is highly personal and exercised solely through the issuance of orders. Authority, on the other hand, is often institutional in nature: it is highly impersonal and rarely exercised imperatively.

In this section, I would like to show how the simple relation of J-authority can be generalized so that it will have the same properties normally associated with institutional authority. My main strategy will be to highlight the various infirmities of J-authority relations that conform to (V) and demonstrate how these weaknesses, in the context of massively shared agency, create pressure towards the institutionalization of J-authority. I will then attempt to illustrate how shared plans enable this institutionalization to take root in shared activities.

A. The Office

In Section II, I attempted to make room for authority in Bratman's models of shared agency by treating it as a mesh-creating mechanism. In this respect, I argued, authority is of a piece with the techniques that Bratman discusses, namely, deliberation and bargaining: each is a method that enables participants to achieve a mesh in their subplans. Authority, however, represents a major technological advance in social ordering. Rather than requiring befuddled or squabbling participants to waste their time and energy arguing or bargaining with each other, authorities can simply cut through the doubts and confusion and impose a solution. Mesh is thrust upon the participants from on high, rather than being stitched together by them at the grass roots.

If authority is best understood as a technology of shared agency, then (V) can only be regarded as a particularly crude prototype.⁴⁰ While vertical interlocking intentions may lead to mesh-creation in shortly-lived, one-off, small-scale activities such as sailing a vacation boat, it will not do in ventures that are highly temporally-extended, recurring, complex or heavily staffed.

For example, (V) does not allow for “impersonal” authority relations, in that it requires participants to submit to the authority of a particular person. In many instances, reliance on purely personalized authority is an awkward arrangement because the normative relation does not survive the departure of the authority and hence the activity in which the relation is embedded must temporarily come to a halt.⁴¹ In order to reestablish the relationship and resume the activity, subjects will have to agree on a new authority figure. This selection process may engender a good deal of deliberation, negotiation and bargaining, precisely the sorts of activity that authorities are supposed to obviate. It, thus, may be more convenient if participants are able to commit themselves to the directives of a person that satisfies certain qualifications. The friends might adopt a plan, for example, that confers captain status on the most experienced sailor, or the one who rented the boat, or was the captain last time, or whose birthday is closest to the day that they set sail ... etc. In this way, when one captain quits or is removed, another one is ready to fill his place.

With the advent of the impersonal relation of authority comes the possibility of an “office.” Offices are relatively stable and persistent positions of power where turnover in occupancy is not only possible but expected. The Presidency of the United States, for example, is an office because it persists from term to term and its normative character does not change merely because one president vacates and a new one assumes power. Presidents come and go but the Presidency remains.⁴²

Impersonal authority relations allow for the possibility of offices because the normative relations are not tied to any particular holder of offices, but rather to the offices themselves. Someone can accept a plan committing himself to follow the orders of anyone who satisfies the qualifications appropriate to the office (e.g., was elected by a majority in a national election). This relation persists across turnover

⁴⁰ The term “technology” I borrow from Chris Kutz.

⁴¹ See Weber’s discussion of the routinization of charisma in *Economy and Society*, ed. Guenther Roth and Claus Wittich (Berkeley: University of California Press, 1978), 246-54.

⁴² Offices can exist despite the lack of an occupant, e.g., papal interregna.

in office-holders and hence does not require participants to figure out who satisfies those qualifications and to reestablish their commitment to them. Authority is also rendered more stable, because the commitments are not based on idiosyncratic attributes of the wielders of power.

The persistence of authority across different office-holders does not yet guarantee the continuity of norms created by past office-holders. Whether the captain of one trip must heed the orders issued by the captain of the previous trip may be an open question and can only be decided by reference to the intentions of the crew. Nevertheless, the persistence of an office renders the continuity of norms instituted by past office-holders easy to establish. To do so, participants can accept a plan requiring them to follow the orders of any office-holder, even if those who issued the orders are no longer in office. When a new president takes office, for example, it is not necessary for him to renew every single one of his predecessor's orders and directives. To be sure, the current president may repeal previous edicts, but the fact that he has to repeal, rather than just ignore, them is a manifestation of their normative inertia.⁴³

B. Authority Ain't in the Head

Shared plans that specify the qualifications and powers of an office-holder help remedy another infirmity associated with (V), namely, that it relies excessively on the intentions of participants in the generation of the J-authority relation. According to (V), J-authorities are J-authorities in part because they intend to be treated in this way. It follows that in order for a participant to determine whether someone has J-authority, he has to determine whether that person possesses the right intentions. He has to establish, in other words, whether their intentions actually interlock.

⁴³ One can get quite fancy in defining offices. The shared plans that define offices may set out eligibility conditions (e.g., the President of the United States must be at least 35 years old, a natural born citizen and have resided in the United States for at least 14 years), selection criteria (e.g., the President must receive a majority of votes in the electoral college) and accession requirements (e.g., the President must pledge an oath of allegiance to the Constitution before taking office.) The definition of offices may also delimit the jurisdiction of the office-holder, regarding either subject matter (e.g., Congress has the authority to regulate matters affecting interstate commerce), personality (e.g., State courts have personal jurisdiction over any person within the state or with those outside of the state but have "minimum contacts" with the state) or both. Plans may also identify the normative position of the office within a structure of authority, thus specifying how conflicts are to be resolved (e.g., the Supreme Court's rulings are binding on all lower federal and state courts on questions of federal law, but not binding on state courts on matters of state law).

In normal face-to-face interaction, it is often unproblematic to ascertain whether an individual possesses a certain intention. However, when dealings are no longer direct, it becomes more difficult to figure out what another is thinking. To be sure, the intention to be treated as an authority can be manifested precisely in the issuing of orders, for only someone who believes that they are entitled to order can sincerely order. Nevertheless, when dealing with multi-membered authoritative bodies such as a legislature, it may be impossible to ascertain whether a shared intention exists to take their orders as authoritative. Indeed, in many cases, legislators intend just the opposite. In strategic voting, for example, legislators vote for certain legislation intending that the legislation is defeated.

Notice that when participants accept rules specifying that someone has authority over particular matters and persons, it is no longer necessary that authorities intend to have their orders taken in a certain way.⁴⁴ They have J-authority by virtue of the plans conferring authority. In such circumstances, J-authority, like many other social phenomena, ain't in the head. This is not to say that intentions are completely irrelevant to the generation of J-authority. Quite the contrary, most participants must accept the plans conferring that authority. What have become irrelevant are the intentions of any particular participant. A participant has J-authority regardless of what she intends.

We might call the process of uncoupling particular intentions from the generation of normative relations as the *institutionalization* of normativity. The actions of any participant may have normative significance in an institutionalized practice despite the fact that they acted without certain intentions. They may significance in the practice, in other words, because the *other* participants intend to give them such significance. The institutionalization of authority, therefore, detaches the intentions of authorities from their normative status. By the same token, a participant may be subject to another's authority even though they lack the requisite commitments to the authority.⁴⁵ The "subject" is "J-obligated" to obey in part because the plans purport to obligate him to obey.

⁴⁴ The discussion that follows was stimulated by a question posed to me by John Campbell.

⁴⁵ When the subject is not committed to the plans of the shared activity, J-authority may not even constitute rational authority for the subject.

C. *Replacing Orders*

The technology set out in (V) not only relies heavily on the intention of authorities to generate the authority relation, but also in the generation of the authoritative directives themselves. According to (V), the directives created by the authority are *orders* and orders are solely the product of intentions: A has ordered S to f only if A intends the expression of his intention that S f as a conclusive reason for S to f. As such, the problems of ascertaining the intentions of authorities just mentioned are applicable here as well.

Thus, in order to identify some utterance as an order, subjects must determine whether the expression of the utterance was intended to be treated as a reason for them to engage in a particular action. The problems of ascertaining collective intent for multi-membered bodies will be particularly acute here. Furthermore, in order to ascertain the content of an order – whether some order is an order to f – subjects must figure out whether the authority intended that they f. This will likely to be a burdensome task even when dealing with individual authorities, for it is always possible that they will have incomplete, vague or confused understanding of what they are doing. Indeed, in the case of modern legislation, legislators often have little or no idea about the actions they are forbidding, permitting or empowering. Requiring subjects to make inquiries into the subjective intent of authorities in order to determine the content of their directives will likely be extremely demanding.

The main method for alleviating this burden is, again, to increase the level of institutionalization in the practice. With respect to identifying the existence of directives, plans can be devised and adopted that specify when power has been properly exercised. Thus, when the king affixes his seal, the Pope speaks *ex cathedra*, the legislator says “Aye” after the bill has been read aloud twice or the bank manager marks the loan application with the “Approved” stamp, these acts have normative significance for their subjects regardless of the intentions with which they are performed. With respect to identifying the content of directives, rules of interpretation can be adopted that specify how to interpret the utterances that have been issued or texts that have been approved. These rules, of course, might direct subjects to examine the subjective intent of the authorities, in which case the informational burden will not lessened.

VI. Conclusion

In this paper, I have tried to illustrate how the philosophy of action might accommodate massively shared agency. In order to show that authority and cooperation are compatible, I suggested that authorities be understood as mesh-creating mechanisms. Their function in shared activities is to obviate the need for costly negotiation and bargaining through the imposition of mesh from above. Authorities are efficient technologies for resolving conflict and ensuring that everyone involved is on the page.

In order to show how alienation and cooperation is possible, I began by criticizing hypercommitted models of shared agency. The philosophy of action must abandon the view that all, or even most, of the participants in a shared activity are committed to contributing to the success of that very activity. I claimed instead that shared intentional activity can take place among alienated participants when a shared plan exists, most participants play their part, resolve their disagreements peacefully and openly, and all of this is common knowledge. Shared plans can exist, I maintained, even if participants are not committed to its success. They must simply accept their part of the plan and be committed not to interfere with the others playing their parts.

Finally, I argued that shared plans can explain how institutional authority is possible within shared activities. When plans accord power to some participants over others participants and those plans are accepted and followed by most, institutional authority is created. In other words, normative relations between an authority and a subject will obtain even though either, or neither, intends the relation to obtain just as long as the other participants take the normative relation seriously.

While this account shows how massively shared agency involving institutional authority is possible, it is important to note one considerable limitation. Consider the legal system. I think it is plausible to regard legal activity as an instance of shared agency, namely, that the creation and application of the law is something that legal officials do together. However, our model cannot handle the legal case. The reason is that the account of authority developed thus far has been confined to those instances where authority is claimed and exercised over

participants. To have J-authority as I have explicated it is to have authority over those who J. Yet, legal authority is exercised not only over legal officials but over ordinary citizens as well. Since ordinary citizens are not involved in the creation and application of the law, they are not strictly speaking legal participants and, hence, the full relation of legal authority cannot be expressed in the model.

A natural thought is that legal authority can be established using shared plans. Imagine that the fundamental rules of a legal system are shared plans and these plans not only accord legal power to legal officials over other legal officials, but over non-officials as well. When officials accept these plans, therefore, they will accept certain authority relations that hold both between themselves and between officials and non-officials. Thus, when some legal official issues a directive that applies to an ordinary citizen, other officials regard that directive as valid. In this way, legal authority over non-participants may be created. Just as the intentions of particular participants are irrelevant to the generation of J-authority, so too are the intentions of non-participants. To be sure, the non-participant may accept the subject designation, but this is not what makes the subject a subject. It is the fact that legal participants accept the shared plans that treat him as a subject, coupled with the general conformity of the non-participants to the plans of the legal participants, which establishes his subordinate legal status.

Whether shared plans can be made to pull this much weight is a matter beyond the scope of this paper. But if they could, an exciting possibility would open up. It would mean that the philosophy of law (and of social institutions more generally) could be understood as a branch of the philosophy of action. Since institutional activity is shared activity, the philosophical resources developed to study shared agency can be exploited for the analysis of these complex phenomena. In this paper I have tried to fashion one more tool in the hope that it may soon come in handy.⁴⁶

⁴⁶ For an attempt to exploit recent advances in the philosophy of action to understand legal phenomena, see my *Legality*, chapters 5-7 (Cambridge, Mass: Harvard UP, 2011).