



A Planning Theory of Acting Together*

ABSTRACT: *We have the capacity to act together in shared intentional and shared cooperative ways. This lecture argues that our capacity for the plan-based, mind-supported cross-temporal organization of our individual activities, together with certain further elements, suffices for our capacity for the mind-supported, small-scale social organization characteristic of acting together. These two fundamental forms of human practical organization—diachronic and small-scale social—are for us grounded in a common core: our capacity for planning agency.*

KEYWORDS: shared intentional action, shared intention, intentions, plan rationality, mesh in sub-plans, mutual responsiveness, persistence interdependence, interlocking, mutual obligation

We have the capacity to act together in shared intentional and shared cooperative ways. Think about string quartets, or our painting the house together, or our having a conversation, or together engaging in a scientific research project, or singing a duet, or dancing together. Such shared intentional/shared cooperative activities play fundamental roles in our lives. They matter to us both instrumentally and non-instrumentally: they are in many cases how we get things done, and in many cases, we see them as intrinsically valuable aspects of our human lives. They are of fundamental interest in many areas of humanistic and social scientific research. Can we articulate theoretical structures that can help guide our understanding of and further research concerning these ways in which we live together?

When we act together in these ways, we are not each simply acting in light of expectations of the actions of others while knowing that those actions of others depend on their expectations of our actions. To draw on Margaret Gilbert's example, merely publicly walking alongside each other on a crowded sidewalk

* These articles are the first in a special series devoted to winners of the Dr. Martin R. Lebowitz and Eve Lewellis Lebowitz Prize for philosophical achievement and contribution awarded by the Phi Beta Kappa Society in conjunction with the American Philosophical Association. Michael Bratman and Margaret Gilbert are the 2020 Lebowitz Prize awardees.

This is a lightly edited version of my contribution to the Dr. Martin R. Lebowitz and Eve Lewellis Lebowitz Prize Symposium at the 2021 Pacific Division Meetings of the American Philosophical Association as part of a debate with Margaret Gilbert on the topic 'What is it to act together?'. My lecture is based on ideas developed over the past three decades (Bratman 1987, 1999, 2014); please see those works for further details and references to the important work of others. My lecture was given together with Margaret Gilbert's contribution to this symposium, 'A Simple Theory of Acting Together'. I respond briefly to Gilbert's lecture in 'Two Faces of Our Idea of Acting Together (also in this issue). Gilbert and I have been involved in a shared cooperative dialogue about these issues over many years, a dialogue from which I have learned a great deal. I take this occasion to thank Margaret for these many fruitful and engaging exchanges.

without colliding, while involving complex forms of mutual responsiveness, is not yet walking together in a shared intentional way. Can we articulate conditions that go beyond such strategic interaction and are sufficient for and illuminating of our acting together?

In our acting together, our minds support distinctive forms of social organization. Taken together, our individual actions tend to track and to support relevant joint activity. The practical thinking of each is sensitive to and supportive of the associated practical thinking of others in ways that tend to track and to support the joint activity. Each sees realization of the joint activity as a common standard of success; and this shapes their thought and action.

How does this work? How do our minds together support such remarkable social organization? I begin by first turning to an analogous question about how our minds support our individual agency over time.

Individual human actions are normally embedded in striking forms of mind-induced cross-temporal organization. Think about growing food in a garden or taking an extended trip. Now, not all goal-directed agents are planning agents; but we are. For agents like us, these temporally extended activities are typically organized by plans and planning. My plan is, say, to visit my sister in Boston next month. This shapes my thinking about what to do in the nearer future, for example get an airline ticket. And it filters out certain other options, for example, getting a ticket to Paris instead.

My plan also allows me to leave unsettled until later certain matters I know I will need to settle but can trust myself to settle later in a timely fashion—for example, how to get from the airport to my sister's house. This illustrates a general point: we normally settle in advance on such prior plans without settling on a fully worked out plan. After all, we humans are limited in our mental resources for figuring out such details in advance. And given limits in our knowledge, we are limited in how much it makes sense to figure out such details in advance rather than await more information. A basic feature of human prior plans is that they are typically partial. And downstream planning and practical thinking, including the weighing of reasons, is then structured by these prior, partial plans.

These prior, partial plans engage rational pressure to fill them in as time goes by—a failure to do that would threaten criticizable means-end incoherence. They engage rational pressure to fill in one's different plans in ways that are by one's lights co-possible—a failure to do that would threaten criticizable plan inconsistency. And they engage rational pressure in favor of stability over time, in part by way of rational pressure against reconsideration of prior plans in the absence of new information.

These pressures of plan rationality lie behind basic roles of such plans in the psychological functioning that supports much of our mind-shaped individual, diachronic practical organization. Our prior, partial plans pose problems of means for further practical reasoning. To avoid means-end incoherence I need, for example, to figure out how to get to Boston. These plans engage consistency constraints on that reasoning and its outputs, thereby filtering out certain options that might have been solutions to problems of means. And rational pressures for stability help support the role of such plans in cross-temporal organization. In

these ways, the rational functioning of our prior partial plans provides a background framework that structures downstream practical thinking.

We have then a sketch of the way in which the distinctive, rational functioning of plans and planning can help explain how our human minds support the cross-temporal organization of individual human action. Intentions are elements in this planning system.

Return now to our acting together in shared intentional ways. How do our minds together support the social organization of thought and action involved in, say, our painting the house together? My conjecture is that we can see this mind-supported social organization as in relevant cases emerging from the cited plan-based, mind-supported diachronic organization of individual human action.

The idea is not that simply by having the capacity for plan-based cross-temporal organization of one's own activities one thereby has the capacity for the small-scale social organization characteristic of acting together. There may well be creatures who are individual planning agents but do not thereby have the capacity for shared intentionality (though they are good at strategic interaction). But we humans have both capacities. So, we can ask whether our capacity for plan-based, mind-supported cross-temporal organization of our individual activities, together with certain further elements, suffices for our capacity for the mind-supported small-scale social organization characteristic of acting together.

And my conjecture is that we can in this way provide illuminating and broadly individualistic sufficient conditions for important forms of acting together. This would suffice to show that we can understand these forms of acting together without appealing to an element that is fundamentally discontinuous from the metaphysics of individual, temporally extended human agency. To do this, we need not sort out the further question of whether we thereby provide necessary conditions for acting together. Ours is a strategy of sufficiency.

To explore this conjecture, I begin with the idea that just as individual intentional action is normally tied to guidance by individual intention, shared intentional action is normally tied to guidance by shared intention. Our shared intentional activity of painting the house together is normally tied to guidance by our shared intention to paint it together. But how should we understand such shared intentions?

In the case of individual intention in favor of one's own action, we focused on roles in plan-shaped organization of thought and action over time. Proceeding similarly with our shared intention to paint together, we can say that its characteristic functional roles include the interpersonal coordination of our action and planning in pursuit of our painting together, and the provision of a shared framework for relevant forms of thinking together, including bargaining and shared deliberation. And we can say that this functioning of shared intention normally satisfies basic constraints of social, interpersonal consistency in plan and social means-end coherence in plan (where that involves timely specification of means across the social web). After all, relevant failures of social consistency and means-end coherence would normally block successful guidance by our shared intention.

We now ask whether we can articulate a structure of interconnected plan states of involved individuals that would, when functioning in the rationality-guided ways

characteristic of individual cross-temporally organized planning agency, play these characteristic roles of shared intention in part by supporting conformity to constraints of social consistency and social means-end coherence. This would be a plan-theoretic construction of shared intention, one that shows how the characteristic functioning of shared intention can emerge from the proper functioning of the elements in that construction. And this would support the idea that our capacity for plan-based, mind-supported cross-temporal organization of our individual activities, together with certain further elements, suffices for our capacity for the mind-supported, small-scale social organization characteristic of acting together.

The idea then is to construct shared intention out of building blocks that come primarily from our planning theory of individual diachronically organized activity. And the idea is to demonstrate that the proper functioning of that construction—functioning in accord with norms of individual plan rationality—will normally induce the kind of social organization of thought and action characteristic of our acting together. In this way, the mind-supported social organization characteristic of our acting together emerges from the plan-based, mind-supported diachronic organization of individual action.

Suppose that you and I share an intention to travel together to Boston. What would such a plan theoretic construction of this shared intention be? I proceed to articulate a quintet of building blocks, adding commentary along the way.

First: each intends that we travel together to Boston. Here we understand “travel together” in a way that is neutral concerning shared intentionality. (It includes, for example, scenarios in which we each are simply traveling to Boston on the same plane.) And we understand my intention that we travel together as doubly reflexive in the sense that it involves both my intention to act accordingly, and a specification of who is, from my perspective, included within the referent of ‘we’.

But can I intend not just to act but, further, that we act? Yes. I will explain why once we get more building blocks within view.

Second, each intends that we travel together to Boston in part by way of the corresponding intention of the other. In this sense our intentions referentially interlock. This precludes a case in which, for example, each intends that the joint travel proceed by way of throwing the other into the car trunk.

Third, each intends that we travel together by way of sub-plans of each that mesh with each other, and by way of mutual responsiveness of each to each, mutual responsiveness that tracks our joint travel.

Are these second and third building blocks too demanding for limited minds like ours? Well, these complex contents may be only tacit or implicit. The conscious and explicit contents of relevant intentions may be simply that we do this together. But these intentions might support dispositions to track and to be responsive to various roles of intentions of others. This can include a disposition to respond to one’s failure to support those roles of others with a kind of ‘oh darn it!’ reaction. Such a web of dispositions may support the attribution of the cited more complex, but only implicit, contents.

We need also to remember our strategy of sufficiency. We can allow that there can be different psychological mechanisms that do the work in our acting together that

these building blocks do in our model of acting together. In articulating these building blocks, we articulate what this work is.

Fourth, each person's intention that we travel together to Boston is set to persist, other things equal, if but only if the other's corresponding intention persists. In this sense there is, other things equal, persistence interdependence. There are different forms of such persistence interdependence. There may be this interdependence because we each know that the desirability of the trip depends on its being a joint trip—it is no fun going alone. There may be this interdependence because we each know that the realistic possibility of success depends on the participation of both—perhaps we each need the other's help in boarding the plane. Or there may be this interdependence because there has been a kind of interaction between us (perhaps, an exchange of promises) that supports relevant moral obligations to continue so long as the other does. Such promissory obligations are not essential to shared intention. Nevertheless, in certain cases moral obligations in this neighborhood can support the relevant interdependence. And given interdependence in any of these three forms (or perhaps others), I am in a position to some extent to trust that you will continue with our shared activity if I do.

Fifth, and finally: these multiple building blocks are to a large extent out in the open: they are each epistemically accessible—with respect to knowledge of others, in evidence-based ways—to each of the participants.

We can now return to the very idea of my intending that we act. A worry is that I am not in a position to control whether we act, so I cannot intend that we act. In special asymmetric cases I can be in a position to control how another acts, and so can be in a position to intend that he so acts. But how in shared intention can each of us have the control needed for each of us to intend our joint action?

My answer is this: Given persistence interdependence, I can, in intending our joint action, be sufficiently confident that, in part because of my intention, you will also so intend. (Which is not to say that you will have lost control over your intentions.) And I can be sufficiently confident that, other things equal, our intentions will together tend to lead to our joint action. And analogously for you. Each of us is in a position to some extent to trust that the other will continue participating in our shared thought and action if she herself does. So, each of us can in this way be sufficiently confident that if she continues to intend that we act then we will. So, each of us can in this way be sufficiently confident that she controls whether we so act. And we can each be right about this. So, we can each intend that we so act.

This interdependence and interlocking will get more complicated in cases of larger groups. But I put these complexities aside here.

And now my basic claim is that if this construction of interrelated plan states functions properly—in accordance with the cited norms of plan rationality—we will normally be acting together in a shared intentional way. The functioning characteristic of shared intentional activity is thereby grounded in this plan-theoretic, broadly individualistic construction.

In partial support of this basic claim: Suppose that I intend that we travel together by way of your corresponding intention, mutual responsiveness, and mesh in sub-plans. Given this intention of mine, I will be set to play my part in our joint travel. I will also be set to be responsive to your efforts to play your part. In both

cases, this is a matter of pressures of means-end coherence and consistency applying to my intention, given my beliefs about you. Further, if I know that you need my help—with the luggage, say—the requirement of means-end coherence as applied to my intention induces rational pressure on me to help you. After all, one aspect of what I intend is your participation in our joint activity. And the requirement of consistency as applied to my intention induces rational pressure on me to avoid ways of acting that would be incompatible with your playing your role. Further, given that I know that there is persistence interdependence between our intentions, I have reason to assume that so long as I continue so to intend so will you. So, I have reason to assume that for me to satisfy rational demands of means-end coherence and consistency on my intentions I will need to work with you and your intentions and plans.

When the rational pressures of individual plan rationality apply to the complex intention of mine that is built into this construction—an intention that has a distinctive, social content and is embedded in a distinctive social context—they induce rational pressures on me in the direction of support of social means-end coherence, social consistency, and social effectiveness. These rational pressures thereby support my role in our joint activity. Given that my intention has these contents, in these contexts, we explain how standard forms of coordination-and-effectiveness-tracking plan-theoretic thinking on my part can rationally support our shared intentional activity. And analogously for you. So, if all goes well the proper functioning of our shared intention, as constructed, normally issues in our shared intentional activity.

Shared intentional activity will, then, emerge from the rational dynamics of this broadly individualist plan-theoretic construction. And this supports our conjecture that our capacity for plan-based, mind-supported cross-temporal organization of our individual activities, together with certain further elements (in particular, the cited social contents and contexts of intentions of each), suffices for our capacity for the mind-supported, small-scale social organization characteristic of acting together.

This construction allows for the possibility that we each participate in the joint activity for different reasons. Even if we were each to have different reasons for travelling together to Boston, the cited construction of our shared intention would normally induce the social coordination and organization characteristic of shared intentional activity. Such sharing of intentions can support our sociality even given significant pluralism in and divergence of background reasons.

At the heart of this approach to our shared intentionality are norms of individual plan rationality: norms of consistency, coherence, and stability. There are here difficult issues about the significance of these norms. Why does it matter whether I am being rational? For present purposes, however, I assume that if you are a planning agent your conforming to these norms matters and your violation of these norms is substantively criticizable (Bratman 2018). So, in the context of our shared intention, if I knowingly fail to support relevant social consistency, coherence, and effectiveness I am, other things equal, criticizably irrational. And similarly, for you.

Suppose that we share an intention to go together to Boston but that I act intentionally in ways I know to be incompatible with this—perhaps I do not get

on the plane, or I fail to help you with your luggage. I am, other things equal, criticizably irrational. Have I thereby violated an obligation to you?

Well, we have noted that in some cases the interdependence involved in our shared intention is grounded in interactions—perhaps an exchange of promises—that induce relevant moral obligations of each to each. But we have seen this as a special case; and we have noted that this interdependence might instead be grounded in other ways.

Might it nevertheless be true that there is a quite general connection between shared intention and mutual obligation? One idea here is that certain mutual obligations suffice for shared intention. A second idea is that shared intention quite generally involves obligations of each to the other. I will focus on this second idea.

Recall that to support a necessary connection between shared intention and pressures of practical rationality we appealed to norms of plan rationality and their application to conditions of shared intention. Analogously, to support a necessary connection between shared intention and mutual obligation, we would need to appeal to, and defend in normative reflection, substantive norms of obligation and the ways they apply to the conditions of shared intention. What norms?

An answer would need to confront hard cases. What if the shared activity is trivial and short-lived, as when we walk together for a brief moment? What if the shared activity is immoral, as when we together engage in torture? What if your participation in the shared intention is coerced by me? What if your participation in our shared intention arises from significant deception by me? What if each has been reading Ayn Rand and, though we participate in the sort of shared intention I have described, each insists: ‘but, no obligations’?

Further, it seems that the basic explanatory role of shared intention does not quite generally depend on such obligations, or beliefs about such obligations. A string quartet involves striking forms of cooperative mutual responsiveness. And it seems that in many cases what lies behind this cooperative mutual responsiveness need not be a matter of obligations of each to each, or even beliefs about such obligations. In a common case, I sync my cello part with your viola part because I intend that we play the quartet by way of meshing sub-plans and mutual responsiveness, relevant matters are out in the open, and I am (implicitly) guided by norms of consistency and means-end coherence. And analogously for the others in the quartet. In such cases the interlocking, interdependent intentions and plans of the participants, taken together with extraordinary skills, explain (at a relevant level of explanation) these remarkable forms of social organization. The sufficiency of these explanatory conditions for such social organization does not depend on the presence of mutual obligations even if such there be. Nor need it depend on beliefs about such obligations. Our fans of Ayn Rand can play together in a quartet.

I infer that the explanatory structures that directly underlie such cases of acting together involve interconnected intentions and plans, and that these explanations need not in general directly appeal to interpersonal obligations, or beliefs about such obligations, even if such there be. Further, when we try to say what interpersonal obligations do enter into our shared intentional agency, we need to go beyond norms of plan rationality and appeal to further normative premises

that themselves need normative defense. In light of these observations, my proposal is that in pursuit of a model of shared agency that can support wide-ranging research in philosophy and the social sciences, we treat our planning model as basic, while leaving room for further reflection on ways in which these planning structures will in many cases engage, or be thought to engage, further potentially defensible norms of obligation.

That said, such obligations will be common within our human shared activity and in some cases support relevant interdependence. There are in the neighborhood broadly moral pressures to be helpful and courteous. In many cases, there will be moral obligations grounded in specific forms of interaction, such as assurances or intentionally induced reliance. And there may be at work other kinds of moral obligations. In each case, however, the claim that there are these obligations needs to be defended in substantive normative theorizing. This will involve defense of substantive norms that go beyond norms of plan rationality.

There is, then, within our theory of acting together, a contrast between the status of such substantive norms of interpersonal obligation and the status of norms of individual plan rationality. These norms of plan rationality are at the heart of our model of human planning agency. They are, on this view, non-separable aspects of our individual planning agency and its characteristic diachronic organization. They are thereby non-separable aspects of our shared intentional agency and its characteristic social organization. Norms of interpersonal obligation—including especially moral obligations—play fundamental roles in our human lives and are frequently engaged by our shared intentional activities. But they are nevertheless separable from the underlying structure of our shared intentionality.

This takes us to a final idea. We have seen how our capacity for planning agency can play basic roles in our capacity for practical organization over time, and thereby in our capacity for shared agency. Given our strategy of sufficiency, we cannot strictly deduce that this is what in fact underlies these forms of human practical organization. But we can proceed by way of inference to the best explanation and defeasibly conclude that these two fundamental forms of human practical organization—diachronic and small-scale social—are for us grounded in a common core: our capacity for planning agency.

MICHAEL E. BRATMAN 

STANFORD UNIVERSITY

bratman@stanford.edu

References

- Bratman, Michael E. [1987] (1999) *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press. Reprint Stanford: CSLI Publications.
- Bratman, Michael E. (1999) *Faces of Intention*. Cambridge: Cambridge University Press.
- Bratman, Michael E. (2014) *Shared Agency: A Planning Theory of Acting Together*. New York: Oxford University Press.
- Bratman, Michael E. (2018) *Planning, Time, and Self-Governance: Essays in Practical Rationality*. New York: Oxford University Press.