

Book Symposium

Open Access

Abraham Sesshu Roth

Practical Intersubjectivity and Normative Guidance: Bratman on Shared Agency

Abstract: In an important new book on shared agency, Michael Bratman develops an account of the normative demand for the coordination of intentions amongst participants in shared agency. Bratman seeks to understand this form of normative guidance in terms of that associated with individual planning intentions. I give reasons to resist his form of reductionism. In addition, I note how Bratman's discussion raises the interesting issue of the function or purpose of shared intention and of shared agency more generally. According to Bratman, the function of shared intention is to promote interpersonal coordination of intention and action. I suggest that power sharing amongst participants must also be included as a function of shared intention.

Keywords: Shared agency; Practical intersubjectivity; Intention; Normative guidance; Power sharing.

DOI 10.1515/jso-2014-0048

1. Some coordination in action is necessary for individuals to act *together*. We are not walking together if we trip all over one another. We fail to travel together if one of us is so late getting to the airport that the other takes the flight by herself. It's also necessary for shared or joint activity that there be sufficient coordination or (to use Bratman's term) mesh in the relevant *intentions* of the participants. There is a problem for our prospective drive to Vegas if I intend that we get there by leaving in the morning, and you intend that we leave in the evening.

This condition on intentions is not merely that there is shared agency only if there is some level of coordination in intentions. There is the further thought that the activity-related practical thought of participants is in some sense – perhaps explicitly, perhaps implicitly – guided and regulated by a demand for coordination. The compatibility of one's intentions with those of fellow participants is, we

Abe Roth, Ohio State University – Philosophy 350 University Hall 230 N. Oval Mall, Columbus, Ohio 43210-1365, USA, Tel.: + 6142927502, e-mail: roth.263@osu.edu

 ©2014, Abraham Sesshu Roth, published by De Gruyter.

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 3.0 License.

might say, normative for each participant. That's to say that there are interpersonal norms relating the intentions of participants in shared activity.

Calling the coordination of intentions a normative requirement does not entail that one *should* so coordinate and do one's part in shared activity. Sometimes one should not participate, such as when the activity is directed toward ends that are morally heinous. Depending on the situation, there might be no reason all things considered – nor even *pro tanto* – to coordinate one's intentions with fellow participants. Nevertheless, if what one is doing is acting *with* the others, then one is guiding one's practical thought toward interpersonal coordination of intention.

It's in this sense that the intentions of fellow participants have a significance for one's practical thought unlike the intentions of third parties, and akin to the significance one's own intentions have for one's practical reasoning. Just as I guide my practical thought so that my own intentions are coordinated with one another, so a participant guides her practical thought so as to secure coordination of intentions with fellow participants. This is what I have in mind when I speak of *practical intersubjectivity* holding between participants in shared activity.¹

One of the many virtues of Bratman's *Shared Agency* (OUP 2014) is that it recognizes practical intersubjectivity as an important element of shared agency.² Although I don't agree with the broadly reductive story Bratman tells, his is a powerful theory. The book covers various issues with clarity and care. I cannot do justice to all of it. I'll raise questions about the part concerning practical intersubjectivity, and end by showing how Bratman's approach raises an interesting new issue for the literature on shared agency.

2. For Bratman, practical intersubjectivity comes by way of the idea of shared intention. Bratman sees shared intention as the difference between joint or shared activity and cases that amount to little more than each individual acting on his or her own. The activity of various individuals is shared if it results from and is explained by a shared intention.

But what is shared intention? Bratman's insightful suggestion is that we understand it functionally, as something playing a distinctive social role. He looks to individual intention as a model. According to Bratman's influential planning theory, individual intentions are embedded in a structure of planning states and associated forms of reasoning. The intention to ϕ leads one to stop thinking about alternatives and to turn instead to planning on how to go about ϕ -ing. The

¹ I will set aside the possibility that practical intersubjectivity should be put in stronger terms, where interpersonal principles have genuine normative force and the participant is not just guided by a demand for coordination of intentions, *but also ought or has reason to*.

² All page references unless otherwise noted are to this work.

cognitive structures that embody these tendencies help secure an inter-temporal coordination of one's actions that allows for more complex, temporally extended forms of agency (15).

Corresponding to the planning roles are *norms* of intending. One such is means-end coherence, requiring one to intend the necessary means to an intended end. Another requirement is intention consistency. Part of what it is to be an intending agent in Bratman's planning sense is that one is guided (be it explicitly or implicitly) by such norms.

Analogously, shared intention helps explain *our* activity and *our* planning; the central idea here seems to be how shared intention allows for *each individual's* intention and action to be coordinated with the rest (27). As with the individual case, the functional role of shared intention comes with associated norms, such as those of consistency and coherence of relevant intentions, this time *across* individuals (27).

With these and other social norms, we have arrived at the point in Bratman's picture capturing the interpersonal norms important for practical intersubjectivity. Part of what it is to *act together or jointly* is that such norms have a significance for the practical thought and action of the participants.

So, the norms of individual intention figure in individual planning agency, and Bratman likewise sees the social norms as significant for shared agency (28). But the story here, at least in the basic case, is importantly different. The significance of the social norms is understood in terms of the significance of the *individual* norms. The norms of planning agency accepted by the agent, qua planning agent, are more basic, and serve as material from which the social norms are constructed.

This comes out when Bratman turns to what it is that plays the role constitutive of shared intention. While in principle shared intention may be multiply realized, Bratman devotes much of the book to showing how a shared intention might be instantiated in a complex of individual attitudes spanning across the participants. At the core of this complex is an intention had by each participant of the form *I intend that we J*, where J is some cooperatively neutral specification of the shared activity. Thus, at least in the basic case, as a participant in shared activity J, I am guided by the norms of individual planning agency that come in play when I form this intention. This will, in effect, secure the *interpersonal* consistency of my intentions with those of the other participants. Thus, if I intend that we drive together to Vegas, and I intend that we leave in the morning and you intend that we leave in the evening, then it's unlikely that my intention that we drive to Vegas together will be satisfied. So the individual norm of means-end coherence that governs my intention that we drive to Vegas together will ensure that I form intentions about how we go about it that are compatible with your intentions. At least, that is the idea.

It's in this way that Bratman's approach is reductionist. Bratman says,

The idea, roughly, is that the social-norm-assessable social functioning characteristic of shared intention emerges from the individual-norm-assessable and individual-norm-guided functioning of relevant structure of interrelated intentions of the individuals, as those intentions of individuals are understood by the [individual] planning theory. (32)

Bratman thinks that the step from individual planning theory to shared intention and modest sociality is a "fairly conservative step" (31), and he challenges approaches that posit "fundamental discontinuities" between individual and shared agency (8–9, 32). The right individual intentions in the appropriate circumstances will, Bratman hopes, serve as a bridge linking the norms of individual planning agency with the interpersonal norms. The "bridge intentions" thereby secure the special significance the intentions of fellow participants have for one, and do so in a way that is "conceptually and metaphysically conservative" (26). Thus:

We seek...a construction of interconnected intentions and other related attitudes of the individuals in appropriate contexts that would, when functioning in the norm-guided ways highlighted by the planning theory of intentions of individuals, play the roles characteristic of shared intention. And we try to see conformity to central norms of social rationality characteristic of shared intention...as primarily emerging from guidance by norms of individual plan rationality that apply directly to the relevant interrelated structures at the individual level. (32)

3. At this point, one may begin to wonder about Bratman's picture of the significance of the interpersonal norms. The fact that these norms are not violated is *incidental* to a participant's being guided by individual norms. The interpersonal norms are not, it seems, what guide the participants' practical thought.

In a sense, this is precisely the point of Bratman's reductive approach. Guidance by interpersonal norms and practical intersubjectivity should not appear in the set of sufficient conditions Bratman is aiming to offer (33). Rather, guidance by the norms of individual planning intentions with the specific contents and in the circumstances set out is what *guidance by interpersonal norms and practical intersubjectivity* amount to.³

If there is a worry for Bratman's analysis, we need to consider whether the conditions Bratman supplies as sufficient for shared intention do in fact add up to guidance by interpersonal norms. To that end, we need to look more closely at what

³ Bratman might say that non-violation of interpersonal norms is all that's required, and simply deny *guidance* by them as a condition for shared agency. Some of his remarks suggest this (see 32, 33, but see 28). If Bratman wants to go this way, the counterexamples below can be modified to show that some form of guidance by interpersonal norms is necessary for shared agency.

Bratman thinks occupies the social roles constitutive of shared intention. I've already noted that it is an interrelated set of individual attitudes, including the intention that we J. Is there something here that establishes not just the satisfaction of interpersonal norms, but also their normative significance as guiding for the participants?

On Bratman's account of the two person case, I intend that we J *in part by way of your intention that we J* and our meshing subplans (85). Wouldn't this show that your intention matters to me in a way that some third party's intention does not? Wouldn't I then be concerned with the compatibility of your intention and plans with mine, and guide my practical reasoning in accord with this?

Not necessarily. Consider the case of the *fungible partner*, where my intention that we J in part by way of your intention that we J is entirely compatible with a regard for your intention that is purely opportunistic and instrumental. If you plan to implement your intention in a way at all incompatible with how I plan to implement mine, then I would sooner dump you as a partner in favor of someone else than acknowledge that you have a say in the matter and try to accommodate your way of doing things. So it doesn't appear that my intention *that we J in part by way of your intention* suffices to capture the significance your intention is supposed to have for my practical thought.

Another element in Bratman's proposal is the idea of mutual responsiveness. For example, Bratman requires that each participant intends that that we J "by way of relevant mutual responsiveness in sub-intention and action" (85). And joint or shared activity J occurs when shared intention leads to the J-ing via a mutual responsiveness in sub-intention and action that tracks the J-ing. No such mutual responsiveness was present in the fungible partner case where I unceremoniously dump you. So perhaps the mutual responsiveness is what's crucial for practical intersubjectivity and guidance by interpersonal norms.

Explicating mutual responsiveness, Bratman says "...each is responsive to the intentions and actions of the other in ways that track the intended end of the joint action – where this is all out in the open" (79; see also 80). I think that Bratman would have to count the following case of *playing chicken* as mutual responsiveness, although it pretty clearly doesn't amount to practical intersubjectivity. Suppose I intend that we will J by means M1, and you intend that we J by some incompatible means M2. Our intentions don't satisfy interpersonal consistency. I hold on to my intention confident that you will back down, and you see what I'm doing and eventually back down – not because you are convinced that this is how best to resolve the difference but simply because I don't change my intention. You certainly are responsive to my intentions in a way that tracks our end that we J. But arguably so am I: if I'm confident that you will relent, then am I not sufficiently "responsive to the intentions and actions of the other in ways that track the end of joint action"?

The upshot, then, is that Bratman's proposed conditions are not sufficient for the guidance by social norms and the practical intersubjectivity essential to shared agency.

4. Let me turn to a different matter. Bratman thinks that the explanatory significance of the norms of *individual* planning agency – the fact that they play the guiding role that they do – is due to its acceptance by the agent. Thus,

Norms enter indirectly into a standard explanation of the normal ways in which these plan states play these roles in planning agency. ...In particular, implicit or tacit acceptance by planning agents helps explain how these plan states play these roles... (16, 32)

We find this notion of acceptance in Gibbard, whom Bratman cites. The idea is that acceptance is realized in the psychology of the individual and is “meant to figure in an explanatory theory of human experience and action.”⁴

We've seen that Bratman's story for the *interpersonal* social norms of shared intention is different. They are not accepted in the way that the norms of planning agency are, at least not in the basic case. The significance of the interpersonal norms is constituted by the norms of individual planning along with the suitable “bridge” intentions. But the Gibbardian picture of norm acceptance raises questions about the Bratman's approach to interpersonal norms.

First of all, one may wonder whether the apparatus of norm acceptance is best thought of in individualistic terms. Why assume that norm acceptance is a personal affair? Why not think of it as a manifestation of shared agency or as somehow presupposing it? The worry for Bratman, of course, is that his account would then be circular; the normativity of planning agency upon which shared agency rests would itself presuppose shared agency. As it turns out, Gibbard's picture of norm acceptance is social and is naturally interpreted as presupposing shared agency. Acceptance for Gibbard is connected to normative governance:

Working out in community what to do, what to think, and how to feel in absent situations... must presumably influence what we do, think, and feel when faced with like situations. I shall call this influence *normative governance*. It is in this governance of action, belief, and emotion that we might find a place for phenomena that constitute acceptance of norms... When we work out at a distance, in community, what to do or think or feel in a situation we are discussing, we come to accept norms for the situation. (Gibbard *op. cit.*, 72-3; see also 74, 75.)

⁴ Allan Gibbard, *Wise Choices, Apt Feelings* (Cambridge, MA: Harvard University Press, 1990), 46.

The activity described is arguably one of the more sophisticated forms of shared agency. At the very least, the burden of proof is on someone who denies this. Bratman understands that he risks circularity if he were to avail himself of Gibbard's particular form of acceptance, and holds out hope for a more individualistic form (16).

One question to raise, then, is whether Bratman can resist the considerations that move Gibbard toward a social conception of acceptance. Gibbard's view is shaped by an attempt to understand the difference between being in the grip of a norm and guidance by a norm. For example, there may be a norm that enjoins not telling someone what might upset them. Imagine that on an occasion the need to convey vital information is overriding, but that I cannot do so because I fear it would be upsetting to my audience. I am in the grip of the norm not to upset the addressee (Gibbard *op. cit.*, p. 58). Here I don't exhibit the sort of control characteristic of normative guidance.⁵ The control necessary for normative guidance must be something that can potentially be exercised from a detached perspective, some distance from the heat of the moment or the social encounter. Gibbard's conception of acceptance as tied to normative discourse would be one way to supply what's missing in this case. The question for Bratman is whether he might offer an alternative. Bratman's planning conception of intentions might seem like a possibility here, because this is precisely something that allows us to exercise some control over what we do in situations at some remove from the deliberative context. But recall that we're talking about the acceptance of norms of planning intention, and so we can't avail ourselves of this form of agency at this point.

A second concern with Bratman's invocation of the idea of norm acceptance is that it opens up the possibility that norms other than those of planning intention might be accepted, thus challenging the priority of planning agency over shared agency. Bratman says that "Built into the kind of construction I seek is that in the basic case the relevant norm acceptance is that of the individual participants, *and the norms accepted are, in the first instance, the rationality norms of individual planning agency*" (33 emphasis added; see also 16, n40). But why not accept interpersonal norms alongside of the individual planning norms? If one accepts a norm for consistency of one's intentions, couldn't one also accept a norm for consistency of intentions across participants in shared activity? Indeed, why not construct the individual planning norms on top of the interpersonal norms? After all, isn't this something like what happens in a competitive game?

It's not that Bratman doesn't recognize the possibility of accepting an interpersonal norm, not to mention the possibility that on some occasions, acceptance of a norm might itself be an instance of shared activity (see 26, 33, and

5 See Gibbard (*op. cit.*, 56) on normative control systems.

Chapter 7). It's a strength of his view over other forms of reductivism that it has resources to describe these different possibilities for agency, beyond the basic case where only planning agency is assumed and no interpersonal norms are invoked.

What, then, is the rationale for the fundamental place of planning agency in Bratman's scheme? The two considerations he offers in response to a closely related question are not conclusive. First, Bratman asks about the possibility that non-planning but nevertheless purposive agents (who engage in goal directed behavior but don't form intentions allowing for temporally extended agency) might in effect accept interpersonal norms (26). Bratman acknowledges this possibility, but responds that since we're *planning* agents it's of more interest for us to ask how planning agents might share agency. Bratman is right that we're planning agents. But we can admit this and wonder whether it's possible that planning agency might come along side of shared agency, rather than be something from which shared agency is constructed. This alternative (amongst others) is entirely compatible with our being planning agents.

The second reason offered for the primacy of planning agency is that appealing to the resources of planning agency rather than mere purposive agency makes for a richer model of shared agency, but in a way that is not radically discontinuous from planning theory. But if the distinctive norms of planning theory stem from their being accepted, it's no longer clear what sort of theoretical advantage is to be had by according primacy to planning agency and constructing shared agency from it. In what sense is the acceptance of interpersonal norms of shared agency somehow metaphysically or conceptually problematic as a starting point whereas acceptance of norms of planning theory is not? The acceptance of interpersonal norms is, arguably, on a par with the acceptance of planning norms; that is to say that both planning agency and shared agency are continuous with a form of agency that encompasses the possibility of acceptance.

5. Finally, return to the function of shared intention. Recall that Bratman thinks of shared intention to J as serving a social function – facilitating interpersonal coordination of action and planning toward J-ing, and structuring the related bargaining and shared deliberation of the participants. Bratman then seeks to offer a reductive set of conditions that are sufficient for playing this social coordinating role. But is Bratman right about the function that is supposed to be constitutive of shared intention? It's hard to deny that social coordination of intention and action is an important part of the function of shared intention. But there may be more to what shared intention is supposed to do, and this can make an important difference for what counts as shared intention.

What other function might there be for shared intention? I suggest that one important purpose or role for shared intention is to allow for individuals to have some power or control over what the group does. Shared intention might have a function of distributing the authority or power to settle practical issues facing a group. The alternative to Bratman's proposal, then, is that the function of shared intention extends beyond interpersonal coordination to include *power sharing* amongst the participants.

This is not to say that nothing like this issue emerges in Bratman's rich discussion. Part of Bratman's account of shared agency involves the idea that each participant has an intention of the form *I intend that we J*. Velleman has objected that this implies that intention entails discretion over the matter, and it's not clear how I could have discretion over our J-ing if you have this discretion as well.⁶ Bratman, Velleman, and others have given answers to this question.⁷ But my point is that Bratman's discussion here (64ff.) is concerned simply with showing how it's possible for one to have this intention while another has it as well. Bratman doesn't consider whether a fundamental role or purpose of shared intention might be to *share* power or discretion.

Thinking of power sharing as a function of shared intention will have a number of implications for how intentions are formed, revised, and transmitted amongst participants in shared activity. I only have space to consider implications for the sorts of scenarios that are compatible with shared activity. If the point of shared intention is power sharing, this constrains what counts as shared activity. Power imbalances take on a significance they would not otherwise have. Gilbert allows that some forms of joint action can occur even as a result of coercion.⁸ And Bratman (102) thinks that some significant power imbalances between participants in shared activity are possible. For example, he thinks that there may be shared activity between POW's who are forced to join with their captors in building a bridge – so long as the background coercion does not interfere with the coherent meshing of intentions/plans that he thinks is crucial for shared intention. But if the point of shared agency is not merely coordination but power sharing, then it's not obvious that the POWs share agency with their captors. And that, I think, reflects our pre-theoretic intuitions about the case.

6 J. David Velleman, "How to Share an Intention." *Philosophy and Phenomenological Research* 57, 1990.

7 My own attempt is in "Shared Agency and Contralateral Commitments," *Philosophical Review* 113, 2004.

8 Margaret Gilbert, "Shared Intention and Personal Intention" *Philosophical Studies* 144, 2009. Gilbert thinks that this shows that the special obligations between participants are fundamentally different from moral obligations.

In closing, Bratman's account of shared agency doesn't quite capture practical intersubjectivity. But the central importance of Bratman's discussion for the philosophy of action and related areas is not in doubt. And the articulation of shared intention in functional terms raises interesting new avenues for theorizing about shared agency and collective intentionality more generally.

Acknowledgements: Thanks especially to Michael Bratman for generous conversation. Thanks also to Frank Hindriks and Kirk Ludwig.