

Abraham Sesshu Roth  
Department of Philosophy  
Ohio State University

### **Prediction, Authority, and Entitlement in Shared Activity<sup>1</sup>**

Abstract: Shared activity is often simply willed into existence by individuals. This poses a problem. Philosophical reflection suggests that shared activity involves a distinctive, interlocking structure of intentions. But it is not obvious how one can form the intention necessary for shared activity without settling what fellow participants will do and thereby compromising their agency and autonomy. One response to this problem suggests that an individual can have the requisite intention if she makes the appropriate predictions about fellow participants. I argue that implementing this predictive strategy risks derailing practical reasoning and preventing one from forming the intention. My alternative proposal for reconciling shared activity with autonomy appeals to the idea of acting directly on another's intention. In particular, I appeal to the entitlement one sometimes has to another's practical judgment, and the corresponding authority the other sometimes has to settle what one is to do.

---

<sup>1</sup> Thanks especially to Michael Bratman, Seana Shiffrin, Lisa Downing, Tyler Burge, Dan Farrell, and Barbara Herman. Thanks also to Sarah Stroud, Dave Hilbert, Marya Schectman, Justin D'Arms, Don Hubin, Margaret Gilbert, Guy Rohrbach, Eric Marcus, Christine Korsgaard, Luca Ferrero, Kirk Ludwig, Richard Samuels, Gary Ebbs, an anonymous referee at *Noûs*, Christine Tappolet, Sara Chant, Zachary Ernst, Jon Garthoff, Kyla Ebels-Dugan, and Peter Vallentyne. I have also benefited from presenting different parts or ancestors of this paper to audiences at Pacific and Central APA symposia on shared agency, a Northwestern University Society for Ethical Theory and Political Philosophy Conference, an Auburn Philosophy Department colloquium, a meeting of the Centre de Recherche en Éthique de l'Université de Montréal, a Kline Chair Workshop on shared agency at the University of Missouri, and seminars at Ohio State University.

Philosophical reflection suggests that shared activity involves a distinctive interlocking structure of intentions, wherein each individual participant intends and is committed to the joint activity. On Michael Bratman's view, a participant in shared activity J has an intention of the form *I intend that we J*. And John Searle holds that each individual has an intention of the form *we intend to do such-and-such*.<sup>2</sup>

But how is such a structure of intentions established? Well, how does shared activity get started? Sometimes it's like this: you begin by doing something on your own, and end up acting *with* another, but there was never a moment when you *decided* to act with anyone, let alone this person. You just fell into it or grew into it, and it might even come as a revelation that you're not acting on your own. Here the structure of intentions is not the direct outcome of any exercise of agency on the part of the participants. Very often, however, individuals do get together and decide to act jointly. That is, shared activity often is simply and knowingly willed into existence by the individuals who will participate in it. In this sort of case, the structure of intentions is established because each individual frames the intention that is his or her contribution to it.

The problem is that it is not obvious how each participant can do this. If the standard understanding of intention is to serve as our guide, the intention (philosophers contend) I form as my contribution to shared activity apparently settles not only my own actions, but those of my partners as well. But do I have the power or authority to determine what others will do? Is such an intention compatible with respecting the autonomy of fellow participants? Can this intention allow the activity genuinely to be shared, or does it require a problematic element of passivity or subordination in others?

---

<sup>2</sup> Bratman (1992, 1993), Searle (1990: 401,7). Gilbert's notion of joint commitment seems also to involve an interrelated structure of intentions, although she does not think it is reducible to this structure. See Gilbert 2003, 2009. Velleman (1997a), following an earlier suggestion of Gilbert (1990) (a suggestion she has since disavowed), appeals to a structure of interrelated conditional intentions. See below.

The issues here have not gone unnoticed.<sup>3</sup> But the problem is more difficult to solve than has been appreciated. I will focus on a sophisticated response that would have one *predict* rather than determine what one's partners will do. If I can predict what you will do, then I will be able to form intentions about what we will do without pretending to supervise or control you, rendering moot the concerns about authority and autonomy.<sup>4</sup> This Predictive Strategy does not work because, I will argue, it gets wrong the relationship between practical and theoretical reasoning. I will suggest that the perspective from which one engages in practical reasoning will constrain one's theoretical perspective and to an extent preclude taking a theoretical perspective on what one will do. The implication for the Predictive Strategy, I will argue, is that the prediction meant to secure the possibility of practical reasoning leading to the requisite intention undermines that reasoning instead. Thus, I hope to show that some insights concerning individual agency, specifically about the nature of commitment and the relation between practical and theoretical reasoning, have implications for discussions of shared agency. I will also take some initial steps exploring how the authority required for intention formation might be compatible with activity that's genuinely shared.

## 1. Participatory commitment and the Intention Thesis

Suppose that several individuals undertake some shared or joint activity, call it *J*. Theorists who differ on a host of issues nevertheless agree that a satisfactory description of the situation will involve each participant having an intention concerning *J*. Call this the *Intention Thesis*.

---

<sup>3</sup> David Velleman (1997a) has raised an influential form of the problem. See also Baier (1996), Bratman (1997), Schmid (2008).

<sup>4</sup> The Predictive Strategy is inspired by but not identical to a view defended by Bratman (1997). See below.

What motivates the Intention Thesis? One argument starts with the idea that intention involves a sort of commitment, and that this is reflected in the rational norms or requirements associated with intending. These norms distinguish an intention to A from a mere desire to A.<sup>5</sup> I will assume that this is correct. Now, the Intention Thesis says that a participant in some shared activity J has an intention with respect to J. It would follow that each participant in shared activity J is *committed* to it. This seems to be just what we should say about shared activity. The Intention Thesis would thus yield an account of an important feature of shared activity – that of *participatory commitment*.

But why think that each participant in shared activity is committed to it? There are a number of reasons. First, shared activity typically involves a significant amount of coordination, where each participant relies on each of the others to do his or her share of the activity. This reliance would not make much sense unless each individual can recognize others as at least to some extent committed to the undertaking.<sup>6</sup>

Second, some theorists recognize a special obligation participants in shared activity have to one another.<sup>7</sup> But it's hard to see what this obligation could require, if it did not at least require one to do what one reasonably can to see the activity through. Implicit in the mutual obligations between participants, then, is some sort of commitment to the shared activity.

Shared activity also involves rational constraints of consistency and coherence between participants. If you and I are J-ing together, then my J-related intentions are subject to demands of consistency and coherence with your J-related intentions. It's less clear that the same sorts of interpersonal rational requirements hold between our other intentions, or between our intentions and those of a non-participating third party. What would account

---

<sup>5</sup> Bratman (1987), Harman (1976, 1986).

<sup>6</sup> Bratman (1992).

<sup>7</sup> Gilbert (1990, 1996), Roth (2004).

for the special status your J-related intentions have for me that your other intentions, or those of a third party, do not? The suggestion is my commitment to *our* J-ing, since our J-ing depends in part on your J-related intentions.<sup>8</sup>

So a number of considerations suggest that participatory commitments are essential for shared activity. And these commitments, in turn, suggest the Intention Thesis – that is, an understanding of shared activity that attributes to each participant an intention concerning the activity.

The Intention Thesis might also be motivated more directly. Attributing intentions to each participant serves to distinguish shared activity from other social or collective phenomena. In shared activity, we think of the individual participants as exercising their agency so as to do one thing or another *together*; acting together is something aimed at, even if not always an end in itself. The contrast is with certain collective or social phenomena where it is theoretically or explanatorily useful to describe a collection of individuals as a single agent or entity, or if not that, to describe events associated with a collection of individuals in terms of action, activity, or intention. But this might be useful for reasons that have nothing to do with the goals of the constituent individuals. For example, the interactions of various individuals might result in some prevailing market price, and we might speak of the market (the suppliers and consumers) setting the price, but no one aims for the market to have such a price. Whereas, in the social phenomenon that concerns us here – shared agency – part of what individual participants are aiming at is doing something together. And it is natural to understand this aiming in terms of the participant's intention with respect to the activity.

This connects with something noted at the outset – that shared activity can be brought about through the participants simply deciding to act together. But if decision is a matter of forming an intention, then each participant must have intentions with respect to the

---

<sup>8</sup> See Bratman (1992) on meshing subplans, and Bratman (ms), Ch. 2.; see also Roth (2003).

activity. This suggests that the Intention Thesis is correct, at least for this important class of decision-based shared activity.

## 2. The problem, and some responses

Although the Intention Thesis is crucial for understanding shared activity, it has been noted that there are aspects to intention that pose a challenge for invoking it in a theory of shared activity. Intending is something I do to settle a deliberative issue: weighing several options, I decide on A-ing, and thereby intend to A. My intentions provide fixed points to structure subsequent practical reasoning and action: what means to A-ing should I take? What else must I intend to secure those means? How may I pursue other ends in ways that are compatible with A-ing?<sup>9</sup> This suggests the Settling Condition<sup>10</sup> that I can only intend what I take to be up to me to decide or settle. It is a violation of a rational requirement to intend something I don't think I can settle, and thus regard my ensuing plans and actions as likely coming to grief. Applying the point to collective action, to say that I intend for us to be dining together presumes that whether we're dining together is something for *me* to settle. But the idea behind *shared* activity and intention is precisely that it's *not* entirely up to me what we do. You have a say in the matter; at the very least what *you* do should be up to you. Our problem, then, is that shared activity would seem both to demand and to disallow one and the same intention on the part of each participant.

---

<sup>9</sup> Harman (1986); Bratman (1987).

<sup>10</sup> Following Bratman (1997), who calls it the "Settle Condition." For some, the Settling Condition implies belief: the sense in which I'm settled on A-ing when I intend to A entails that I believe that I will A (Harman 1976, Velleman 1989). Even if one doesn't have such a strong belief requirement on intention, there remains the worry of how to reconcile the respect one should have for the agency of fellow participants with the idea that one's intention settles the matter. Thus, Bratman takes this problem seriously, even though he doesn't endorse the strong belief condition on intention.

The general issue here might be articulated in more than one way; or perhaps we might say that there is a family of problems here. We might put the concern more **objectively**: by intending that we J, I'm actually settling what we all do. This suggests that I have a causal power over what we all do, one that deprives other participants of the ability to settle what they do. A version of the objective concern might be put in more normative terms: in intending that we J and settling what we all do, I have and exercise an authority over others that is incompatible with the activity being shared. The settling concern might be put in more **subjective** terms as well, where this also has causal and normative versions. In intending that we J, I must take myself as settling what we all do; but how could I count as being engaged in shared activity if I believe or take myself to be exercising this causal power or authority over other participants? Although I would want to allow for the possibility of acting with others even if one doesn't quite realize it, it does seem that merely thinking that one is exercising this power or authority, let alone actually doing so, does not sit well with genuinely shared agency.<sup>11</sup>

One response to this settling problem (or family of settling problems) is to adopt a version of the Intention Thesis that appeals to conditional intentions.<sup>12</sup> Each participant's intention regarding the activity (or his contribution thereto) is explicitly conditioned on the other participants intending likewise. The suggestion is that each participant has the intention of the form

C1 I intend that we J, so long as you intend likewise.

---

<sup>11</sup> I suspect that the version of the problem that most animates Velleman (1997a) is objective and causal. But the version involving authority also seems to be relevant. In any case, the debt to Velleman here should be obvious. Bratman's discussion here is also important. See sections 5 & 6 in Bratman (1997).

<sup>12</sup> The idea is most fully articulated by Velleman (1997a). Velleman sees himself as developing a position held by Gilbert when she speaks of a *pooling of wills*. Indeed, Gilbert also seems to talk about interdependent conditional intentions (as Velleman recognizes); see her (1990: 8). In any case, Gilbert disavows the idea of conditional intentions in more recent work, e.g. (2003: 51-2). For discussion, see Roth (2004).

Unlike a categorical intention, a conditional intention doesn't presuppose that the intender can by himself settle what the others will do. The matter of our J-ing is settled, but only with the satisfaction of a condition outside the scope of my agency, involving the intentions of others. When every participant has the appropriate conditional intention, the condition in each intention is satisfied, an intention is shared, and shared activity undertaken.

I don't think that this will work.<sup>13</sup> Interdependent conditional intentions fail to settle anything, not even conditionally, and so are not properly speaking intentions at all. If my intention to go for a walk is conditioned on yours, and yours is likewise conditioned on mine, then my intention is, in effect, conditioned on itself: I intend to go for a walk if I so intend. But then the question of whether to go for a walk hasn't yet been answered: do I or do I not intend to go for a walk? As it stands, an intention that's conditioned on itself, whether directly or through interdependence with those of another, doesn't settle anything, and commits me to nothing. The structure of interdependent conditional intentions regarding some activity A is as compatible with our doing nothing as it is with our A-ing. Thus, it is not at all clear why such interdependent conditional intentions lead to any activity getting under way, nor is it clear that they could capture any sort of participatory commitment of the agents involved.<sup>14</sup>

One response to the above criticism appeals to a revised understanding of the conditional intentions involved in the proposal.<sup>15</sup> For the case of walking together, the intentions are of the following form:

C2 I intend that (we walk, so long as you intend likewise)

---

<sup>13</sup> Roth (2004).

<sup>14</sup> It does not suffice to add that each individual has a desire to engage in the activity. A desire to A does not embody a commitment to A.

<sup>15</sup> I would like to thank an anonymous referee for pressing me on this matter.



The parentheses indicate that the condition lies within the content of the attitude. This was not explicit in the initial formulation. The thought here is that this intention could be had independently of whether the other person has the corresponding intention.<sup>16</sup> So if we each have the intention independently of the other, I can't so to speak get rid of your intention simply by getting rid of mine. Thus, given that the condition specified in the content of my intention is or will be satisfied – i.e., you maintain your corresponding conditional intention – then the only way for me to live up to the rational constraint imposed by my intention is to go ahead and act in ways appropriate to our walking together. So it seems that on the revised understanding, the conditional intentions do retain the element of commitment to the shared activity that we were hoping to capture.

But this solution comes with costs. Recall that the intention is meant to capture the sense in which each participant is committed to the activity. But these participatory commitments, and hence the corresponding intentions, are usually not independent but *interdependent*. I wouldn't have an intention concerning our walk unless you also had such an intention.<sup>17</sup> This interdependency might be for different reasons: perhaps it's not possible to perform the action in question unless the other does so as well, or perhaps there is no point to doing so if the other doesn't. Indeed, I think that recognizing this interdependency is the insight that motivates the conditional intention approach in the first place.<sup>18</sup>

---

<sup>16</sup> I set aside the worry that the intentions on this proposal may be semantically ungrounded. Though I don't share his optimism in thinking that this worry has been adequately addressed, Velleman (1998, 45 n26) has a nice discussion of the problem. Kit Fine has also expressed skepticism about Velleman's solution in "Response Dependence and Joint Intention" presented at the Collective Intentionality VIII Conference in Manchester (2012).

<sup>17</sup> Gilbert 2009, 180. Of course, things wouldn't get started unless the one who in fact starts were entitled to think that the others would join in. So the onset of the intentions needn't be interdependent. But at the very least, *persistence* in the commitments (and intentions) are interdependent. (Bratman, ms 79ff; Bratman 2009, 159)

<sup>18</sup> In his review of Bratman, Velleman (2001, 120) holds that the intention remains conditional even as its conditions are satisfied. Whereas, it is more natural to think that the intention becomes categorical when the

There may be some cases where the participatory commitments are not interdependent. But this is not typical. Thus, even if the form of the conditional intentions is as in C2 above, if for each participant such an intention embodies his or her commitment to the activity, then in most cases of shared activity these intentions would be interdependent, and this would be common knowledge. In such cases of interdependence, our problem of accounting for participatory commitment returns. Let me explain.

Normally, the satisfaction of the conditions for conditional intentions are not within the scope of the agent's will, at least as it is exercised on that occasion. For example,

I intend (to go to the beach, so long as it's sunny this afternoon).

I have no control over whether it's sunny this afternoon, so if it is sunny, then the only way for me to satisfy the constraint imposed by the intention is to go to the beach. There is a point to framing such an intention; I can settle a matter, committing myself to the extent that I can given the cooperation of the world, and move on to thinking about other things. In contrast, consider

I intend (to give you \$1000 dollars, so long as I don't snap my fingers in the next minute).

One way to live up to the rational constraint imposed by having this intention is to pay up. But I could just as well snap my fingers and not pay (I'm assuming here that it's completely up to me whether I snap my fingers in the next minute). Given that there is such an easy way for me to get out of paying, it's not clear what the point of framing this intention is.

---

condition is satisfied. Why insist on maintaining the condition? The answer, I think, is that the intentions being conditional is, on Velleman's proposal, crucial for capturing the interdependence of the commitments.

Likewise, there are not one but two ways for me to live up to the rational constraint that is imposed by the intention in C2. First, I could go ahead and act in ways appropriate for our walking together. Alternatively, I could see to it that the condition in question is not (or is no longer) satisfied, in which case the intention so conditioned no longer constrains me. How might I accomplish this? Given interdependence, I could simply drop my intention. Then you would not maintain the intention, and so the condition on my intention is not satisfied. Thus, dropping the intention triggers the condition that releases me from having to act. But can I just drop an intention? Dropping an intention for no reason would, normally, amount to a violation on a rational constraint on intending that reflects the commitment embodied in intention. However, dropping the intention in this case is permitted precisely because the condition on the intention is no longer satisfied. This is why conditional intentions are not meaningfully conditioned on something over which one has such straightforward and immediate control. So dropping the intention is a legitimate option in this case. I conclude that the intention in C2 does not really embody any commitment to the activity.<sup>19</sup>

I should note that another attempt to reconcile the Intention Thesis with the Settling Condition on intentions invokes the intention *merely to do one's part in our J-ing*, instead of the intention that we J.<sup>20</sup> A commitment to one's part in our J-ing is not quite a participatory commitment to our J-ing as a whole. But might it be enough for the purposes of coordination and reliance? To see why not, consider walking together.<sup>21</sup> We might describe my part as walking at a certain pace. But intending to do that is entirely

---

<sup>19</sup> Notice that in general, there is no point in retaining the conditional intention if one knows that the condition won't be satisfied (and I think this is true even for intentions of form C2). Thus, if I get rid of my conditional intention, and you know this (we're assuming common knowledge of each other's attitudes), then you would not retain your conditional intention. Thus, we are able to get the same result without even appealing to any interdependence of intentions peculiar to shared activity.

<sup>20</sup> Intentions of this form figure in many papers by Tuomela. For a recent statement, see Tuomela (2006), 39. See also Kutz (2000).

<sup>21</sup> Gilbert (1990).

compatible with undermining your contribution, for example by intentionally tripping you. Such an intention would not reflect a commitment to the shared activity. Suppose instead that we avail ourselves of some robust conception of *part*, so that each participant intends to do his part in shared activity, *as such*. This would appear to rule out attempts to undermine a partner's contribution. Unfortunately, it's not all that clear that I can actually will my part so understood, if I don't regard your contribution as forthcoming. My intention would seem to require that your contribution and our J-ing be a settled matter; but when did we settle it? How? This is what we were trying to account for in the first place.<sup>22</sup> So it seems that if some robust conception of one's part is playing a role to capture a commitment to shared activity, then it's not yet clear that we can intend it; and if 'part' is not understood robustly, then we may very well intend it, but we haven't captured participatory commitment. More could be said about this strategy.<sup>23</sup> But in this paper I'd like to focus my critical remarks on yet another approach.

### 3. The Predictive Strategy

If my intention that we J doesn't settle the matter for you, how could I be in a position to form it? One important proposal is that I can *predict* that you will have the requisite intention.<sup>24</sup> Many conditions or circumstances are not under my control. But given that it's reasonable for me to believe with some confidence that they'll be in place, I can go ahead and settle some practical matters by forming intentions that presuppose these circumstances. For example, my intention to get a tan by sunning myself at the beach this afternoon requires that it not be cloudy. I don't have any control over the weather, but so

---

<sup>22</sup> For the related worry about circularity, see Kutz (2000), and Petersson (2007).

<sup>23</sup> On an alternative strategy for characterizing the intention to do one's part, see Gold & Sugden (2007), which I discuss in Roth (forthcoming).

<sup>24</sup> This view is inspired by Bratman (1997). Bratman's account is subtler and stronger; it evades the criticism I lodge against the simple Predictive Strategy. Although I don't refute Bratman's more sophisticated version, discussion of the simpler view will help to articulate my concerns with his view.

long as I can reasonably predict that it'll be sunny, I can go ahead and have this intention. Now, our J-ing requires certain intentions on your part. But if I can *predict* with some confidence that you will have the appropriate intentions, I can go ahead and form the intention that we J.

This proposal is attractive because the appeal to prediction allows one to represent the intentions and actions of another in the content of one's intention, but in a way that avoids requiring that one think of oneself as deciding or settling what another will do. It doesn't follow from my being able to predict what you do that you cannot form your intention through practical reasoning in the usual way. There seems to be nothing inherently coercive or otherwise problematic about making predictions about what others will do, and then going on to form intentions or plans in light of those predictions.<sup>25</sup> So it appears that we have an explanation for how it is that I, as a participant in shared activity, can form the intention that we J without settling in the relevant sense what my partners will do. One need not appeal to any special authority one has over the other in order to do this, just as one needn't have any authority to settle all sorts of conditions that are presupposed for the success of one's intentions, such as the weather not being inclement.

But might there be something worrying about regarding a fellow participant in this predictive mode? It need not be coercive, but it doesn't seem to be an attitude toward individuals that is consistent with seeing them as partners or prospective partners in shared activity.<sup>26</sup> As a preliminary illustration, and to get some of the issues on the table, suppose that you and I are both concerned about some social or political matter. I know this about

---

<sup>25</sup> Bratman (1997).

<sup>26</sup> I'm willing to admit that there's *some* sense in which shared activity can occur even when one individual is coercing another. But surely it is not fully shared; to some extent the sharing is compromised. Compare the case where I form an intention to A later, and I subsequently A, but only because I've rigged the circumstances to ensure that the costs of not following through are prohibitive (e.g. I hire a thug to beat me up if I don't keep my appointment). Then I think that in some sense my agency is compromised. An individual who regularly relies on such measures is not exercising genuine diachronic agency.

you, but I figure that left to your own devices, you wouldn't ever get around to joining a demonstration downtown on just this issue. But I predict that if we have lunch downtown, and schedule a movie screening a couple hours later, then we'll have a block of time in between that won't suffice to get any meaningful work done but would be just enough to join the demonstration nearby. We agree (without coercion or manipulation) to have lunch and see a movie. Having lunched, we kick around downtown and “come across” the protest. I had predicted that you would join in, and it is on the basis of this prediction that I was in a position earlier to intend that we join the protest. But such an intention does not seem to represent a participatory commitment to shared activity.<sup>27</sup> We do not here jointly attend a protest – or if we do, the participatory commitment involved is not captured by *this* intention. The Predictive Strategy I employ in so intending subjects you to opportunistic use for my purposes. However benign or benevolent those purposes may be, my attitude would seem to preclude this from being an instance of shared activity.<sup>28</sup>

Raising this worry about the predictive view does not commit one to thinking that what I'm doing in the example is *morally* problematic. For example, Korsgaard worries that if the predictive view is correct, then each individual would be intending to treat fellow participants as means to the end of our J-ing.<sup>29</sup> Be that as it may, the point of the current

---

<sup>27</sup> But is taking the predictive attitude what's problematic here? One suggestion for an alternative explanation is that I have a dubious ulterior motive in a number of my interactions with you leading up to the activity in question. But in this case the motive itself might be impeccable. I just want to get you to attend the protest with me, something that you wouldn't have any objection to; it's just that you'd never get around to it on your own. Perhaps what's problematic is that I am not entirely forthcoming about why it is that I'm so interested in having lunch and seeing a movie with you. But just how forthcoming about one's motivations must one when acting with others? If that's why this case is problematic, many cases that we'd like to count as shared activity wouldn't. See Bratman (1999, 122).

<sup>28</sup> I concede that if having joined the protest you continue shouting and marching up and down side by side with me, despite there being real opportunities for you to withdraw, it would be hard to deny that eventually we come to be engaged in genuine shared activity. I don't think that acknowledging this affects the argument.

<sup>29</sup> Korsgaard raised this concern with Bratman in discussion at the Collective Intentionality VII conference in Basel in 2010. See Bratman (ms) for discussion and response.

argument doesn't rely on it being a moral requirement not to treat others as means. Rather, it is simply that my intention, based as it was on my predictions about what you will do, doesn't adequately reflect your contribution to what we're doing, and it puts in doubt that our going to and *joining* the protest is shared activity. That is, treating someone as a means, whether or not it is moral, is not genuinely to act *with* them.

All this is compatible with the thought that one can in moments of reflection step away from participation, and the attendant shared practical perspective, in order to take on a more individual, predictive stance with respect to fellow participants. Indeed, an analogous intellectual maneuver is possible with respect to one's own actions. But the point is that this is not, fundamentally, how the agent looks upon what she is doing; taken too far, it's symptomatic of a self-alienation that undermines genuine agency.<sup>30</sup> The predictive stance would seem to occupy a similarly precarious position with respect to shared activity.

A *preliminary* worry for the Predictive Strategy, then, is this. On the one hand, shared activity involves individual participants with intentions for example of the form *I intend that we J*. To have such an intention, I need to be able to regard my intention as settling the matter. On the strategy outlined just above, this involves my making a prediction about your intentions. But on the other hand, the predictive attitude appears to be the wrong perspective to have on fellow participants. In acting together, you and I are sharing a practical perspective. I don't enter into my own practical or agential point of view fundamentally as an object to be predicted. Likewise, when I take on this shared perspective, you shouldn't figure in it in a way that a third party very well may – as fundamentally subject to prediction.

#### 4. Intention and prediction

---

<sup>30</sup> See the next section.

Recall how we got to the view that would have us regard fellow participants in the predictive mode. The original problem with shared agency concerned how it was that I could take myself as settling what you will do while regarding what you do as up to you. The Predictive Strategy is meant to solve this problem. I don't have to settle what you do; I can predict what you do instead. I've just suggested that in the context of shared activity, there may be some concerns with my taking a predictive attitude with respect to your intentions and actions.

But surely these concerns pale next to those stemming from my *determining or settling* what you are to do. Not everyone has the intuition that shared activity requires each individual to have the *respect* that would preclude adopting a predictive attitude regarding fellow participants; for some it hints of preciosity. If dropping the respect requirement is the price for solving the initial problem, many would be happy to pay it. For all that's been said, shared activity might differ from individual agency in being more hospitable toward taking the predictive stance – at least towards one's partners if not towards oneself. So we're at a bit of a stalemate, and I do not want to rest too much on the claim about respect or lack thereof in my criticism of the Predictive Strategy. Let's think a little more about this Predictive Strategy to try to see what else, if anything, might be wrong with it.

The proposal we're considering is that for me to intend that we J, I must be able to predict that you will also have an intention with a similar form. But that cannot be all. If I am to rely on your intention, or the prediction of your intention, as a basis for forming mine, then I have to think that your intention is *effective* and can be relied upon. I can't think that your intention is a passing fancy, somehow misguided or confused, or based on false or unfounded presuppositions. It would, for example, be problematic for me to plan around your intention to go shopping when I believe that your car is broken and that you have no other means of transportation (unless, of course, I'm planning around your *failure*). Similarly, without some warranted confidence in the effectiveness of your intention



regarding our J-ing, I would not be able to rely on it in shared activity; I would be unable reasonably to intend and commit to our activity.

If I have to assess the effectiveness of your intention, I will need to be confident in its presuppositions. One presupposition of your intention that we J is that I also so intend. Our J-ing requires my contribution after all, and in shared activity that contribution is not forthcoming unless I have the right intention. So, in assessing your intention as reliable, I have to think that your prediction or presupposition of *my* intention was correct. Thus, in order to form my intention, I would be predicting it.

I've been arguing that (1) the Predictive Strategy entails self-prediction, i.e. the prediction of one's own intentions in order to arrive at them. My strategy for criticizing the Predictive Strategy will be to argue, further, that (2) this self-prediction is problematic. Before I do this, I want to address several worries that may be raised about the argument just given for thinking that the Predictive Strategy entails self-prediction.

(a) The argument for (1) presupposes that in order to assess the reliability of your intention, I need to assess whether a condition upon which the reliability of your intention rests – namely, the condition that I will form the appropriate intention. The first worry is that this assumption (that I have to assess whether I will form the intention) is problematic in that it relies on the following implausible general principle: For any condition C, if the reliability or effectiveness of an intention rests on C, properly assessing whether the intention is reliable will require assessing whether C is in place.<sup>31</sup> I agree that properly assessing the reliability of an intention cannot require that you check *every* such condition; that would set an impossibly high standard, and we would never be able to assess the effectiveness of any intention. So the argument for (1) is flawed if it were to rely on such a principle. But I don't think it does. The argument relies instead on the idea that assessing the reliability or effectiveness of an intention requires that I investigate whether some

---

<sup>31</sup> Thanks to an anonymous referee for *Noûs* for raising this issue.

salient or relevant conditions for its success are in place. And I submit that in many cases of shared activity, whether one will oneself form the requisite intention is a salient condition for the reliability or effectiveness of the other's intention. So I think that the Predictive Strategy does entail self-prediction, and that the entailment does not depend on the aforementioned implausible general principle.

(b) The argument for (1) assumes that assessing the reliability or effectiveness of your intention requires me to assess whether your intention is grounded on warranted beliefs about conditions for success, in particular, the warranted belief that I will intend. The second worry is that it doesn't matter whether your belief or presupposition that I will intend was warranted or not. I can, the suggestion goes, *make your presupposition correct*, by intending as you think I will. Thus, simply going ahead and intending will ensure that I will be able to rely on your intention, which is a condition for that very intending of mine. All this, without having to predict my own intention. The problem with this suggestion is that it is not an implementation of the Predictive Strategy, which is supposed to make use of a prediction – a sort of belief – in order to justify the forming of an intention. But no reason for believing in or predicting the effectiveness of the other's intention is given. An attitude that doesn't meet minimal standards of rational prediction – indeed is formed typically as a result of a mere wish to do something with another – is not properly a prediction.<sup>32</sup>

(c) A further worry with the argument for (1) stems from the idea that the Predictive Strategy isn't necessarily committed to the view that one explicitly makes a prediction every time one forms the intention that we J. On this construal, the Predictive Strategy is engaged in some rational reconstruction: it tries to make sense of intention in this context by showing how the rational norms of intending *would* be satisfied were the relevant prediction to be made, and this makes for the rational legitimacy of the intention even though the prediction isn't actually made. If this is right, then the Predictive Strategy

---

<sup>32</sup> A similar point is made by Grice (1971) against a proposal due to Hampshire and Hart (1958).

wouldn't necessarily result in self-prediction. In reply, I suggest that successful rationalization will include at least implicitly a further prediction concerning one's own intention. But we'll see shortly that this self-prediction in this context is rationally inadmissible, and therefore that the prediction of the other's intention is not warranted. A rational reconstruction that invokes (even if implicitly) rationally inadmissible or unwarranted predictions does not succeed as a rationalization. So even if the Predictive Strategy does not generate actual self-prediction, self-prediction nevertheless poses a problem for the sort of rationalization being offered by the Predictive Strategy.

In sum, the Predictive Strategy entails self-prediction because it tries to show how the intention that we J satisfies a rational requirement on intention – the Settling Condition. Given the concern with rationality, it would be a problem if the prediction in question were never formed in a minimally rational fashion. It's therefore incumbent on the Predictive Strategy to address the rationality of the belief concerning the effectiveness of the other's intention. And given the interdependence of intentions, that brings the agent to the point of having, rationally, to predict his or her own intention.

So much for the defense of (1), the claim that the Predictive Strategy entails self-prediction. Now let us turn to (2) and ask, what's so bad about predicting one's intentions? As was noted above, I can on occasion think about myself from a detached, third-personal perspective, and come to conclude that I will intend or do so-and-so because of my character and such-and-such circumstances. But the troublesome aspect of the present case is that the detached perspective enters right into – and, I contend, undermines – the practical reasoning that's supposed to lead me to act. I'm suggesting that there is a tension between predicting one's intention and forming that intention through practical reasoning that relies on that prediction. The point is most perspicuous in the case of forming the intention *right now*, where every participant forming or having formed the intention by now is required for our J-ing: it's now or never. This requires that I have to believe, *now*, that you intend likewise. But in coming to believe, in these circumstances, not only that you

intend that we J but also that your intention is effective, I will have *discovered* that I already have the intention (for how could your intention be effective if I don't also intend?). No reason, then, to *form* the intention through practical reasoning.

Not all cases are like this. Your intention (or the intention you will acquire) might be premised on the *prediction* that I will intend, not on my already having the intention. So the only theoretical conclusion that I would run in to in assessing the reliability or effectiveness of your intention is that I will intend sometime in the future. And it's not yet clear that this is problematic. At least, it doesn't seem as bad a result as the one where an attempt to reason toward intending leads me to conclude that I already so intend. If I only predict that I will intend, could I not then go ahead and form the intention at my leisure?

Well, what would this amount to? One forms intentions for the future in order to settle some practical matter. But, given my prediction, it seems that there is nothing about the world for me to settle. Or, if it is not something in the world but my thought that is to be settled, I've already done so by predicting. So what would be the point of intending this conclusion? Intending (or deciding, for that matter) is not supposed to be an endorsement of one's fate or destiny. From one's point of view, the intention is to be taken as settling the matter. But in reasoning toward the intention, one realizes that the matter is already settled, in that one has predicted the intention.

Goldman has argued that nothing precludes one from having inductive knowledge of one's future acts – and, presumably, intentions – knowledge that is not based on intending or deciding on those acts. He criticizes Hampshire, who insists that knowledge of one's own voluntary actions can only be practical, based on one's firm intention or decision to act.<sup>33</sup> One of Goldman's counterexamples to Hampshire's thesis involves the agent forgetting the

---

<sup>33</sup> Goldman (1970, 192-4). He cites Hampshire (1965, 54). See also Hampshire and Hart (1958, 5) for something close to the claim that one cannot have predictive/evidence-based knowledge of voluntary action (assuming that a voluntary action is one based on a decision).

prediction and subsequently intending and acting as predicted. In another case, the agent starts with the prediction of an action – in this case, suicide – but then gradually over the course of five years acquiesces to the idea, so that in the end he possesses the intention to kill himself.

Perhaps Goldman's argument is effective against Hampshire's thesis as it was formulated.<sup>34</sup> His counterexamples nevertheless betray *some* sort of incompatibility between the practical and theoretical perspectives, a point seemingly conceded in Goldman (1970, 194). In neither instance is the practical matter settled through any act of will or decision undertaken in the context of a prediction regarding that very matter. In the first case, the prediction is screened off from having any impact on decision because it is simply forgotten or not attended to. In the second, there is no decision; one's will is structured passively, over an extended period of time. At no point in this example does one decide on suicide.<sup>35</sup>

One factor underlying the incompatibility of prediction and decision/intention is that the former undermines the point or motivation for the latter. If the issue of what I will do is already settled for me, then what purpose is served for me to *decide* so to act? We can make the point by assuming that forming an intention is an instrumental act, a means we take toward an end. In this case, the act would be without purpose, since the matter is

---

<sup>34</sup> It is less clear to me that the formulations in Hampshire and Hart (1958) are susceptible to the counterexamples.

<sup>35</sup> The discovery or prediction of one's intention to A would, of course, leave some room for practical thought. For example, one might yet opt for and settle on certain means rather than others, or decide to realize or execute the intention in some specific way. Perhaps this is a way to show how prediction needn't undermine intention. But notice that in the case of shared activity, I cannot *by myself* settle how we all carry out an intention. If I were again to avail myself of the Predictive Strategy, I would then have to predict not only your intention that we J, but also your intention that we J by our K-ing. But then I'd be making a prediction regarding our intention to K, which I cannot do; the same problem of predicting one's intention has recurred at the level of means. What then is left for me to intend? Perhaps something like that we J by *my* K-ing. But there is no reason to think that my K-ing is a legitimate way for me to contribute to our J-ing. That would depend on what other people would do, which in turn depends on whether I'm K-ing. And the problem emerges yet again.

already settled. But even if we don't think of intention formation as an act, the point remains that it would be an attitude without a function or role to play in reasoning. For that role would already be played by the prediction.

A related consideration has to do with the deliberation that leads to intention. The mental activity of recognizing, assessing, and weighing reasons cannot count as deliberation or practical reasoning if one doesn't regard the activity and the decision/intention that is its outcome as something that will settle what one will do. If on the basis of evidence or theoretical considerations one already regards it as settled that one will A, it's hard not to regard the consideration of reasons for or against A-ing as idle, and not genuine deliberation.<sup>36</sup>

So, the suggestion is that the Predictive Strategy undermines the point of intending. Part of the practical reasoning necessary for intention involves an assessment of what is in one's power or capacity to intend. Indeed, genuine deliberation and reflection on the reasons for acting will in some way encompass the issue of whether one can A, for one doesn't have the relevant sort of reason to A if one cannot A. We saw above that when I reason toward the intention that we J, I will need to establish whether I can so intend, and this would lead to predicting my own intention. In attempting to engage in practical reasoning in the manner the Predictive Strategy would have me, I would be led into a theoretical or evidential relationship with my own intention. And that's not compatible with the intention being a practical conclusion reached through genuine deliberation.

In sum, although Goldman establishes the possibility of predicting one's own intentions, his result is restricted to cases where prediction lies outside the context of practical reasoning toward intending. Hence it does not bear on the situation that is our concern – that of willing shared activity. If we adopt the Predictive Strategy in our account of shared activity, theoretical prediction cannot be screened off from practical intention formation.

---

<sup>36</sup> Moran (2001).

The theoretical perspective would enter into and derail the practical reasoning that was supposed to culminate in intending shared activity. The claim here is not that intentions are incompatible with the prediction of the intentions, or that such intentions cannot be accompanied by predictions. What I am questioning is whether we can arrive at these intentions by way of such predictions, or while also being required to make such predictions.

It might be objected that the following case illustrates the compatibility of predicting an intention with forming it through practical reason.<sup>37</sup> Suppose I am an arms smuggler and that I am about to enroll in a course in applied ethics. You know that the course is highly effective and will lead to my forming the intention, on moral grounds, to quit the smuggling business and to work against it. When you tell me about this, I confidently predict – indeed know – that I will form the intention to quit smuggling and to work to undermine the efforts of others in the business. I do not yet have the reasons to do these things, and so am not in a position to intend to quit or to work against the business. That is, from my point of view, the matter is not settled. Thus, the objection goes, there *is* a point to my later forming the intention on the basis of practical reasoning, once I possess the relevant reasons. The upshot, then, is that a predictive attitude toward one's own future intentions is compatible with subsequently forming the intention through practical reasoning.

In reply, I want first to concede again that it's possible to take a predictive attitude with respect to one's own intentions. My point is that adopting and maintaining the predictive attitude interferes with and undermines one's being able to form the intention through practical reasoning. This example doesn't really speak to this point; I think that the scenario described is possible, but only if we assume that when I go on to form the intention I have forgotten that prediction, or don't attend to it. *This* is why the matter seems open and not settled from my point of view, and explains why I can draw the inference through practical reasoning. In contrast, the objection assumes that the matter can seem open, and be open,

---

<sup>37</sup> I want to thank an anonymous referee for *Noûs* for pressing me on this matter, and to whom I owe the example.

simply because one has not formed the intention against smuggling. This ignores the point that some matters can be settled in one's perspective as a result of one's beliefs and theoretical considerations, and not because of one's intentions and practical reasons. Given that one predicts that one will intend to quit smuggling and to work against it, and one is very confident in this prediction, it strikes me that the matter IS settled, even though one has not yet intended. After all, various alternatives are ruled out: suppose I predict with confidence that after taking the course I will acquire at time  $t_1$  an effective intention to quit smuggling at some later time  $t_2$ . Then I cannot maintain my current intention to (continue to) smuggle at  $t_2$ , even though I now (before taking the course) have reasons to smuggle. That's because I cannot intend to do something (smuggle at  $t_2$ ) if I believe that I won't do it (because I predict acquiring at  $t_1$  the effective intention not to smuggle at  $t_2$ ). Likewise, if I predict that I will work against the smuggling, I cannot form intentions to do things that are incompatible with that. I conclude that the prediction of intentions will undermine practical reasoning toward the intention by undermining the point of it. Unless, that is, somehow the prediction is screened from the practical reasoning, e.g. when one forgets or is distracted. And such a screening off is of course not available to the Predictive Strategy, which employs the prediction precisely to get into a position to form the intention through practical reasoning. So I conclude that the Predictive Strategy, in this form, is not compatible with forming the relevant intention through practical reasoning.

Before moving on, I should note that the Predictive Strategy as presented so far is not the view defended by Bratman (1997). According to Bratman, what allows me to form the intention that we J is not my outright prediction that you will likewise intend. Rather, I believe that you will intend, *if I intend*. The conditional form of my belief relieves me from having to predict my own intention. I only need to be assured that if I were to intend, you would have the intention, and that it would be reliable and effective. So, the issue of our J-ing is therefore still an open one, but one that I now can settle by going ahead and intending that we J. Likewise, you might have the conditional belief that my intention would be effective if you were also to intend. So when you recognize that I've intended,



you can think that the issue is an open one until you form your intention. Thus, you can make use of your conditional belief to form your intention and settle the matter. In general, any one of the participants could think: if I were to intend that we J, then every participant would intend effectively that we J. Armed with this belief, the participant could take him or herself as settling the matter of what to do by going ahead and intending that we J.

But what is not at all clear on this sophisticated version of the Predictive Strategy is whether I could regard *you* (or anyone besides myself) as settling what to do by intending that we J. Indeed, on this view, it seems that I *can't* think that you intend in just the way that I do. To see you as settling the matter of what to do by intending that we J requires that I regard your intention as effective, hence its presuppositions well-founded. That in turn requires me to predict my own intention. But we've seen that it is problematic to do so on the way to forming my intention. So in intending and settling the matter of what to do by intending that we J, I never see you in quite the same light as I see myself, as settling the matter.

Now, this is not a refutation of Bratman's proposal. Bratman shows how it can be (i) that each participant can have an intention regarding our J-ing and thereby take himself to settle that matter of what to do by so intending, and (ii) that each regards fellow participants as figuring in the action as intentional participants so that our activity comes about in part because of the intentions of fellow participants as well as one's own.<sup>38</sup> But (ii) is not quite to hold that each sees fellow participants as intending in just the way that she herself does. The point is that on Bratman's sophisticated version of the Predictive Strategy, I can never see you as I often see myself in shared activity – as settling what to do by intending that we J. For seeing you in this way would entail self-prediction on my part, which I've argued is problematic. It is interesting that a theory that seeks to depict shared

---

<sup>38</sup> This is condition (ii) in Bratman's account, e.g. in Bratman 2009, 157. Bratman (ms.) describes this as a semantic interconnectedness.

activity as involving a highly symmetric structure of intentions would be forced to attribute to each participant a significantly asymmetric conception of what is going on.

This is to say that although Bratman seems to account for the objective version of the settling problem, he seems not to address the subjective version. Recall that this was the worry that seeing oneself as exercising authority or power over fellow participants is at odds with the activity being shared. A very natural way to try to address this concern is to suggest that seeing one as exercising this authority is compatible with holding that one's partner also exercises this authority by settling the matter. But that's what I've suggested one cannot do on the predictive view, on pain of self-prediction.

There may be further avenues to explore in defense of the Predictive Strategy, not to mention the proposal in terms of interdependent conditional intentions.<sup>39</sup> But if I am right that there are costs associated with taking these strategies, it might be worthwhile at this point to get an alternative on the table.

## 5. Acting directly on another's intention

It seems that shared activity can be willed into existence. But so far it's not clear how. We've seen that forming the requisite intention involves one participant exercising some authority to settle what another will do, and that doesn't seem to be compatible with the

---

<sup>39</sup> One might be tempted to appeal to some notion of *trust* instead of prediction, hence avoiding all the worries associated with the latter. I intend my part trusting you will do yours. Or I simply trust that you will act appropriately, and hence I can go ahead and intend that we J. But what is meant by trust here? Is it a license to form something like a prediction, hence allowing us to pursue the Predictive Strategy but without being saddled with the associated epistemic responsibility? If so, how would this work? (If we're appealing to testimony, then that presupposes shared activity.) If the invocation of trust is meant to hint at an alternative to the Predictive Strategy, then we must not rest content with the mere label. Indeed, the strategy I undertake below might be construed as just such an approach. I will nevertheless proceed without attempting in any explicit fashion to capture the contours of what we might mean when we use the term *trust*.

activity being shared. The strategy of predicting the intentions and actions of fellow participants sought to circumvent this problem by proposing a way to form the intention that doesn't require any such authority. But simpler versions of the strategy entail taking a predictive attitude with respect to oneself, which is not compatible with the agential or deliberative perspective necessary for forming the appropriate intention. And I've suggested a worry for a more sophisticated version of the strategy due to Bratman. If the Predictive Strategy won't allow us to evade the issue of authority, perhaps we will make some progress on the problem of willing or intending shared activity by admitting that authority is an issue, and addressing it straight on.

Let's start with a scenario that ought to be familiar enough. Suppose I consent to or agree to your deciding some matter – say where we're to go for dinner.<sup>40</sup> (Perhaps it's your birthday, or I decided last time.) You then decide that we should go to an Indian restaurant, and off we go. What intention am I acting on? One thing I'm doing is acting on *my* intention to let you decide the matter. But that couldn't be the whole story, for if I'm to succeed in acting on my intention to let you decide, it must be that in some robust sense I'm acting on *your* intention as well. Otherwise you wouldn't be settling the matter.

In what sense am I acting on your intention? Notice that you didn't merely *suggest* that we go to an Indian restaurant. That would be a very different scenario. Our case is one where you get to decide. If you've merely proposed something for me to entertain and agree to, that wouldn't amount to you deciding and settling the matter. This suggests the need to relax the individualistic assumption that only one's own intentions can rationally settle what one does. Perhaps your intention that we J can settle what I'm to do, just as my own prior intentions can. The idea is that it is possible for one to take up and act *directly* on another's intention. And surely it must be possible, given that something like what's depicted in this scenario happens all the time.

---

<sup>40</sup> Although she touches on consent in the context of defending a view about promising, I have found Shiffrin's (2008) discussion to be quite illuminating also for the issue of shared activity.

But what exactly does it mean to act *directly* on an intention? In the individual case, when it comes time to act on the prior intention, the practical issue is supposed to be settled (defeasibly, of course). I decided earlier to go for a bike ride now, and in so deciding, committed myself to this course of action. This commitment is reflected in some general norms or requirements of practical reasoning with intentions. For example, when one intends to ride, one should not intend or act in ways that are incompatible with riding, such as spending the afternoon in the office. And one is subject to some sort of demand for intending means that promote one's riding, such as making sure the tires are inflated. The decision and commitment are not compatible with my treating the matter of whether to ride as open – for example, by actively considering going to the movies instead. The matter is settled for me, and I should act directly on the prior intention.

Our example suggests that something like this can happen between people. You issue the intention that we J – e.g. that we go to the Indian restaurant. This intention settles the matter of our J-ing. It wouldn't do so if the matter of our J-ing is, for me, an open question. Part of what enables you to issue the intention is that I take the intention to settle the matter of what I, and therefore we, will do. I (and you) act directly on an intention you've issued for the both of us. Indeed, the suggestion is that if, in the appropriate circumstances, I were not to take up the intention in such a way that it settles what I will do, then I would be violating norms of commitment associated with these interpersonal intentions. I would be guilty of something akin to the irrationality I exhibit when failing to act on my own (appropriately formed, undefeated) intentions.

To understand the normative significance of your intention in my reasoning, it's essential to see here that I intend because the intention you've issued is preserved in me. Suppose instead that I act on your intention, *but not directly*. That is, I currently intend because I have just now independently settled on acting in a way corresponding to how you've intended for me to act (perhaps, for example, I've *now* decided independently to

take on your intention). If that's the case, then in what sense have you really intended and settled the matter? Think about the individual case. When acting directly on my prior intention, I now have the intention to do A but not because I've *now* just decided to A. For then the prior intention would not be doing what it's supposed to do, which was to have settled *earlier* what I am to do *now*. My prior intention must, rather, be preserved in me now.

My acting directly on your intention does not involve any mysterious psychological action at a distance. Your intention does not have some magical immediate causal impact on my psychology. It must in one way or another be communicated to me, just as my prior intention must be transmitted via memory in order to figure in my current thought and action. So when I act directly on your intention, the intention is as realized in my psychology as any that I've issued for myself. The sense in which I can act directly on your intention is that your intention figures in my practical thought in the way that my own prior intentions do: unless somehow defeated, the intention is not a *pro tanto* consideration to be weighed against rival considerations, but serves as a (defeasible) fixed point guiding further practical thought and, ultimately, action.

## **6. Authority and shared activity**

We now have a straightforward story about how it is that shared activity can be willed into existence. One of us simply issues the intention that we J, and when it's time we both act directly on it. In our example, we intend to go to the Indian restaurant because you just decided that we would.

Unlike the Predictive Strategy, the proposal here involves one individual exercising some sort of authority over other participants in shared activity, settling what they will do. A view like this is mentioned by Velleman:

There is nothing problematic about first-person-plural intentions in themselves. One person can decide or plan the behavior of a group, for example, if he holds authority or control over the behavior of people other than himself. If you will do whatever I tell you to do, then what you'll do is up to me, and I am in a position to make decisions about it. As your boss or commanding officer or master, then, I am in a position to decide what you and I will do together, and so I am in a position to form intentions about what "we" will do. (Velleman 1997a, 205)

As the scare quotes suggest, Velleman dismisses this view because he thinks that one person having or exercising this authority over another isn't compatible with shared activity.<sup>41</sup> He is surely right if the authority is as he characterizes it ("you will do whatever I tell you to do").<sup>42</sup> But is this the only way in which one person might have or exercise authority over another?

One feature that distinguishes our case of the restaurant outing from the more dictatorial authority described by Velleman is the element of consent or agreement: it's your birthday, so I'm *letting* you decide for us both. But it might be problematic to appeal to consent as a general explanation of the compatibility of authority with shared activity. That's because consent might *presuppose* forms of shared activity. For example, consent is normally secured through conversation, itself shared activity. There is the further concern that not all instances of shared activity appear to involve consent. Is there normally anything like consent or promising when people come together to play chess or softball?<sup>43</sup>

---

<sup>41</sup> It leads Velleman to the interdependent conditional intentions approach mentioned above in section 3.

<sup>42</sup> Note that the premise here is not fundamentally a moral intuition. The argument is not that shared agency is good, this situation is not good, so this couldn't count as shared agency. Rather, the intuition concerns the exercise of agency. The worry is that the exercise of such a dictatorial authority by one individual puts into question the other's ability to exercise agency in the relevant matter; so much so that it can't be said that they act together.

<sup>43</sup> See Hume's *Treatise* (3.2.2.10 SBN490) example of the rowers for a closely related point.

So, although consent and agreement might be involved in modifying relationships or arrangements to allow for new ways in which individuals might act directly on the intentions of one another, they do not appear to offer the resources for a general account of the authority in question.<sup>44</sup>

But there remain other respects, besides consent, by which the authority an individual exercises in shared activity can be distinguished from a more dictatorial authority that would undermine the autonomy and agency of fellow participants. One involves defeasibility. Some considerations might defeat your intention for us to J: perhaps the restaurant is closed for a private party, or perhaps we learn that a new restaurant that we'd been wanting to try is finally open. What's significant here is that it wouldn't only be *your* awareness of these considerations that subjects the intention to revision or defeat. *I* could become aware of such information that makes a big difference in whether we ought to J. Or, I could learn something about how you arrived at your decision, something that compromises the authority with which you issued the intention. Maybe I discover that when you decided that we should J, you suffered from a certain deliberative bias, blinding you to the significance of a range of considerations important for making a good decision. Or maybe the problem is not so much with missing considerations that make for a good decision as with missing considerations that make for the activity being shared. If the sort of bias you have in deliberation is one that *simply ignores* my interests or the sorts of considerations that I would take to be significant for making a decision on the matter in question, then this too would undermine your authority to settle what I do.<sup>45</sup> Discovering

---

<sup>44</sup> The notion of joint deliberation likewise may modify the form and distribution of interpersonal authority necessary for shared activity, but for obvious reasons of circularity would not serve as a general account of shared activity and how it can be willed into being.

<sup>45</sup> For different proposals or ideas about the constraints on the reasons and deliberation that may figure in shared intention and activity, see Bratman (2007), Westlund (2009), Korsgaard (2009, Ch.9), and Shiffrin (2011). The conditions they discuss can be adapted as proposals for conditions on the authority that is compatible with and necessary for shared activity.

these defeaters, I would now be in a position to revise or reject your intention.<sup>46</sup> Thus, in shared activity, the locus and scope of responsibility relevant for the possible defeat and revision of the intention extends beyond the individual who issues it. This goes a long way toward differentiating the authority to issue shared intentions from a dictatorial authority that has no place in shared activity.

Another distinguishing feature of the authority exercised in shared activity is reciprocity or symmetry. For example, your intention for us to E settles a practical issue for me. But in taking up your intention, I might intend certain means M for us all to take toward E. Thus I will in turn be settling some matters for you. More generally, suppose that there is no significant asymmetry in the authority each of us has to settle what the other does. Then the authority underlying acting directly on another's intention would be very different from the sort that issues in Velleman's dictates.

## 7. Entitlement and Authority

So I don't think that the sort of authority implicit in one individual acting directly on the intention of another is necessarily incompatible with shared activity. But this doesn't yet establish that there is such an authority. In this final section, I'd like to take some preliminary steps in this direction. A promising defense of A's authority to issue an intention for B to act on, comes, I think, by way of B's *entitlement* to the reasons and the deliberation that went into A's issuing the intention. And the resources for a defense of that entitlement are, in turn, to be found in a proper understanding of what it is for agent of the sort that we are to act for a reason. I will argue that sometimes one is entitled to the reasons underlying the intention or decision of another. (This is an argument *that we*

---

<sup>46</sup> The rational dynamics of shared activity depicted here are quite different from those described by Gilbert (1990, 2009), who insists that any sort of modification of shared activity must be done through explicit consent from every other participant.



(sometimes) have this entitlement, not an explanation of *how* it is that we have this entitlement.)

Why think that one person could be entitled to the judgment and reasons underlying the intention or decision of another? Can the reasons acquired through honest practical thought and deliberation accrue to my acting on the resulting decision, even if I wasn't the one that did that thinking? Well, let's consider first the case of individual intention. When I act directly on my own prior intention – say, to go for a bike ride – my action is not merely that of conforming to a prior intention. I am setting out on a bike ride for whatever reasons I had for so deciding, and in the face of competing reasons to pursue other activities or projects. I don't necessarily have those reasons at hand, for not all reasons for A-ing, let alone the rest of the considerations that figure in deliberation, need be available and attended to at every moment during the course of an action in order to guide it through. Constantly attending to the further considerations or reasons in favor of what I'm doing can be distracting, hindering or even undermining practical reasoning and performance of the task at hand.

But though these further reasons in favor of A-ing are not attended to, I am in some sense acting on them. The intentionality of the action is richer than that of merely acting so as to conform to the prior intention. The reasons that prompted the decision and intention explain the action, and serve as a standard for assessing it (e.g. whether the act is completed, or how well it was executed). Moreover, it is an explanation and a standard that the agent does or would invoke, and thus plays a regulative role for the agent.<sup>47</sup> Take an example: I might intend to move the sofa over there, and do so. Perhaps I decided to move the sofa because I wanted to clear a path in order to replace the fridge. When I get around to acting, I might no longer attend to why it is that I'm moving the sofa; I just have in mind to move the sofa over there. But the reason still has a role to play: I've moved it this far so as to be able to get the fridge through; or if the fridge doesn't fit, then that's a

---

<sup>47</sup> Railton (2006, 9ff).

problem with my execution of the sofa movement. So I submit (and I don't think that this is controversial) that I am moving the sofa *in order to clear a path for the fridge*. It is in this sense that the intentionality of what I am doing is richer than just that of acting on the intention to move the sofa, and encompasses (at least some of) the reasons that figured in my decision to move it.

What follows from the fact that I don't have these further reasons in mind (in that I'm not attending to them), but nevertheless am acting *for* those further reasons? On an initially attractive understanding of the relation between the agent and her reasons in intentional action, a reason for the action specifies an end, goal, or aim of the action, and the agent (or some causal mechanism) guides her movements or sub-actions so as to realize or attain the specified aim, often compensating for interfering forces.<sup>48</sup> But we are now considering a case of acting for a reason where the further reason for the action specifies an end that I'm not at the moment explicitly attending to, and which doesn't actively figure in the guidance.<sup>49</sup> Just to put a label on it, I will speak of this alternative relation the agent has with respect to the further reasons for which he or she is acting as that of *entitlement*: I must be entitled to these further reasons insofar as the action described in terms of those further reasons is intentional, even though they didn't figure in active guidance when the action was performed. It's this entitlement that makes it possible for me to act for these further reasons even when not attending to them, and thus underwrites the richer intentionality we noted at the outset of the argument.<sup>50</sup>

---

<sup>48</sup> Frankfurt (1978, 47-8). See also Wilson (1987).

<sup>49</sup> The case I have in mind is one where something like *getting the sofa over there by the wall* is what actively guides the pushing of the sofa, not *making way for the fridge*. (Perhaps it's only after the action is complete that I step back, assess in light of latter, further reason, and depending on the assessment and other factors have another go at it.)

<sup>50</sup> For Burge (1993, 2003), a subject need not be able to articulate his or her epistemic entitlement. I haven't yet given a characterization (let alone a full account) of entitlement (in the practical case) except to contrast it with the more standard case of conscious/attentive guidance toward an end.

So, as I understand it, to be able to act for reasons in this way – that is, in a way distinct from that involved in conscious, attentive guidance toward an end – is part of what it is to be able to act directly on prior intentions. This discussion is meant to articulate what I have in mind when one acts directly on a prior intention. But it also suggests an argument *for* the normativity of doing so. If acting for reasons in this way is an inevitable and unproblematic aspect of one's agency (and I think it is), then it must be that one has an entitlement to such reasons, and to acting directly on one's prior intentions.

Turning to the interpersonal, in acting directly on your intention, I take up the intention you've issued as I do my own prior intentions. What might serve as a starting point for defending an entitlement to rely in this way on the practical judgment and intention (formation) of another? I suggest that it is something like this: I have some sort of reason for what I'm doing that goes beyond what I actually attend to in acting. What I do has a richer intentionality, for in many cases the further reasons entailed by the richer description of the action serve for the purposes of guidance and assessment of my performance, and are standards that I myself do or would endorse. Some one or some others together have formed the requisite intention, or engaged in the deliberation that culminates in an intention to A. I may have never actually done so myself. But granting that I nevertheless take up the intention and am acting for those reasons, I must, then, be entitled to those reasons and the practical judgment that went into the intention.

Here's an example: I'm driving in an unfamiliar city, trying to get to a conference. My guide – who is a lousy driver but has studied the route or perhaps is recollecting from an earlier trip – tells me to make a left, and that is what I do. What's the reason for turning left? Arguably, to get on Dunedin in order to take High Street downtown. The suggestion is that I am acting *for this reason*. And this, even though I don't have any sense of the route we're taking. I'm just attending to executing this turn; perhaps I have in mind that *this* (turning left) is in order that we eventually get to the conference. But I myself haven't taken in the host of other reasons for taking the left, reasons based on consideration of the

route we're taking (e.g., in order to get on Dunedin, so that I can get to High St., etc.). I think that such reasons nevertheless figure in a description of the action as intentional. So I must be entitled to these reasons, if I am indeed acting for them.

It might be objected that there is no need to posit a special entitlement here since I do not act for the route-based considerations that underlie your instruction to turn left. Rather, the objection continues, I'm merely turning left in order to comply with your instructions, since I take it that following your directions will get me to the conference. I don't deny that this can happen. But I don't think that this is typical. Indeed, it would be something of an achievement, requiring no small amount of training and discipline, to narrow one's overall practical outlook so that one is turning left merely in order to comply with the instruction (even if one is following the instruction in order ultimately to get to the conference). Think of elite military training. The normal situation, at least if we are at all tempted to characterize it as an instance of shared agency, is so to speak to look through the instructions or commands to what their point is. I think of compliance with your instruction, not so much as a *reason* for turning left, but as something enabling me to be responsive to the route-based considerations or reasons that you've thought through. Thus, if it is clear how to comply with your instructions, but it seems to me that we're going in the wrong direction, or I'm not getting positive feedback ("Aha, there's a sign for downtown via High Street."), then it would be appropriate for me to start asking you questions: *do you realize that turning left will put us on Dunedin? It looks like we'll then be heading east, but didn't you say that downtown was south?* Whereas, if I'm turning left just to comply with your instruction, I am more sensitive to doing as you say than to what the route looks like. And, given the simplicity of the instruction in this case (it's a simple left turn), it would be less appropriate for me to raise questions about the route – at least not yet.<sup>51</sup>

---

<sup>51</sup> And if someone in the back seat were to ask me why I am turning left, I wouldn't say in order to follow the navigator's instruction. I would, rather, turn to the navigator and say something like, *why am I turning left here? Is it to get on the highway?*

Of course, I am not always so entitled. I might not understand the further reasons for the intention, or the significance of the larger course of events of which my doings are a part.<sup>52</sup> Or, it might be that I would never endorse them. If there is no possibility of communication with those who formed the intention that would somehow allow me to secure the comprehension or understanding necessary for the possibility of endorsement, then that would seem to preclude my entitlement to those reasons. If there were no way for me to simulate comprehendingly the reasoning of those who have decided so as to capture the further reasons for what I'm doing, then that too would challenge an entitlement to those reasons. But that is to be expected, for we would not want to say that one could be acting for these further reasons in those circumstances.

So, one's entitlement to the reasons underlying the intention of another – an entitlement that would underwrite acting directly on that intention – is by no means straightforward. Nothing guarantees that the conditions for entitlement are always satisfied. But if we take seriously the range of cases properly described as acting for reasons, I think that a case can be made for thinking that these conditions are sometimes satisfied.

Turn now back to authority. A natural thought would be that one individual's consent would secure for another the relevant authority. In the case of shared agency, each participant consents to conforming to the relevant intentions of his or her partners. I've already expressed some reservations about appealing to consent. But now we may add a further worry. Suppose that I consent to let you decide. This doesn't explain how, when I act on your intention, I am entitled to the reasons that go into your decision. Without that entitlement, the intentionality of my action is as it were blinkered. I would not be doing what I'm doing for the further reasons that led you to form the intention; I would just be doing your bidding. By itself, my consent may secure for you the authority to settle what I

---

<sup>52</sup> Here, my use of 'entitlement' diverges from that of Burge.

do; but it doesn't necessarily secure for me any entitlement to reasons underlying the decision you make for me or us. Thus, I wouldn't have the agential resources or status to count as making enough of a contribution to the proceedings for it to be *shared* agency. It shouldn't be surprising that consent here fails to secure a robust intentionality and a more full-blown agency. After all, consent might be given to someone to treat me in ways that don't involve exercising my own agency (e.g. I might consent to be operated on, or to cede authority on some matter).

If there is no consent, why think that such an authority may nevertheless exist? I believe that we can extend the conclusion of the previous section. There we came to the view that for example I can be *entitled* to the reasons that went into the intention you issued (e.g. for me to turn left). But this suggests that at least in principle you have the authority to settle what I'm to do by issuing the intention – in the driving example, the instruction to turn left. In brief, we're working with the notion of an entitlement to act for reasons that went into a decision/intention. It's hard to make sense of this sort of entitlement if there were never any authority to form or issue such intentions. In the individual case, how could I, in acting on a prior intention, count as acting for the reasons that led me so to intend if I didn't have in the first place the authority to settle what I subsequently do? In the interpersonal case, I would be without any entitlement if you were not in a position to issue the intention. So you must, at least when I would be so entitled, have a corresponding authority to issue the intention.

## **8. Conclusion**

I do not pretend to have given anything like a complete defense of the sort of authority that underlies the possibility of acting directly on another's intention. My suspicion is that the normative issues here are similar to, and connected with, those that arise in the case of individual intention. Some have noted insightfully that any account of intention and agency

must allow that one can determine what one will do later, without compromising the autonomy of one's later self.<sup>53</sup> I agree that some account of this diachronic authority must be given, that it is an issue that can't be evaded.

Indeed, it is unclear how one could be self-governing if one didn't have the ability to act directly on one's prior intentions and thus the authority to form such intentions. Part of what's involved in autonomous agency is acting in accord with a conception of what is good or right or worth pursuing, a conception that is sensitive to what reasons there really are for acting. But to a significant extent, that will require some foresight: often it is difficult to articulate or discern what is worthwhile, and to figure out how to pursue or implement it. It's not something that can be accomplished on the spot – not always, at least. One important way to avoid overwhelming our deliberative and reasoning capacities when it comes time to act is to plan ahead of time. If at any moment I reach some decision about what to do, I should be able to take for granted that decision at least for some time. Acting directly on one's prior intentions affords the sort of stability required for one's more reflective judgments to have an impact on one's actions. The moral is that acting directly on an intention is not, in itself, at odds with autonomy. It seems, rather, to be a condition for it. (This is, of course, something that could be manipulated or taken advantage of. But no one said that we should be immune to that.)

So, pretending that no such diachronic authority exists would do too much damage to our conception of how our agency is exercised over time. I take a similar view with the interpersonal authority underlying the possibility of acting directly on another's intention. A satisfactory understanding of our agency will require us to understand how this authority

---

<sup>53</sup> Anderson (1996, 452) and Velleman, (1997b, 45-6). See also L. Ferrero (2010). This is not to say that I agree with how these authors think it's possible. They claim that the only way to settle what one will do later without compromising one's future autonomy is to hold that one's prior intention to A gives one's future self a reason to A. I wonder whether this captures the sense in which an intention to A involves a (defeasible) commitment to A-ing. Nor is it clear how, on such a proposal, the reasons that go into deciding and intending to A are preserved so that when one acts on the intention, one acts for those reasons.

can be exercised in the context of shared activity. This is not generally recognized. I've discussed a number of approaches to characterizing the intention underlying participatory commitment in shared activity: the intention to do one's part, interdependent conditional intentions, and – most prominently in this paper – the Predictive Strategy. What unites them is that each in its own way is an attempt to *evade* the issue of authority in shared activity. On these views, there is no need to defend this authority to settle what another does because it isn't required for an account of shared activity; all that I will in shared activity is my own part in it, or what I will is merely conditional on what others intend, or based only on a prediction of what they will do. I don't think that these maneuvers will work. And I hope that what I've said here suffices to show that the phenomenon of acting directly on another's intention is familiar enough, and the underlying authority is innocuous enough, to dispel worries about its compatibility with shared activity. Far from evading the matter of authority in an account of shared activity, we must figure out how best to embrace it.

## Bibliography

- Anderson, Elizabeth. 1996. Reasons, Attitudes, and Values: Replies to Sturgeon and Piper. *Ethics* 106: 538-554.
- Baier, Annette. 1996. Doing Things with Others. In *Commonality and Particularity in Ethics*, ed. Alanen, Heinämaa, and Wallgren. New York: Macmillan.
- Bratman, Michael. 1987. *Intention, Plans, and Practical Reason*. Cambridge: Harvard University Press.
- 1992. Shared Cooperative Activity. *Philosophical Review* 101: 327-41.
- 1993. Shared Intention. *Ethics* 104 (October 1993): 97-113.
- 1997. I Intend that We J. In *Contemporary Action Theory, Volume 2: Social Action*, edited by Raimo Tuomela and Ghita Holmstrom-Hintikka, 49-63. Dordrecht: Kluwer. Reprinted in Bratman, *Faces of Intention*.



- 1999. *Faces of Intention*. Cambridge: Cambridge University Press.
- 2009. Modest Sociality and the Distinctiveness of Intention. *Philosophical Studies* 144:149–165.
- ms. *A Theory of Shared Agency*.
- Burge, Tyler. 1993. Content Preservation. *Philosophical Review* 102: 457–88.
- 2003. Perceptual Entitlement. *Philosophy and Phenomenological Research* 67: 503-548.
- Ferrero , Luca. 2010. Decisions, Diachronic Autonomy, and the Division of Deliberative Labor. *Philosopher's Imprint*. <http://hdl.handle.net/2027/spo.3521354.0010.002>
- Frankfurt, Harry. 1978. The Problem of Action. *American Philosophical Quarterly*, 15: 157-62.
- Gilbert, Margaret. 1989. *On Social Facts*. London: Routledge. (Reprinted Princeton: Princeton University Press, 1992)
- 1990. Walking Together: A Paradigmatic Social Phenomenon. *Midwest Studies in Philosophy* XV:1-14.
- 1996. *Living Together*. Lanham, Maryland: Rowman & Littlefield.
- 1997. What Is It for *Us* to Intend? In *Contemporary Action Theory, Vol. 2, The Philosophy and Logic of Social Action*, edited by G. Homstrom-Hintikka and R. Tuomela, 65-85. Dordrecht: Kluwer Academic Publishers. Reprinted with minor revisions in Gilbert 2000.
- 1999. Obligation and Joint Commitment. *Utilitas* 11. 143-63. Reprinted with minor revisions in Gilbert 2000, 50-70.
- 2000. *Sociality and Responsibility*. Lanham, Maryland: Rowman & Littlefield.
- 2003. 'The Structure of the Social Atom: Joint Commitment as the Foundation of Human Social Behavior', in F. Schmitt, ed., *Social Metaphysics*, Lanham, MD: Roman and Littlefield.
- 2009. Shared Intention and Personal Intention. *Philosophical Studies* 144. 167-87.
- Grice, H. P. 1971. Intention and Uncertainty. *Proceedings of the British Academy* 57:263-79.
- Gold, Natalie & Robert Sugden. 2007. Collective Intentions and Team Agency. *Journal of Philosophy* CIV:3

- Goldman, Alvin. 1968. Actions, Decisions, and Books of Life. *American Philosophical Quarterly* 5:3.
- 1970. *A Theory of Human Action*. Prentice-Hall (1970).
- Hampshire, Stuart. 1965. *Freedom of the Individual*. New York: Harper & Row.
- Hampshire, Stuart & H.L.A. Hart. 1958. Decision, Intention and Certainty. *Mind* LXVII: 1 - 12.
- Harman, Gilbert. 1976. Practical Reason. *Review of Metaphysics* 79:431-63.
- 1986. *Change in View*. Cambridge: MIT Press.
- Hume, David. 1739-40/1978. *A Treatise of Human Nature*, edited by L.A. Selby-Bigge, 2<sup>nd</sup> edition edited by Ph. H. Nidditch, Oxford; Clarendon Press.
- Korsgaard, Christine. 2009. *Self-Constitution – Agency, Identity, and Integrity*. Oxford: Oxford University Press.
- Kutz, Christopher. 2000. Acting Together. *Philosophy and Phenomenological Research* 61:1-31.
- Moran, Richard. 2001. *Authority and Estrangement*. Princeton: Princeton University Press.
- Peterssen, Björn. 2007. Collectivity and Circularity. *Journal of Philosophy* CIV:3
- Railton, Peter. 2006. Normative Guidance. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 1*. Oxford University Press.
- Roth, Abraham Sesshu. 2004. Shared Agency and Contralateral Commitments. *Philosophical Review* 113, no. 3, 359-410.
- . 2003. Practical Intersubjectivity. In *Socializing Metaphysics: the Nature of Social Reality*, Fred Schmitt, ed. Rowman & Littlefield. 65-91.
- . Forthcoming. Team Reasoning and Shared Intention. Springer.
- Searle, John. 1990. Collective Intentions and Actions. In *Intentions in Communication*, edited by P. Cohen, J. Morgan, and M. Pollack, 401-415. Cambridge: MIT Press.
- Shmid, Hans Bernhard. Plural Action. *Philosophy of the Social Sciences*. 38:1. 25-54.
- Shiffrin, Seana. 2008. Promising, Conventionalism, and Intimate Relationships. *Philosophical Review* 117, 481-524.
- 2011. Immoral, Conflicting, and Redundant Promises. In *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon*. Oxford: Oxford University Press.

- Tuomela, Raimo. 2006. Joint Intention, We-Mode and I-Mode. *Midwest Studies in Philosophy*, XXX.
- Velleman, J. David. 1989. *Practical Reflection*. Princeton: Princeton University Press.
- 1997a. How to Share an Intention. *Philosophy and Phenomenological Research* 57:29-50.
- 1997b . Deciding How to Decide. In *Ethics and Practical Reasoning*, edited by Garrett Cullity and Berys Gaut, 29-52. Oxford: Oxford University Press.
- Westlund, Andrea. 2009. Deciding Together. *Philosopher's Imprint* 9:10.
- Wilson, George. 1989. *The Intentionality of Human Action*. Stanford, CA: Stanford University Press.