

without proper discrimination. Thus, in the law in question, the individual participants may be treated equally, and equally harshly. Kutz cites an English case, *Regina v. Hyde* (pp. 230–32). The white-washing worry envisages an alternative scenario in which those who think a group bears guilt for its actions think it is not appropriate to punish individual group members at all for what they did as participants in the group's action. That judgment seems more concordant with the "naturalization" or "reification" of groups.

52. For some remarks on this see Gilbert (2002a, p. 187).

From *Socializing Metaphysics: The Nature of Social Reality*, F. Schmitt, ed. Rowman & Littlefield 2003. 65-91

From *Socializing Metaphysics: The Nature of Social Reality*, F. Schmitt, ed. Rowman & Littlefield 2003. 65-91

3

Practical Intersubjectivity

Abraham Sesshu Roth

The intentions of others often enter into your practical reasoning, even when you're acting on your own. Given all the agents around you, you'll come to grief if what they're up to is never a consideration in what you decide to do and how you do it. There are occasions, however, when the intentions of another (or others) figure in your practical reasoning in a particularly intimate and decisive fashion. I speak of there being on such occasions a practical intersubjectivity of intentions holding between you and the other individual(s). I try to identify this practical intersubjectivity and to take some preliminary steps toward giving a philosophical account of it.

Occasions of practical intersubjectivity are usually those where individuals share agency, or do things jointly, such as when they walk together, kiss, or paint a house together. I do not assume that all instances of practical intersubjectivity are instances of shared agency. But the converse is true: any instance of shared agency involves a practical intersubjectivity holding between the participants. An account of shared agency (or related notions like shared activity, joint action, etc.) is inadequate if it fails to handle practical intersubjectivity.

The chapter is structured as follows. In the first section, I present an example to illustrate this idea of practical intersubjectivity, at least as it appears in the context of shared agency. Practical intersubjectivity is a normative phenomenon, and it is on this basis that in the next section, Intersubjectivity and the Coordination of Intentions, I distinguish it from the mere coordination of intentions some have recognized as essential for shared activity. The task of the third section, *Why Not Intersubjectivity on the Cheap?*, is to show how practical intersubjectivity cannot be adequately described in terms of ordinary intentions familiar from the study of individual agency. Such approaches fail

to handle the rational dynamics of intention revision when practical intersubjectivity is in place between agents. Finally, the last section, Accessibility, lays the groundwork for the revision in our understanding of intention necessary for adequately describing practical intersubjectivity. An important challenge to understanding practical intersubjectivity in terms of intentions as I do here stems from the idea that, fundamentally, one can only intend one's own actions. I contend that this stricture should be relaxed. An advantage of so revising our understanding of intentions is that it yields a satisfying explanation of the social phenomenon of commands.¹

PRACTICAL INTERSUBJECTIVITY

Intentions as Rational Constraints

Suppose that you and I have decided to drive to Vegas together to try our luck at blackjack. We have agreed to take your car, to set out after the morning rush hour, and to take turns driving. One of the details we haven't yet worked out is who will drive when. It turns out that you had to drive a cab all night to pay off some of your gambling debts not covered by your graduate student stipend and so could use a break from driving. So you go ahead and decide to drive the second leg of the trip, figuring that I wouldn't have any objection to this. Moreover, you're correct in your supposition. When you express your intention to me, I don't object and I go ahead and take the driver's seat to drive the first leg.

Why haven't I objected? There are two interestingly different sorts of stories we can tell to fill out our scenario to answer this question. First, upon hearing of your intention, I might consider the matter of who is to drive when and come to my own conclusion about it. Perhaps I prefer to drive the first leg—my salary as faculty member has (so far) covered my gambling debts; I haven't driven a cab all night and am well rested. Or, perhaps I don't much care when I drive, and because you want to drive the second leg, I decide to drive the first leg. There are other possibilities for this sort of story. What's important is that *I resolve the matter for myself*. Since my conclusion fits with yours, there is no reason for me to object to your intention.

On the second story, I don't object because I simply don't concern myself with trying to figure out who should drive. I consider it a matter that has already been settled—by your decision. Given that you intend to drive the second leg, I act accordingly.

In the first story, your intention, if it does figure in my practical reasoning, will only figure as a *consideration*. It is a factor in my decision, one consideration among a number that I might weigh in trying to answer the practical question of who is to drive when. On the second story, I do not face this

practical question because it has already been answered by you. What is left for me to do is simply to act in accordance with how you've answered the question; that is, I act in accordance with your intention. If there is any further reasoning for me to do, it will take your practical conclusion for granted. Your intention will serve as a defeasible *rational constraint* on my subsequent practical reasoning.² I will only consider intending and acting in ways that are consistent with or cohere with your intention. In particular, I will intend to drive the first leg of the trip, thereby making explicit what was implicit in your intention to drive the second leg.

The way in which your intention figures in my practical reasoning in the second story—as a rational constraint rather than a mere consideration—is of particular interest for me here.³ I want to emphasize that it is not an unusual way of reasoning. This kind of thinking is perfectly natural and happens all the time. Some may not feel that this example is a very good illustration of a common and natural way of thinking, perhaps because they feel that the issue of who will drive when is controversial and not something that can be so easily settled in the way I've described. But deciding who drives when needn't be a controversial matter. Consider a modification to our example: you decide that you're going to drive the second leg not because it's your preference to do so, but because you think that I have some preference for driving the first leg. Maybe I'm reticent and haven't expressed my preference, but you know me well enough to have reason to think that I'd like to get behind the wheel sooner rather than later. So now you've decided to drive the second leg, and I take the matter of who will drive which leg to have been settled by you. More generally, take cases where individuals are fairly confident that any decisions made by one will not run roughshod over the preferences and expectations of the others. This confidence might be due perhaps to the familiarity of the individuals with one another and the good will there is between them. One's confidence in the intentions of others might also be based on the nature of the activity engaged in, its circumstances and environment, or the institutions within which it is carried out.

So in our example, though I may have a preference about what to do, I do not get to the point of making a decision on the matter. You get there first and settle the issue. There is nothing left for me to do but to take your intention as a rational constraint and to reason and act accordingly. I think it fair to say that many deliberative or practical matters are settled in this fashion.⁴

The way in which an intention serves as a rational constraint on practical reasoning should be familiar from discussions of individual practical reasoning. When I form an intention to spend the whole day at the beach, I can take it for granted that I will not be spending the day at the library (assuming the library is not on the beach).⁵ My intention to go to the beach is not a mere consideration in favor of the beach option as opposed to the library option.⁶

Rather, it serves as a constraint on my practical reasoning and the deliberative problems I take up. So long as I maintain the intention to go to the beach, going to the library is not an option for me.⁷

My suggestion is that just as one's own intentions serve as rational constraints on one's practical reasoning, it seems that there are natural forms of reasoning that allow one individual's intention to serve as a rational constraint on another's reasoning.

Conflicts between Intentions

Though your intention might have the status of a rational constraint on my practical reasoning, we should not assume that there will never be occasion for a conflict between your intention and one of mine. After all, my own intentions serve as rational constraints on my practical reasoning, and sometimes they conflict with one another. Thus, you may have gone ahead and formed an intention to drive the second leg of the trip and not the first, not realizing that I have similarly intended to drive the second leg and not the first. Our intentions conflict. Given background facts that can be taken for granted by us, these intentions cannot both be satisfied. We cannot drive the same car at once, nor can the car drive itself. We'll get nowhere at this rate.

So one problem with having conflicting intentions is that it is unlikely that we'll have the coordination in behavior necessary for driving together to Vegas. But there's something worrisome about conflicting intentions in this context, even if by chance our behavior ends up being coordinated. You and I might have conflicting intentions about who will drive which leg of the trip, but I might fail to act on mine. Lack of sleep causes me to forget to hand you the keys at the start of the trip. I absent-mindedly hop into the driver's seat, and I continue to drive for several hours until you finally announce that it's your turn to drive.⁸ There is the suspicion that something is amiss here, even though there's no failure of coordination in behavior.

The idea of intentions as rational constraints helps us articulate this worry. Consider the individual case first. My own intentions serve as rational constraints on my practical reasoning. My intention to *A* is, in virtue of its content, intelligibly related to other intentions that I have or could possibly have. There is, for example, the intention concerning the means to *A*-ing. And there's the intention *not* to *A*. And there's the intention to *B*, which I can recognize would preclude *A*-ing, given my beliefs and background assumptions. In virtue of these intelligible relations and my intention to *A*, I rationally ought or ought not to have certain other intentions. Thus, coherence requires that I intend the means to *A*-ing.⁹ And consistency requires that I refrain from adopting the intention not to *A*, or some other intentions the satisfaction of which would preclude my *A*-ing. To the extent that I fail to live up to these principles of consistency and coherence, I am liable to the charge of irrationality.¹⁰

I suggest that the lingering worry with the conflicting intentions in the interpersonal case is akin to the problem in the individual case. The conflict between our intentions regarding who will drive when opens each of us to the charge of irrationality. The rational tension might not be so evident if it is unclear that we're doing something together. But so long as this is kept in mind, then it is quite obvious that we have here a kind of incoherence. This is evidenced by the evident tension in the following assertion: "We (you and I) intend to drive to Vegas together, but I intend to drive the second leg, and you also intend to drive the second leg." This violates a norm of rationality, much in the way that the following does: "I intend to spend the whole day at the beach today, and I intend to stay inside this library all day."¹¹ So the worry with having conflicting intentions about who will drive when is that you and I have intentions that, in the context of our intention to drive to Vegas together, seem to subject us to *rational* criticism. Our conflicting intentions in the case of driving to Vegas together amount to something like the kind of inconsistency in intentions for which an individual might be criticized. The difference, however, is that the inconsistency appears to hold across two (or more) individuals.

So just as my own intentions are intelligibly related to one another in virtue of their contents, so too are the intentions of different individuals. That is, we might speak of the intention of one participant in some activity being consistent or inconsistent with the intention of another participant. (Likewise for coherence.) Moreover, I am entertaining the possibility that these intentions might rationally *engage* one another. This is to say that the intentions of one individual might be rationally relevant for another, serving as rational constraints on the latter's practical reasoning. In some circumstances, if you form an intention, principles of consistency or coherence of intentions may *require* me to take up, discard, or otherwise revise my activity-related intentions.¹² If I fail to meet these requirements, I am subject to a form of rational criticism.

So suppose that there is this sort of rational engagement between the attitudes of certain individuals, where the intentions of each serve as rational constraints for the other(s). I will speak of there being a practical intersubjectivity holding between these individuals. It is an intersubjectivity in light of the symmetry or equality of authoritative status: each individual is in a position to issue intentions that serve as rational constraints for the rest.¹³ The intersubjectivity is practical because it is defined in terms of a rational engagement of *intentions*, as opposed to the epistemic or theoretical intersubjectivity of beliefs.¹⁴

Limitations in Scope

The practical intersubjectivity I have in mind need not be universal in scope. While everyone most likely falls under some instance of practical

intersubjectivity, I don't assume that there is some single instance of practical intersubjectivity that holds between everyone.¹⁵ I have suggested that when you and I drive to Vegas together, the activity-related intentions of each of us will serve as rational constraints for the other. Correspondingly, you and I are subject to a rational demand not to have inconsistent or conflicting intentions. The same cannot be said of other individuals who are not engaged in that same activity. Let me illustrate.

Suppose that on our drive to Vegas we come across a roadblock. It has been set up by a man on a mission to prevent access to that city, perhaps for moral or religious reasons. His intention conflicts with ours. Not all of the intentions of the three of us can be satisfied given the circumstances. In this sense, there's an inconsistency that holds across all three of us. But the charge of irrationality that might have been made against you and me (if, e.g., we had conflicting intentions about who will drive when) does not similarly hold against us and the man who seeks to block our way. My own intentions must not conflict with one another on pain of irrationality. And I take your activity-related intention as one with which my intentions should not conflict on pain of irrationality. But the intention of the blocker does not so present itself to me; I do not take his intention to be a rational constraint on my practical reasoning. A practical intersubjectivity holds between you and me; it does not hold between me and the blocker.¹⁶

Though the intentions of the blocker do not have the status of rational constraints for my practical reasoning, there are nevertheless ways in which his intentions can be relevant for what I do. First of all, the blocker's intention might pose an obstacle for me, much in the way that a dust storm or mudslide on the highway might. Such obstacles might force me to reconsider my intention: if I think that I cannot *A* because of some obstacle, then many philosophers hold that I cannot intend to *A* and will have to revise my intention. Even if, contrary to this view, I can maintain my intention to go to Vegas, it might no longer be worth the trouble. Given what the blocker is up to, I'll need to revise my intention and think of something else to do. But so long as I think that the blocker will fail to act as he intends, this sort of consideration will not force me to change my intention. (Even if I think that the blocker will succeed and so change my mind about going to Vegas, I am not so much seeking consistency between our intentions, but rather just modifying my own intention in light of new information regarding what I think I can accomplish or new information regarding the benefits of the project.)

Of course, the blocker is an agent acting for reasons, not a mere obstacle like a mudslide on the highway. This gives us a second way in which his intention can be relevant for my practical reasoning: his reasons might serve as reasons for me as well. For example, his reason for setting up the roadblock might be that gambling is a sin or that one should not act in a way that benefits the mob.¹⁷ Either way, his reason to set up the roadblock might also

serve as a reason for my not going to Vegas to gamble. And the blocker's actions might present me with reasons against going to Vegas, and thus be an occasion for me to reconsider what I'm doing. So the blocker's intention, in virtue of the reasons underlying it, might get me to change my intention. Now, if I do revise my intention, it was not because his intention served as a rational constraint on my practical reasoning, that is, as something with which my intentions should be consistent. Rather, if I accept his reasons as relevant for what I do, these reasons only serve as *pro tanto* considerations against going to Vegas. That's to say that I might have these reasons while also having conflicting reasons in favor of going to Vegas and gambling. If I revise my intention in accord with the blocker's, it's not because the blocker's intention settles the issue for me but because I was faced with these opposing reasons or considerations—for and against going—and decided in favor of those reasons that happen to agree with the blocker's. Had his intention served as a rational constraint for me, I would not consider whether to continue on to Vegas. I would have taken it for granted that I would not.

So, the intentions of the blocker might have an effect on my practical reasoning and get me to revise my intention, either by presenting me with mere obstacles or with reasons against what I'm doing. In neither case is the blocker's impact that of what I have defined as a rational constraint on my practical reasoning. This is in contrast with the status of the activity-related intentions of a fellow participant in shared activity. It is in this sense that the practical intersubjectivity in the Vegas example is limited between you and me and does not extend to the blocker.

Given that the blocker's intentions don't serve as rational constraints while those of a fellow participant do, we can draw the contrast between the blocker and fellow participant in the following manner: the blocker's intentions are subject to undermining or circumvention whereas those of the fellow participant are not.

Consider cases where the blocker's intention conflicts with mine.¹⁸ Here, there doesn't seem to be any rational proscription against circumventing or undermining the blocker's intention. If his reasons are inadequate or misguided, and if he insists on his course of action, it might be rational for me to try to circumvent or undermine his intention.¹⁹ The decision is, after all, mine to make. If I've already satisfactorily taken into consideration the relevant reasons behind his actions and he and I are still at odds, then the only way that the blocker's intention has any further relevance for me is as a fact about the world or my environment I need to face up to in pursuing my course of action. Perhaps his intention poses an insurmountable obstacle, or one not worth surmounting, and I should give up what I'm doing. But it might instead be an obstacle that can be cleared away or gotten around. The fact that it has the status of an obstacle entails that at least in principle it is subject to this sort of strategy.²⁰ Thus, it *might* be that the most reasonable

course of action available to me in the circumstances is to circumvent or undermine the blocker's intention.

In contrast, the intention of a fellow participant is not similarly subject to circumvention or undermining by me. As a rational constraint in my practical reasoning, his intention is not presented to me as an object or state of affairs that, if it stands in the way of what I'm doing, is subject to undermining or circumvention. A rational constraint is supposed to settle what it is that I'm to do. If there is no proscription against undermining or circumventing certain intentions, then those intentions will not settle anything for me and cannot count as rational constraints in my practical reasoning. So assuming that a fellow participant's intentions do serve as rational constraints for me, they are not subject to undermining or circumvention.²¹

So, a practical intersubjectivity of intentions holds for example between participants of shared activity but does not extend to nonparticipants. The intersubjectivity consists in the rational engagement between the activity-related intentions of the participants. There is a rational demand to maintain consistency and some level of coherence between the relevant intentions. This is what it means to say that the intentions of fellow subjects of practical intersubjectivity serve as rational constraints on one's own practical reasoning. To the extent that one fails to treat his or her intentions as rational constraints, one will be subject to a form of irrationality. If a practical intersubjectivity holds among individuals, the intentions of each are not subject to circumvention or undermining by any of the others.

I don't mean to have given anything like a conclusive argument for thinking that any practical intersubjectivity actually exists. I've described a case that I take to exhibit this intersubjectivity, but others might want to try to describe the case differently. Still, the example serves to gesture at what I have in mind. A fair bit more needs to be said about practical intersubjectivity. To get a better fix on this idea, I would like to contrast it with certain forms of coordination some of which have been offered as necessary conditions for shared activity and related phenomena.

INTERSUBJECTIVITY AND THE COORDINATION OF INTENTIONS

Let me turn for the moment to shared activity.²² This phenomenon involves coordinated behavior or action toward some common end. This sort of coordinated behavior or action can be traced to the coordination of intentions of the various participants. So it is natural to take the coordination of *intentions* as central to shared activity.²³

Satisfying the intention-coordination condition will require each participant to fill in and otherwise modify his set of activity-related intentions in a way that maintains consistency and coherence not only with his own inten-

tions, but also with the activity-related intentions of the others. This will not require that all participants have the same set of intentions, in the sense that each participant's set of activity-related intentions specify the same satisfaction conditions. Coordination is possible with diverging sets of intentions so long as these sets of intentions are to a sufficient extent mutually satisfiable.²⁴

How is the coordination necessary for shared activity related to practical intersubjectivity which has been our concern so far? Because both are a matter of maintaining consistency and coherence between the intentions of different individuals, we might be tempted to identify one with the other. Would this be mistaken? Would it, for example, be wrong to think that the problem with our conflicting intentions about who drives when is simply that we've failed to satisfy this intention coordination condition?

To see the important distinction between intention coordination on the one hand, and practical intersubjectivity on the other, consider what happens when some conflict emerges between the intentions of different participants. What can we say about this case solely on the basis of the intention coordination condition? One thing is clear: there will be less coordination of intentions. Indeed, the conflict of intentions can be so serious that there won't be sufficient coordination of intentions for the condition to be satisfied, and we'll no longer have shared activity. But if the intention coordination condition is all we have to go on, we cannot say that there *should not* be this conflict of intentions, that the participants *ought* (at least *prima facie*) to modify their intentions in order to avoid this sort of conflict in their intentions.²⁵ There is, correspondingly, nothing in this condition to suggest that a lack of coordination entails the possibility of a *mistake* or failure on the part of one individual or the other (or both). The mere fact that there is a coordination of intentions between individuals does not mean that there is any normative commitment, *prima facie* or otherwise, to sustaining it at that level, or at any level at all. Coordination *per se* seems not to be normative.²⁶

The emergence of any conflict between the intentions of two individuals is, in and of itself, the reduction or undermining of intention coordination.²⁷ Not so for intersubjectivity. Indeed, it is often in circumstances of conflict that it becomes more evident that intersubjectivity is in place, or in force. A rational or normative demand tends to be more evident when it is unmet. Practical intersubjectivity involves the normative demand for interpersonal coordination in terms of consistency and coherence of intentions. This demand goes unmet when, for example, someone makes a mistake or otherwise fails to intend in a way that is consistent with the relevant others. The notion of mistake or failure is essential to that of normative demand.²⁸ Given that practical intersubjectivity involves a normative demand, this intersubjectivity and normative demand in principle must be able to remain in force when less or perhaps even none of the coordination called for is in place.²⁹

So, conflicts of intention point to the normative character of intersubjectivity and the non-normative character of intention coordination. When a practical intersubjectivity holds between several individuals, each is subject to interpersonal norms of consistency and coherence in their intentions. There is some rational demand for each to coordinate his intentions with the others. The intention coordination condition involves no such normative demand on the individuals.

We've distinguished practical intersubjectivity from the intention coordination condition in terms of the normativity of the former and the non-normativity of the latter. Practical intersubjectivity involves a normative demand for coordination. As a characterization of practical intersubjectivity, this is fine as far as it goes. But it is not sufficient. It is, after all, possible to bring about the coordination of intentions through mere causal manipulation, or through threats, intimidation, and coercion. This won't do as a way of satisfying the normativity of practical intersubjectivity, and rational engagement does not countenance coordination being imposed in such fashion. The core of the idea of there being a rational engagement of intention between individuals is that what each intends can serve as a rational constraint on the practical reasoning and intention formation of the other(s). This is what's missing when I coerce you into intending in a way that is consistent with me; if our intentions are coordinated only because I've coerced you, then your intentions never had the status of a rational constraint on my practical reasoning and intention formation.³⁰

WHY NOT INTERSUBJECTIVITY ON THE CHEAP?

In practical intersubjectivity, the intentions of another might serve as rational constraints in my practical reasoning. We've seen that this involves a normative demand to coordinate one's intentions with those of another; in certain situations my intentions rationally *should* engage those of another, so that my practical reasoning and intention formation can be constrained by their intentions. In the example of the drive to Vegas, you intend to drive the second leg of the trip, and given the practical intersubjectivity that holds between us, there is a rational demand for me drive the first leg.

There is a way of characterizing this demand—and thereby giving an account of practical intersubjectivity—that many find compelling, but which I think is misguided. As we'll see, this approach may be described as individualistic. It is worth seeing why individualism will not work.

Individualism provides a simple answer to the question of when practical intersubjectivity is in place: whether intersubjectivity holds between some individuals is a matter of each person having the right sort of individual inten-

tion. Returning to our example, a simple version of individualism holds that the central condition for the intersubjectivity holding between me and you is that I have the intention of driving to Vegas with you and that you have the corresponding intention to drive to Vegas with me.³¹

An individualism of this sort insists that I should intend to drive the first leg simply as a matter of maintaining a consistency among my own intentions. Given the *fact* that you intend to drive the second leg, the only intention for me to form that would be consistent or coherent with my own intention to drive to Vegas with you would be the intention for me to drive the first leg. In these circumstances, no other intention of mine will allow me to realize my intention to drive to Vegas with you. This suggests a reductive strategy: interpersonal consistency and coordination of intentions is reduced to intrapersonal consistency and coordination. The demand to maintain consistency and coherence with another party to intersubjectivity just turns out to be a matter of maintaining a consistency and coherence among one's own intentions.

The problem with this sort of individualism should be fairly clear in light of the discussion in the previous sections. The individualist suggests that in general the only way to satisfy my intention of driving to Vegas with you is by intending to drive the first segment. This mistakenly assumes that your intention to drive the second segment is a fixed point, some inviolable feature of my environment. But it needn't be. My intention to drive to Vegas with you might be satisfied if I were somehow to trick or coerce you into driving the first segment. If all that is demanded of me is maintaining consistency among my own intentions, then nothing would rule out this sort of strategy.

But we've seen that this sort of circumvention or undermining of another's intention is incompatible with that intention's status as a rational constraint on your practical reasoning. In practical intersubjectivity, each agent accords the other a certain authority such that the other's intentions have a special status—that of a rational constraint. Just as one accords to one's own intention a rational and practical authority so that it, along with the norm of consistency, imposes a defeasible constraint on what further intentions one forms, so in the case of practical intersubjectivity another's intention is accorded an authority that, along with an interpersonal norm of consistency, imposes constraints on what further intentions one forms. As it stands, the individualist proposal fails to capture this idea. Each individual regards the intentions of others as rationally inert—lacking the impact on his practical reasoning that his own intentions have, and always subject to undermining or circumvention.

I should reiterate that taking another's intention as a rational constraint does not entail that there is no possibility of a conflict between his intention and yours. Conflicts in the interpersonal case can arise just as my own intentions sometimes make conflicting demands of me. In the individual case,

undermining or circumventing one of one's conflicting intentions exhibits a kind of irrationality.³² Rather, one or both of the conflicting intentions must be revised through rational decision. Correspondingly, it would be a sign of irrationality (of a social variety) were one to have to resort to strategies of circumvention—even if it is a reasonable tactic to adopt when one is acting on one's own. Such strategies involve a failure to recognize that the intention in question has some sort of rational authority over one. In recognizing some element of authority in the other's intention, one acknowledges not only the demand for coordination with the other's intention, but also the possibility that in the event of conflict, rational negotiation may require one to revise one's own intention.

The individualist account as it stands fails to capture the thought that the intentions of one individual can serve as rational constraints for another. Why not then revise the proposal so that it requires as a condition of intersubjectivity that each individual grant the requisite status or authority to the intentions of the other? The central thought of this more sophisticated individualism is that your intention has the requisite status and corresponding impact on my reasoning in virtue of a special meta-intention I form: the intention to coordinate my intentions with yours. I call this a *bridge* intention and this more sophisticated form of individualism the bridge intention proposal.³³

The bridge intention proposal is not without its virtues. First, it affords a straightforward account of the source of the normative demand for coordination. In general, the intention to *A* involves a commitment to *A*-ing, and one is thereby subject to a (defeasible) normative demand to *A*. Thus, the intention to coordinate involves a normative demand to coordinate.

Second—and in an improvement over the simple individualism of the previous proposal—the bridge intention proposal has a story to tell about the status your intention has in my practical reasoning. Take your intention to drive the second leg of the trip. We saw earlier that even if I have an intention to drive together with you to Vegas, your intention cannot serve as a rational constraint for me because there is nothing to stop me from circumventing or undermining it. But now my bridge intention to coordinate with your intentions prevents me from doing this. So on the current proposal it seems that your intention to drive the second leg of the trip to Vegas can serve as a rational constraint for my practical reasoning. Let me elaborate.

My bridge intention to coordinate with your intentions effectively transfers to you some element of discretion over what I will do. Presumably this discretion will not be unlimited. I will not grant every one of your intentions this status, but only those that pertain to our trip to Vegas. And among those intentions, there are limits to what you might get me to do; some of the things you decide might be completely unacceptable for me. Within limits, there is a range of intentions or decisions you may form that I will accept as settling

the matter for me. So long as the range within which I allow you to exercise discretion is not arbitrarily narrowed around what I myself would decide on any particular occasion, my bridge intention will genuinely extend to you and your intentions some authority over my practical reasoning. Thus, while I myself would have decided to *B* on some occasion, you might form an intention that in effect settles that I am to do *A* instead. My bridge intention, then, makes a difference and allows your intentions to have a substantive impact on my practical reasoning.³⁴

But despite its virtues, the bridge intention proposal is unsatisfactory. Static cases, where intentions are filled out or elaborated but not significantly revised, are handled adequately. But the proposal loses its grip when we turn to ordinary dynamic cases where there is more significant intention revision in light of changed circumstances. I proceed by describing a case that poses a problem for the bridge intention proposal.

Presenting the case requires a little setup. Notice first that the practical intersubjectivity that holds between a couple of individuals is often limited in the sense that it concerns only certain sorts of intentions, usually those concerning some activity. Thus, in the example of going to Vegas together, I only take as rational constraints those of your intentions that pertain to the trip to Vegas. Unrelated intentions, such as those involved in your plans for next year, or your intention to smoke a cigarette, will not have a similar status. So practical intersubjectivity is circumscribed so that not all of any individual's intentions are taken as rational constraints by the other individual(s). Only one's intentions that in some way pertain to the activity at hand will serve as rational constraints for others. The bridge intention proposal will have to reflect this fact. Thus, my bridge intention will not be to coordinate with your intentions in general, but to coordinate with your intentions as they pertain to filling out and implementing the intention to drive to Vegas together. More generally, the bridge intention will be of the form *I intend to coordinate with your A-related intentions*, where "*A*" denotes the shared activity at hand, and *A*-related intentions are those intentions we form in filling out and executing the intention to *A*. In the case of driving to Vegas together, the bridge intention proposal would have to say that we each have the intention to drive to Vegas with the other, and moreover, we each have the bridge intention to coordinate intentions pertaining to driving to Vegas together.

A further preliminary. Whenever one has an intention, there are some ways that one would be willing to fill it out in carrying it out and other ways that one would not be willing to fill it out. And there are some circumstances that would lead one to modify the intention (revising it, or even giving it up) and other circumstances that don't prompt any modification. Moreover, how intentions are filled out, modified, or defeated can be different in different people, depending on a host of background factors, such as beliefs, concerns, and character.

Now here's the sort of case that's a problem for the bridge intention proposal. Take any two-person case (the argument can be generalized) where one participant (you) forms an intention that conflicts with some activity-related intention of another (me), and that this new, conflicting intention of yours is a result of how you've filled out, modified, or defeated the overall intention concerning the activity. Thus, whereas the original intention was to drive to Vegas with me, your new, revised intention is to *take a bus* to Vegas with me. Or maybe it's the intention to go with me to *Reno* instead. Whatever the case, let us suppose that you think that this is a perfectly reasonable way to revise the intention, in light of the circumstances. For example, the blocker has done a good job and the 15 is impassable, so you think that we should go to Reno instead, given that neither of us minds a lot of driving. Or perhaps you hear about a great deal for bus tickets, and neither of us was excited about doing the driving anyway. Your revised intention, then, is not crazy, and neither of us thinks that it is.

The bridge intention proposal will not adequately capture the normative demand for coordination that I will face in light of how you revise your intention. Recall that on this proposal, the source of the demand to coordinate with your intention, and hence the source of the status or authority your intentions have in my practical reasoning, is my bridge intention. More specifically, it's my bridge intention to coordinate with your intentions relevant for filling out and implementing the intention to drive to Vegas together. But now you do not have any such intentions for me to coordinate with. Your intention now is to take a *bus* to Vegas, or perhaps it's the intention to drive together to *Reno*. My bridge intention has nothing to say about these intentions of yours. As a result, I no longer face any demand to coordinate my intentions with you. The special status or authority your intentions previously had for me *completely evaporates*. Your intention is now like those of some non-participant. Conflicts between our intentions no longer demand the rational resolution they did before. It is now open to me to somehow circumvent or undermine your intention. So if the bridge intention proposal of the individualist is correct, your intentions lose the special status they're supposed to have for me. No substantive practical intersubjectivity would remain between us when you revise your intention in the manner described.

This consequence of the bridge intention proposal is seriously mistaken, and is a reason, I think, to reject it. I think that the correct and natural view to take about this case is that your intention retains its authority as a rational constraint for me. Reasonable revision of intentions does not straightforwardly undermine practical intersubjectivity in this way. Of course, your revised intention very well may conflict with my intention, for I might still be thinking of heading to Vegas rather than Reno, or I may have revised the intention in a different and incompatible manner. But we've seen that such conflicts of rational constraints do not in themselves rule out practical inter-

subjectivity. Rather, when a practical intersubjectivity is in place between two individuals, and there is a conflict between their intentions, then there is a rational demand for them to render their intentions consistent; and it is not open to one to circumvent or undermine the other's intention.

This is exactly what the advocate of the bridge intention strategy *cannot* say. On the bridge intention proposal, your intentions might regain their status as rational constraint for me only if I were to adopt a new bridge intention to coordinate with your intentions pertaining to the new activity. Why would I do this? Presumably, I will make a *decision* about what is expressed in your intention. And this highlights the authority your intention lacks, for my making this decision is precisely what it is for me *not* to take your intention as a rational constraint. Your intention is no longer the sort of thing that is supposed to settle what it is that I'm to do. In effect, any revision of intention of the sort described here will require that practical intersubjectivity be restarted. And this would undermine what is central to the idea of practical intersubjectivity, viz., the thought that the intentions of one individual might serve as rational constraints in the practical reasoning of another.³⁵

At this point, the individualist might try to defend the bridge intention proposal along the following lines. Suppose that our reason for going to Vegas is to gamble. Thus the relevant bridge intention very likely will be not merely to coordinate intentions pertaining to going to Vegas, but to coordinate intentions pertaining to gambling. Thus, when you revise the intention to one of going to Reno, my bridge intention is broad enough to ensure that your new intention has the necessary authority and will serve as a rational constraint for me.

The individualist's reply assumes that there will be a common further goal beyond going to Vegas together. It's not clear that this need be the case; perhaps my aim in going to Vegas is not to gamble (as it is in your case), but to take in the special kitsch peculiar to Vegas—one that Reno cannot provide. But let us suppose for the sake of argument that we share the same reasons for going to Vegas: we both are going to Vegas in order to gamble. Thus we both have the intention to gamble. It doesn't mean that this intention to gamble is somehow *freestanding*. Gambling is an end that can be satisfied by going to Vegas, but there is no reason to think that it is because of this end that we've decided to go to Vegas. On the contrary, we may have adopted this end in order to make the most of going to Vegas. Or perhaps for us the end of gambling comes packaged with going to Vegas and that we would not otherwise pursue or intend gambling.³⁶ Either way, the end of going to Vegas is as it were an end in itself and not merely a means to gambling. If this is the case—and I stipulate that it is in our example—then my bridge intention (and yours) will be to coordinate with your intentions pertaining to going to Vegas (or going to Vegas to gamble); it will not be the broader intention to coordinate with your intentions pertaining to gambling. If this is right,

then the bridge intention proposal will not have the resources to account for the authority or status retained by your revised intention to go to Reno.³⁷

I turn now to another case of intention revision that spells trouble for the bridge intention proposal. Suppose I revise my bridge intention to coordinate with your intentions. This is presumably something I can do, given that the bridge intention is, by hypothesis, an individual intention. It is therefore mine to reconsider should circumstances arise that I judge to warrant reconsideration. If something does come up that gives me good reason to revise my bridge intention, I may do so—even though you don't think it's a reason to revise this intention. Once I revise the bridge intention, the intentions you have (supposing they are not geared to the revised bridge intention, assuming I even have one) will not have the status of rational constraint for me. I am free to ignore (circumvent or undermine) your intention that would have me *A*. I would be able to do some *B* instead.

But this seems to give me a way of shielding myself from any sort of rational objection to my *B*-ing rather than *A*-ing. In revising my bridge intention, I sweep away any sort of authority you may have had as a party to the intersubjectivity between us. Your intentions no longer have a status that demands my coordination. It was this authority and status that was the basis of your objection that by *B*-ing I am not maintaining the proper sort of consistency or coherence between our intentions. By revising my bridge intention, your intentions no longer can make any claim on me. This doesn't seem right. I should not be able so easily to undercut the special status your intentions have for me.

I do not mean to be suggesting that one may never rationally act against the sort of demand to coordinate that we find in practical intersubjectivity. One may sometimes be able to extricate oneself from the demands of practical intersubjectivity and even do this unilaterally.³⁸ What I'm objecting to here is the possibility that one can simply *dissolve* the status had by the others' intentions so that they are no longer the sort of things that can make a demand on one. The mistake here is to represent (1) a case where there is a conflict of rational demands (and one acts on one demand rather than another) as (2) a case where there is no conflict of rational demands at all. We have in (1) a case that requires a rational resolution between these differing rational constraints. I may not simply ignore (circumvent or undermine) your intention in the manner that the bridge intention strategy would seem to allow.³⁹

I think that I have said enough to raise serious worries about the prospects of individualistic approaches to practical intersubjectivity. I now consider a challenge to understanding practical intersubjectivity in terms of rational constraints.⁴⁰

ACCESSIBILITY

When your intentions serve as rational constraints in my practical reasoning, they settle practical matters for me much in the way that my own intentions do. This suggests that you are often in a position to have intentions concerning what it is that I'm to do. For example, your intention to drive the second leg of the trip to Vegas will require me to drive the first.⁴¹ Your intention may have been explicitly formulated only in terms of what you are to do (drive the second leg), leaving implicit what I am to do (drive the first). But in cases like this where the rationally demanded coordination rules out all but one course of action for me, there is no reason to think that what I'm to do could not figure explicitly in your intention.⁴² Thus, if your intention has the status of rational constraint and thereby settles what I'm to do, there is an important sense in which what I do is *accessible* to your intentions. According to the Accessibility Thesis, whether explicitly or implicitly, when practical intersubjectivity holds between you and me, you may intend my action much in the manner in which you intend your own (and vice versa).⁴³

Practical intersubjectivity and the sort of accessibility it entails are at odds with the idea that, fundamentally, one can only intend one's own actions. The Own Action Stricture⁴⁴ on intending is widespread among philosophers of action and mind. Sellars, for example, says that

Intentions pertaining to the actions of others are not "intentions to do" in the primary sense in which *I shall do A* is an intention to do. Thus, in spite of their superficial similarity, *Tom shall do A* and *I shall do A* do not have the same conceptual structure. The former has the form

(ceteris paribus) I shall do that which is necessary to make it the case that Tom does A

whereas the latter cannot, without the absurdity of an infinite regress, be supposed to have the form

(ceteris paribus) I shall do that which is necessary to make it the case that I do A (1968, p. 184)⁴⁵

Perhaps we can avoid the conflict between the Accessibility Thesis and the Own Action Stricture because the accessibility underlying practical intersubjectivity is not as strong as I may have suggested. Sellars allows that my actions can figure in the content of your intention as the intended consequence (in a causal sense) of something you do. This is not normally the way in which your own actions figure in your intentions. But perhaps this weakened form of access will do for the purposes of practical intersubjectivity.

On this proposal, what I do is related to your intention much in the way that a rock and what it does is related to me when I use it as a paperweight, that is, when I intend that it hold these papers down (where this is understood as intending to do something that causes/leads to the rock holding

down these papers). But surely this is not the sort of accessibility we find in practical intersubjectivity. Notice that the rock is not in any way *responsible* for keeping the papers from blowing about. The rock is, obviously, not subject to a rational demand to behave in any way. If it does not keep the papers from blowing about, it is not itself somehow at fault. (If anything, I am at fault for not placing it properly or for failing to act on my intention at all.) Now, supposing the accessibility of my actions to your intentions is understood along these exclusively causal lines, then your intention would be to do something to cause/bring about my *A*-ing. But then I would not thereby be responsible for *A*-ing; I would not thereby be subject to any rational or normative demand to *A*. If anything, my not *A*-ing would reflect a failure on *your* part; you would have failed to do something that causes me to *A*. We have seen, however, that in practical intersubjectivity, I *am* subject to a (defeasible) rational demand to *A* when your intention serves as rational constraint for me. The sense in which my action is accessible to your intention should mark the rational demand to which I am subject when your intention serves as rational constraint for me. The purely causal sense of accessibility therefore will not do for our purposes.

We cannot accommodate the Accessibility Thesis to the Own Action Stricture by weakening accessibility in the manner suggested. Practical intersubjectivity and the accessibility it entails will force us to reject the claim that fundamentally we can only intend our own actions. In rejecting the Own Action Stricture, I am not suggesting that it's possible for your intentions to have some sort of magical access to my actions. Your intention serves as a rational constraint for me, and I act accordingly. Some process must underlie this, much in the way that some psychological process (involving in part a form of remembering) allows me to exercise my ability to act on a prior intention of my own. Moreover, the process that allows for my acting in accord with the rational constraint you set for me could very well involve certain acts on your part. For example, you might have to perform a speech act to communicate your intention to me. Surely something like this must happen, and rejecting the Own Action Stricture doesn't commit me to denying that it does.

The Own Action Stricture derives some of its appeal from being confused with a different and very plausible idea. This is the thought that typically the only way one's intentions can have the impact on the world that they're supposed to have is through one's actions. Thus, realizing my intention that I have a nice meal tonight will require me to *do* something. Similarly, realizing your intention for me to drive the second leg of the Vegas trip will require you to do something, such as communicate your intention to me. This strikes me as plausible, and it can be accepted without granting what I deny: that your intention for me to drive the second leg of the trip is to be understood as the intention to do something that will cause me to drive the second leg of the trip.

Recall that intention involves commitment. My intention to *A* involves a (defeasible) commitment to *A*. I become subject to certain norms so that I am rationally criticizable if my subsequent practical reasoning and action fail to be in accord with the intention. The Own Action Stricture entails that the only way for me to have such a commitment and to be subject to the corresponding norms is through some intention of my own. I will never be subject to this sort of commitment in virtue of another's intention. This is the idea that our inquiry into practical intersubjectivity forces us to reject.

Once we distinguish the Own Action Stricture from the more plausible point regarding what is necessary for realizing one's intentions, there is little to be said in favor of the stricture. And the existence of practical intersubjectivity gives us a reason to reject it.

There is, moreover, another reason to think that the Own Action Stricture on intending must be mistaken. A common form of interaction between individuals involves one *commanding* another to do this or that. I hold that the notion of command involves the expression⁴⁶ of one person's intention regarding what another is to do. Just as Estelle's intention to go to the store is supposed to settle what it is that she is going to do, so it is when Naomi commands Estelle to go to the store. The intention underlying Naomi's command is supposed to settle what it is that Estelle is to do—as if it were an intention Estelle herself had generated.

Notice that Naomi's intention underlying her command is not the intention to do something that will cause Estelle to go to the store. Such an intention would not capture the sense in which Estelle would be responsible for going to the store. Estelle would not face any sort of rational demand to comply; nothing speaks against her ignoring or circumventing Naomi's intention. If Naomi's intention were simply to do something to cause Estelle to go to the store, this would fail to reflect the way in which the command is supposed to settle and commit Estelle to a course of action.⁴⁷

Insistence on the Own Action Stricture would, therefore, rule out the sort of intention familiar to us from the case of commands.⁴⁸ This gives us another reason to reject this stricture.

It is time to take stock of what has been done and to note what is yet to be done. I have characterized practical intersubjectivity in terms of certain interpersonal norms of practical reasoning. More specifically, practical intersubjectivity holds between individuals when the intentions of each serve as rational constraints for the others. Each is subject to a rational demand to coordinate her intentions with the others. This normativity distinguishes practical intersubjectivity from the coordination condition often seen as a necessary condition for shared activity. Individualistic reductions of the rational demand one faces in practical intersubjectivity do not succeed, failing in particular to capture the dynamics of intention revision. Finally, practical intersubjectivity

entails that it is mistaken to think that one can only intend one's own actions. But this in any case is mistaken, as the case of commands illustrates.

A host of interesting questions remain. Two in particular are worth mentioning, if only to dispel certain confusions that may arise from assuming that somehow I have answered them. First, in practical intersubjectivity, an individual has a kind of authority so that her intentions have the status of rational constraints for others. How, in general, does one individual acquire this sort of authority over another?⁴⁹ Second, and relatedly, just when does practical intersubjectivity hold between individuals? I think it is clear that practical intersubjectivity very often does hold between individuals and that it can be sustained through some revision of intentions. (My arguments against individualism rely on this fact.) But I do not deny that often it might be difficult to discern whether practical intersubjectivity is in place, and I have not offered any formula that yields a simple answer to this question.

There are no easy answers to either of these questions. But their difficulty should not tempt us to adopt a skepticism about practical intersubjectivity. To succumb to this temptation would be to blind ourselves to the social reality that surrounds us and pervades our agency.

ACKNOWLEDGMENTS

I got a lot of help from conversations with or comments from Tyler Burge, Lisa Downing, Bill Hart, Barbara Herman, Seana Shiffrin, Fred Schmitt, David Kaplan, Hans Lottenbach, Chris McMahon, and members of a seminar I taught at UCLA during the spring of 2002. I gave a related paper as a talk at the Philosophy Department at UCSB, and benefited from discussion with the audience there, especially Matt Hanser and Kevin Falvey.

NOTES

1. The story I tell here is incomplete. While the form of intention that I identify overcomes certain important obstacles to understanding practical intersubjectivity in terms of intentions, there are other aspects of this intersubjectivity that call for further developments in the proposal and point to important ways in which the sort of intentions implicated in practical intersubjectivity and shared agency differ from the sort we find in commands. It is therefore more accurate to say that the sort of intention we find in shared activity, and the sort we find in the phenomenon of commands, are species in a common genus.

2. The constraint is defeasible for it might be revised given new information or the discovery of conflicts with other constraints (see below). The term *constraint* is meant to mark how the intention has the status of something that both structures and limits practical thought. That said, I'm not especially invested in the term. One might talk instead of rational commitment, default, or presumption.

3. Besides intentions, someone's wants or desires may also figure as rational constraints in my practical reasoning. But if wants and desires are not subject to norms of consistency and coherence the way intentions are, there is the possibility that when I take up your desire as a rational constraint, I might run afoul of your intentions (in the case where you have decided against acting on that desire).

4. I have in mind a case of relative equals; each individual has the authority or status to issue intentions that serve as rational constraints for the other. Nevertheless, there is no guarantee that each will exercise his or her authority to the same extent. It would be worrying if, between us, you were always the one settling these practical issues. I come to resent the control you exert over what we're doing, and you become irritated at how my passivity forces you to do all the planning and decision making. The possibility of such problematic relationships should not, however, put into question the very idea that one individual's intentions might serve as rational constraints for another. A similar problem exists in the case of individual intentions. Someone might have the tendency to form too many intentions too early on in the course of some activity, with a resulting loss of spontaneity and sensitivity to circumstance. That the ability to form intentions for the future might not be well exercised does not put into question the very idea of forming such individual intentions, which commit oneself to a course of action and serve as rational constraints in one's subsequent practical reasoning.

5. This is a point nicely made and developed by Harman (1976, 1986) and Bratman (1987).

6. If it were, then we'd be led to illegitimate bootstrapping. For discussion, see Broome (2001), Bratman (1987), and Velleman (1997b).

7. This is not to say that nothing will ever come up to prompt me to reconsider my intention to go to the beach. In that case, going to the library can become an option for me.

8. I am not sure whether this case where there is coordinated behavior despite conflicting intentions should count as genuinely shared activity.

9. I don't mean to suggest that the intending of means is the only way to increase coherence. Given that one intends *A*, one might form some intention *B* not because *B* is a means to or otherwise facilitates *A*-ing, but because having *B* as an end would make better sense of *A*-ing; *A*-ing would be more worthwhile in such a context. For example, suppose I had an interest in seeing the desert. Then it might be more worthwhile to go to Vegas if I were to take it as an opportunity to see the desert. Apart from its intrinsic merits, a trip to Vegas has the added benefit of being *convenient* for seeing the desert. So I adopt the intention to see the desert as a matter of coherence, though it is certainly not a means to going to Vegas. Another intention I might adopt is to dress in a rat-pack sort of way, or perhaps don my Elvis outfit, and stay not at a newer family-oriented hotel, but at a seedy old establishment. By introducing an element of kitsch into the activity, greater sense is made of it. Adopting the intention to dress in this way is not a means toward the end of going to Vegas, nor need it facilitate it—it could even make it harder. But this intention, too, is adopted as a matter of coherence.

10. That both intentions and beliefs are subject to principles of consistency and coherence has tempted some to identify intention as a kind of belief. See Harman (1976).

11. It needn't be the case that I am irrational in making these sorts of statements (or in thinking the corresponding thoughts). After all, these expressions might be thought or uttered when one realizes that one has been forgetful or has otherwise made a mistake. There is a rational tension here, and one is irrational if one doesn't see it as a problem or is otherwise complacent about resolving it. (Of course, I might see it as a problem, but there might be more urgent things to attend to, like getting out of this burning house.)

12. Why only activity-related intentions? I might have some intention that is not connected with the activity in question, and this might conflict with some nonactivity-related intention held by another participant. It doesn't seem to me that there is a rational demand for us to eliminate this conflict—at least, there is no rational demand having to do with the activity we share at the moment.

13. The idea that the intentions of one might serve as rational constraints for another is also manifest in more hierarchical situations, where not everyone has the same status to issue rational constraints for others (see below).

14. Sellars discusses the logical engagement of intentions between individuals in the course of defending a quasi-Kantian account of morality. See especially Sellars (1968), but also Sellars (1963b). Also relevant for practical intersubjectivity is the distinctive sort of commitments, identified in Gilbert (1989, 1990), between participants in shared activity. I discuss these commitments in some detail in Roth (2002).

The demand to maintain consistency and coherence with the intentions of another raises issues of authority and entitlement. For *A* to be able to form an intention that serves as rational constraint for *B*, *A* must have some sort of authority to settle what *B* is to do. Correlatively, *B* must have some sort of entitlement to the practical conclusion embodied in *A*'s intention. This issue is raised in my discussion of the idea of acting directly on another's intention, in Roth (2002).

My focus is on cases where one individual forms an intention that serves as rational constraint for another. I have not suggested how individuals may settle practical matters by *deliberating together*. No doubt we sometimes arrive at our intentions in this manner. However, it would be too impractical to settle *every* issue between us by deliberating together. When people do in fact deliberate together, they'll communicate and exchange of ideas regarding what to do until a point is reached where one or the other individual (or both) will be in a position to be confident to form an intention that will serve as a rational constraint for herself and the other. (And particularly complicated or momentous decisions might involve a period subsequent to the initial decision making during which one sees if the decision/intention will stick, so to speak.)

15. Nor do I deny it. I leave it as an open question.

16. At least, the intersubjectivity that holds between you and me does not hold between me and the blocker. There may be some other intersubjectivity that does hold between me and the blocker, one that is broader than the one that holds between you and me. But there need not be, and I will be assuming here that there is not.

17. Thanks to Seana Shiffrin for the latter consideration.

18. In one situation, the blocker's reasons might be accepted by me but outweighed by other considerations in favor of continuing to Vegas. In another situation, I might simply reject his considerations as misguided, confused, and so on.

19. Of course, there may be legal or moral proscriptions against circumventing or undermining the blocker's intention.

20. Again, I am setting aside the possibility that there is something like a moral rule or principle that disallows taking this kind of stance against the blocker. It might be that it is rationally permissible to undermine the blocker's intention but not morally permissible. Alternatively, it might turn out that some sort of practical intersubjectivity does in fact hold between me and the blocker, contrary to what I have stipulated in the setup of the example. If so, I would not be able to undermine his intention, as I presently argue is the case with respect to a fellow participant.

21. Since one's own intentions are rational constraints on one's own practical reasoning, a similar line of thought explains why it would be problematic and a sign of irrationality if one were to try to circumvent or undermine one's own intentions.

Bratman (1992) includes a no-coercion condition on shared cooperative activity. He is moved by the intuition that coercion is not a part of cooperative activity. But we can see now that there is a further consideration behind such a principle, one that might explain Bratman's intuition. To the extent that shared cooperative activity involves a practical intersubjectivity, it seems that the no-coercion condition might be based on the idea that coercion is not compatible with the thought that each individual takes the intentions of the others as rational constraints.

22. Recall that my focus has been on practical intersubjectivity, which I take to be necessary for but perhaps not limited to shared activity. Though the example of driving to Vegas is an instance of shared activity, my purpose in describing it is to illustrate points about practical intersubjectivity. I am not trying to give an account of shared activity or agency, and I am simply assuming that this case of going to Vegas is an example thereof.

Practical intersubjectivity (or something very close to it) might be in place between individuals who are not engaged in shared activity when for example they are engaged in the pursuit of individual goals under the restriction that they not interfere with one another. In accommodating one another, each will take intentions of the other as rational constraints (at least insofar as they respect the consistency condition, if not the coherence condition). But there is no explicit shared intention, such as to build a house together, or to go to Vegas together, and the like. (I find it strained to insist that there is an activity that they are doing together, viz., together accommodating one another. But not much hinges on this.)

23. Some theorists resist distinguishing intention as some psychological attitude distinct from intentional action and so would resist any fundamental distinction between the coordination of intention and the coordination of action. On such a view, then, the coordination of action pretty much just is the coordination of intention. To the extent that the former is central to shared activity, so is the latter.

24. To make use of a term from Bratman (1992), the intentions must *mesh*. For example, I might intend that we stop at a cheap place for lunch and not care about the decor, and you might intend that we stop at a tacky place for lunch and not care about the price. Our intentions are not the same; they specify different satisfaction conditions. But it seems that their coordination is possible, given the background beliefs or assumptions that there is a cheap and tacky restaurant somewhere along our route.

A proper formulation of the coordination condition is rather difficult. While some level of coordination of intentions is necessary for shared activity, it would be mistaken

to require complete satisfaction of all activity-related intentions. Shared activity might be had even while some important though nonessential activity-related intentions go unsatisfied and some conflicts between participants go unresolved. I would go so far as to say that this is the norm. We need therefore to allow that two sets of activity-related intentions *mesh* to the extent that their intentions are mutually satisfiable (this is to revise Bratman's usage, which makes mesh to be incompatible with *any* conflict or mutual unsatisfiability of intentions [Bratman 1992, p. 332]. Two sets of activity-related intentions *clash* to the extent that their intentions conflict. Thus, my intentions relevant for our ride to Vegas mesh with yours to the extent that we agree on when to drive, what car to use, what route to take, and so on. But our respective sets of activity-related intentions clash to the extent that we have a conflict over who will drive which leg of the trip. Then, the level of coordination between our activity-related intentions is determined by two factors: the level of mesh and agreement between the two sets of intentions and the extent to which they clash. The coordination condition for shared activity will require some sufficient level of coordination so understood in terms of mesh and clash. I think that the boundaries between what counts as shared activity and what does not are vague, and so the vagueness of the coordination condition seems appropriate. However, I don't pretend that the proposal is not in need of refinement. (Another source of complexity: How many participants in shared activity can fail to have sufficiently coordinated intentions before there is no shared activity? And to what extent can one's intentions fail to coordinate with the rest before one will no longer count as a participant?)

25. There can, of course, be a hypothetical ought: if there is to be shared activity, there should not be this conflict of intentions.

26. I do not deny that the concept of intention is normative. I reject the idea that it is a purely causal/functional concept. The point is that the *coordination* of intentions is not normative—at least the coordination of individual intentions of different individuals is not.

27. Reducing the level of coordination will not necessarily lead to a failure to satisfy the intention coordination condition on shared activity; whether it does will depend on how stringently the condition is formulated. If it is formulated so as not to allow *any* conflict, then the condition would go unsatisfied. But I think most would agree that, so formulated, the condition is too severe as a condition for shared activity.

28. One is tempted to say that normativity entails the possibility of error. But that would be too strong, for it would seem to render incoherent the notion of infallibility, and that doesn't seem right. Rather, we want to say, roughly, that were someone not to conform to the normative demand, then they would be making a mistake.

29. One might think that *some* coordination of intention is a condition for intersubjectivity. Because it is normative, the very idea of intersubjectivity goes with that of mistake or failure—in this case miscoordination. So we can't require *perfect* coordination of intentions as a condition for intersubjectivity. For as soon as there was miscoordination, intersubjectivity and its normative demand would not be in place. The normative demand of intersubjectivity would become empty: holding only for those for whom it is otiose, and not applying to those for whom it would be substantive.

We will need to address at some point the nature and extent of coordination that would be a condition for intersubjectivity. It might be that coordination plays a part

of a genetic condition for intersubjectivity, and that it is possible for the coordination to disappear even while the intersubjectivity and the normative demand remains.

30. See the remarks on Bratman in note 21 above.

31. The inspiration for this position can be found in Tuomela and Miller (1988), who claim that in what they call joint activity, each participant has a "we-intention," where this is analyzed in terms of individual intentions of the form *I intend to do my part in our A-ing*. It should be noted that Tuomela and Miller are not there defending a view meant to provide an account of practical intersubjectivity.

32. This is not to say that it can't be a rational strategy. But if one needs to resort to this sort of strategy, then one suffers from some form of irrationality. For example, it might be rational for me to try to undermine some compulsive intention. Another case where it might make sense to adopt a strategy of undermining an intention is when I lack a sort of rational unity over time (e.g., my day self doesn't want to eat M&M's, but my night self does—and not because it's a compulsion; at bedtime, I decide in a deliberate and cool-headed fashion that eating M&M's now is a fine thing to do, whereas six hours earlier, and the next morning as well, I think that I really should not eat the M&M's).

33. Bratman (1992, pp. 332–34) defends a view like this, although it is not formulated as an account of practical intersubjectivity. He attributes to each participant in shared activity an intention to mesh subplans.

34. The nature of the sort of reduction being attempted by the sophisticated individualist will depend upon how the content of the bridge intention is formulated. If my bridge intention is to accord to your intentions the authority requisite for their serving as rational constraints for me, then the individualist might be appealing to concepts that can no longer be considered purely individualistic in content. In that case, the individualist is not offering a conceptual reduction of practical intersubjectivity; but the account is individualistic in that nothing besides intentions of individuals is appealed to. On the other hand, the individualist might attempt a more ambitious and controversial semantic reduction and try to specify the content of the bridge intention in arguably individualistic terms. On this view, my bridge intention might be something like the intention to form intentions that are consistent and coherent with the intentions you form. Both sorts of proposals will be susceptible to the criticism I give below.

35. Think of how the very notion of one's own intentions serving as a rational constraint is undermined if every time I were to act on a prior intention, I could not take the previous decision for granted, and would have to decide the matter anew.

I don't deny that there are ways in which an individual might revise an intention that would lead to the dissolution of the practical intersubjectivity holding between them. But the bridge intention proposal has the implausible consequence that the intersubjectivity is dissolved *every time* that someone revises an intention in a way that is not anticipated by the bridge intention.

36. Though gambling may be higher in the order of justification, this does not mean that it is higher in the order of what it is that we're settled on doing, let alone what it is that we're settled on doing together.

37. Another suggestion on behalf of the individualist is to relax the bridge intention in our example to something like the intention to coordinate with your intentions pertaining to our going to Vegas, *or to any similar activity*. But what would count as

sufficiently similar activity? Similarity is notoriously vague and there are many different dimensions along which activities might resemble one another. What activities count as similar to going to Vegas? Would spending the weekend playing blackjack on the home computer count as sufficiently similar? Would going sightseeing in the desert count as sufficiently similar? Or going gambling in Reno? Or going to Grace-land (supposing you had in mind kitsch when heading to Vegas in the first place)? Even if we set aside this worry and grant that we can make sense of similarity here, there is the further worry that it's not at all clear that the rational revision of intentions must lead from the intention to perform one activity to the intention to perform an activity similar to that originally intended. For example, on our way to Vegas, we're held up by a traffic jam (no doubt because of the blocker). So we decide to head home and get ahead of schedule doing research for a paper we're planning to jointly author. Given the circumstances, it could make a lot of sense to revise the intention in this fashion. But this activity really doesn't seem to be similar to that of going to Vegas to hang out in the casinos.

38. Pace Gilbert, who characterizes her central notion of joint commitment in terms of the claim that no one party to such a commitment can withdraw from it unilaterally. See Gilbert (1999). That said, I should add that Gilbert's work has been influential in my thought.

39. It is useful to compare this with an analogous proposal regarding the status one's own intentions or decisions have on what one will do. Suppose that at time t_1 , I decide/intend to send a paper to an editor at some later time t_2 . My decision/intention should settle the matter of what I will do at t_2 . It is now t_2 . Given my decision, and the fact that nothing has since come up that gives me reason not to act on it, I should be sending the paper now. Why? Is it because at t_2 I have what amounts to a diachronic bridge intention to act on the prior intention I formed at t_1 ? That couldn't be it. For nothing would stop me from discarding the intention to act on the prior intention, in which case the prior intention would lose its status to demand any sort of behavior on my part. But, in fact, the intention does retain its status for me. My prior intention *persists*, and isn't merely *represented* in my subsequent practical reasoning and action by the intention to act on that prior intention.

40. The arguments presented against individualism—of both the simple form and the more sophisticated bridge intention proposal—rely on the thought that they fail to reflect what is evident: that practical intersubjectivity persists through some revision of intention. Of course, there is some revision in intention that *would* bring practical intersubjectivity to a halt. The issue of when it is that a practical intersubjectivity holds between individuals is a difficult and interesting one that needs further exploration. Its difficulty is obscured by individualism, which gives an overly simplistic (and as we've seen, mistaken) answer in terms of the special individual intentions ascribed to each party to intersubjectivity.

41. This isn't always the case. Often the intentions of yours with which I'm supposed to coordinate may be such as to leave me significant latitude in deciding what to do. Nevertheless, in principle, and often in practice, your intention can narrow things down so much as to specify just what action I am to take.

42. Once we admit this, there is nothing to stop us from allowing for the possibility that your intention might have explicit content that exclusively concerns what I am to do. For example, your intention might have been for me to do *all* the driving.

43. Accessibility in its purest (and perhaps most controversial) form is manifested when your intentions specifically concern what I'm to do. But accessibility more generally is manifest whenever your intentions serve as rational constraints for me.

44. To adapt Michael Bratman's term for the view.

45. Other proponents include Baier (1970), and more recently Stoutland (1997), who holds

an agent can intend only to do something *herself*. She cannot intend anyone else to do some act, but at best intend to do something herself that gets someone else to do the act. I cannot intend *you* to buy *me* a dinner; I can only intend to do something which might result in that. (55–56)

See also Velleman (1997b) for important discussion of another related thesis that might also raise worries for what I say about practical intersubjectivity here. I hope to discuss that article on a different occasion.

46. Anscombe (1963, p. 5) rightly points out that commands are expressions; there is no such thing as an unexpressed command. Intentions, however, are not themselves expressions, though they may be expressed. Talk of expressions of commands is redundant. The current thesis is that a command is an expression of one individual's intention regarding what some other individual is to do.

47. Often we utter things that sound like commands in contexts where no relations of authority are in place. These are not genuine or substantive commands and do not carry with them the sort of status that induce a commitment in the individual to whom they're issued.

48. One way in which the situation of commands differs from that of practical intersubjectivity lies with the nature of the relations between the relevant parties. In the case of commands, there tends to be an asymmetry: only one individual usually has the authority to settle practical issues for the other. In the case of practical intersubjectivity, authority is symmetrical: each individual has some authority with respect to the other so that each can issue intentions that serve as rational constraints for the other.

49. Corresponding to this question of authority is the question of how it is that one individual is *entitled* to the practical or deliberative conclusions of another. These issues receive some discussion in Roth (2002).