

Shared Agency and Contralateral Commitments*

Abraham Sesshu Roth

My concern here is to motivate some theses in the philosophy of mind concerning the *interpersonal* character of intentions. I will do so by investigating aspects of shared agency. The main point will be that when individuals act together one must be able to act *directly* on the intention of another or others in a way that is relevantly similar to the manner in which an agent acts on his or her own intentions. What exactly this means will become clearer once we understand what it is to act directly on one's own intentions. But I take it to be a fundamental assumption of the prevailing individualism of the theory of action— one at the core of its conception of the separateness of individuals — that one person cannot act directly on another's intention. I agree that there is an important way in which we are or can be separate and autonomous thinkers and agents. But the way the individualist tries to capture this separateness is misguided.

Taking a participant's contribution to shared activity to be an exercise of individual agency, the individualist considers which acts of independent agents add up to instances of acting together, and tries to articulate the conditions under which they do. Acting together — so understood — is constituted by individual acts, and the theory and

* Thanks to Tyler Burge, Lisa Downing, Barbara Herman, and Seana Shiffrin for all their help with this paper. Special thanks to Michael Bratman for generous discussion, and to Margaret Gilbert and two anonymous readers for the *Review* for detailed comments on an earlier draft. Thanks also to Andrew Hsu, David Kaplan, George Wilson, Bill Hart, Gavin Lawrence, Joseph Almog, Terry Parsons, Pamela Hieronymi, Christian Perring, Rahul Kumar, Joel Pust, David Silver, Michelle Casino, Adam Kazan, Eli Hirsch, David Wong, Berent Eng, Gary Ebbs, Scott Shapiro, Tamar Schapiro, Dave Hilbert, Tony Laden, and Marya Schectman. I have also benefited from comments and questions from audience members when ancestors or versions of this paper were delivered at philosophy departments at UCLA and University of Illinois at Chicago (UIC), and at a Social Ethics and Normative Theory workshop at Stanford. This work was supported by a UCLA Faculty Career Development Award, the Faculty Grants Program of the UCLA Academic Senate Council on Research, and the Institute for the Humanities at the University of Illinois at Chicago.

concepts of shared agency would be reduced to those of individual agency.

It is reasonable to consider reductive approaches to shared agency, and indeed there has been lucid theorizing in this vein.¹ However, the view that a participant's contribution to shared activity is simply an exercise of individual agency leads to an implausible description of many cases of shared agency and ultimately cannot be sustained. Even with a more sophisticated version of individualism wherein each participant conditions his or her performance on the contributions of other participants, the underlying picture of social interaction seems to be one of suspicion, featuring wary agents coming to an understanding in an almost painfully diplomatic manner. Thus, in the case of two people going for a walk, we have the following tentative interaction: "I'll go for a walk, *if you will*." "Well, *since you're willing*, I'll go as well." It almost sounds as if these individuals are negotiating an end to hostilities. Perhaps some walks are negotiated endings of hostilities. But when you think about it, most are not.

An investigation of shared agency that departs from this individualistic orientation will have interesting implications for the philosophy of mind and action in general, including individual thought and agency. These issues, it turns out, are not as settled as this sort of reductionism assumes.

These reservations about the reductive approach and hopes for a less individualistic alternative would not be compelling if reductive individualism in some form were to capture everything that we want to say about shared agency. But it does not. My case against individualism will draw attention to certain aspects of shared agency, emphasizing in particular (i) the commitment of each participant to the activity or its end, and (ii) the distinctive pattern of commitments between participants when they act together in the relevant sense. I think that these commitments are best understood in terms of intentions. Much of what I do in this paper is directed toward developing an account of shared activity – and these commitments in particular – in terms of the practical language of intention. It is in developing such an account that we will discern a need for an interpersonal conception of intentions.

But the nature and even the existence of these commitments cannot be taken for granted; indeed they are matters of some controversy. So something needs to be said about them before I can proceed to the implications for our understanding of intentions. This is what I do in the

¹ Michael Bratman's series of papers is a prominent example. Though I defend a rival view, I am in no small way indebted to his work.

next two sections. Section 1 seeks to establish that each participant in shared activity has a participatory commitment to the activity or its end, and this commitment is linked to the concept of intention. Section 2 introduces the special “contralateral” commitments that hold between participants in shared activity. Shared activity and related notions are often given accounts in terms of intentions. But formulating a satisfactory theory in this vein turns out to be rather tricky. Sections 3 and 4 show how certain proposals in terms of individual intentions of one form or another fail to account for the participatory and contralateral commitments. Finally, in sections 5 and 6, I develop the components of an interpersonal conception of intention. Section 5 introduces the central idea of acting directly on another's intention. Section 6 makes use of the authority and entitlement underlying acting directly on another's intention in order to capture the pattern of contralateral commitments that we find in shared agency.

1. Participatory commitment

My goal in this first section is to defend the thesis that each participant in some shared or joint activity such as that of walking together is *committed* to that activity.² Now, it is not all that clear what commitment should amount to in this context, and in some senses of the term the thesis seems false. Mightn't one have a non-committal attitude toward one's walk with someone if, for example, one suspects that person might turn out to be irritable and unpleasant company? So in what sense are participants necessarily committed to the activity they share?

I think that each participant is committed in that at least for now, the participant can answer the question of what he is doing or will be doing by saying for example ‘We are walking together’ or ‘We will/intend to walk together’. I will be assuming that expressions such as these are generally expressions of intention.³ And so, the rather modest sense in which the participant is committed to walking together is the sense of

² The commitment can be regarded as a commitment to the activity, or a commitment to the goal or end of the activity. Depending on the example, it might be more appropriate to speak of one rather than the other. I do not think that much of what I have to say here will require that we carefully distinguish between these.

³ That is, ‘we are walking together’ and ‘we will walk together’ are not usually to be read merely as *descriptions* or *predictions* of what we are or will be doing, or – in the case of ‘we intend to walk together’ – not as a mere description of some psychological attitude.

commitment associated with intention.⁴ That's to say that a decision has been made to walk together – or, at least it's as if a decision has been made. While there may yet be a change of mind and the decision/intention revoked, one has gone beyond merely desiring or having some reason to walk together. Having a desire to walk together is compatible with having a desire *not* to do so; a reason is *pro tanto* – it is a consideration in favor of the action, but there may be considerations against it as well. Whereas, in intending, one has gone beyond the point of weighing considerations for and against, and has committed to a course of action. And unlike desires, intentions are subject to norms of consistency, so that the intention to walk together is not compatible with an intention not to, nor with other intentions that would preclude walking together. So when someone asks me what I'm doing, and I say 'We are walking together' or 'We will/intend to walk together', I'm expressing an intention concerning the shared activity, and thereby also expressing my participatory commitment to that activity.⁵

But why think that the intention is somehow attributable to me in particular, so that *I* have this participatory commitment? Perhaps you and I constitute a group, and *it* is committed to the walk, but that this needn't entail that I myself am so committed. One might reply along individualistic lines and insist that a group is nothing over and above its constituent individual members, and that a group's being committed to X is just a matter of its members individually being committed to X (under certain conditions). But is there any other response that does not so straightforwardly invoke individualism?

To establish participatory commitments on the part of each individual in shared activity, we need to show that somehow the intention and commitment being expressed for example by 'We will walk together' is not, or not merely, the intention of a group comprising everyone encompassed by the 'we'. Rather, the intention must be *distributive*, so that it will attach to each of the individuals referred to by the 'we'. Only then

⁴ For an understanding of intention as being settled on or committed to a goal, see Harman 1976, 1986, and Bratman 1987. Some remarks in Grice 1971 perhaps anticipate this view.

⁵ A number of theorists understand shared or joint activity in terms of intentions, most prominently, Tuomela, Bratman, and Searle. Bratman (1992) explicitly formulates his account of shared cooperative activity in terms of intentions in order to handle, among other things, what I call participatory commitments. Kutz (2000) uses the expression 'participatory intention' to denote a form of individual intention that contains reference to what he calls collective acts. As will become clear below, I do not think that the intention underlying participatory commitment is an individual intention.

would ‘We will walk together’ as thought or uttered by me express a participatory commitment *I* have to the activity.⁶

To see that such an intention expressed with a plural grammatical subject must be distributive so that it expresses a participatory commitment, consider how we are to understand the possibility of irrationality or rational tension in one’s intentions.⁷ Suppose again that I am or will be engaged in some sort of activity with someone else. When asked by a third party what I am up to, I say for example ‘We intend to go to Vegas by car’, or more simply, ‘We are going/will go to Vegas by car’. Given that this intention is reasonable and justified, rational norms of coherence and consistency for intentions require, roughly, intending the best means for going to Vegas by car (or at least means that are good enough). If only the two of us are going, and my partner cannot drive, then a further intention is then required of me: I should intend to do the driving. But suppose that, knowing full well that the only other person in the car cannot drive, I maintain that we intend to go to Vegas by car, but also insist that I will not do the driving. Then it seems that *I’m* being incoherent. When I insist on not doing the driving, I’m expressing an intention not to drive. I don’t merely not want to drive; I’ve resolved not to do so, and am committing myself to this course of action. Given the circumstances this seems to be at odds with what’s expressed when I say ‘We intend to go to Vegas by car’. *I* am being irrational if I insist on these both.

A natural way to understand the tension in my thought here is to see it as involving a conflict between my commitments. I am committed to not driving, but I’m also committed to something that requires me to drive. But to understand the tension this way requires that when I say ‘We intend to go to Vegas by car’, I’m expressing a commitment on my part characteristic of intention – specifically, a commitment on my part to our driving to Vegas together.⁸

⁶ I don’t mean to rule out that the thought or utterance may express *more* than just my own participatory commitment.

⁷ Sellars (1968) emphasizes logical relations holding between intentions in his discussion of ‘we-intentions’.

⁸ One might object that there are better explanations of the tension between ‘we intend to drive to Vegas together’ and ‘I will not do the driving’, and so I cannot conclude that the former is an expression of commitment on the part of the speaker. On one such alternative, the tension amounts to a straightforward inconsistency in belief. The intention to go to Vegas by car implies the belief that I will do so, and given that someone must drive and the other person cannot, I come to the belief that I will do the driving. On the other hand, an intention not to do the driving implies the belief that I won’t drive the car. Thus, the tension between the corresponding intentions is a matter

There is an inconsistency or incoherence in *my* thought when, in the given context, I say or think ‘we intend to drive to Vegas together, but I will not do the driving’. This tension suggests that the intention expressed in the first part of the sentence is distributive and manages to attach to me as opposed to being something that can only be said of the entire group.⁹ Only then will *I* have the conflicting commitments that in this example underlie the charge of inconsistency or incoherence in my practical thought.¹⁰

of these conflicting beliefs. This explanation assumes – controversially – that whenever one intends to A, one believes one will A. Moreover, the proposed explanation cannot generally assume that my doing the driving is the *only* means, and hence necessary for going to Vegas by car. It is, rather, the best means (or the only one that is good enough); but there might be other means, such as inviting another licensed driver to come along for the trip, or perhaps giving one’s partner a crash course in driving. So it doesn’t follow that I can infer the belief that I will drive simply on the basis of the intention and belief that we will drive to Vegas together, even knowing that you cannot drive. Thus, when the means in question is not necessary, we do not generate an inconsistency of belief in the way presupposed by this proposal; it therefore fails to explain the tension between the intentions in question.

One might instead try to appeal to beliefs about likelihoods in order to explain the tension in my thought. Given the intention to go to Vegas together by car, it’s very *likely* in these circumstances (with a non-driving partner) that my means will be to do the driving. So I predict that it is likely that I will drive. But given my intention not to drive, I believe that I will not drive. Hence the tension. But it’s unclear how one arrives at the prediction of likelihood here. What exactly are the circumstances in question? Shouldn’t they include the fact that I intend *not* to drive? If so, then presumably it’s *unlikely* that our means will be that I drive. So, given that I am aware of my intention not to drive, the prediction of likelihood cannot be attributed to me to account for the tension in my thought. Of course, it’s still possible to attribute a generalized belief to the effect that people in similar situations usually will adopt such and such means. But such a belief is entirely consistent with my belief that I will not drive. So unless we have some independent consideration to attribute to me the predictive belief specifically about my own case, this proposal will not account for the tension in my thought.

⁹ My point is not linguistic. Some predicates are distributive. For example, given ‘Bob and Sue are walking’ we have e.g. ‘Sue is walking’. But other predicates are not: ‘A and B and C ...and Z surround the fortress’ will not yield ‘A surrounds the fortress’. My point is that when A says ‘we intend to surround the fortress’, the *intention* is distributive, and attaches to A, and may express A’s commitment to the fortress being surrounded. I trust that my claim concerning intentions and commitments has no implications for the linguistic point. In particular, it does not violate the non-distributivity of the predicate in question.

¹⁰ Sometimes the plural grammatical subject ‘we’ might be used to ascribe to a group an intention that does not necessarily attach to each of its members. In such cases, an individual might express individual intentions conflicting with the intention of the group without being subject to the charge of inconsistency or irrationality. Thus, ‘We (the U.S.) intend to invade Canada, but I will act to ensure the failure of this endeavor’, or, ‘We (the philosophy department) intend to hire Mr. Boe Zoe, but I will do everything in

Having established participatory commitments in this fashion, it should be noted that we have yet to clarify how the participatory commitments attributed to each of the participants are related to one another. The participatory commitment stems from an underlying intention, but I have not said anything about whether or not the underlying intention is an individual intention of the sort that in principle could be had independently of what the other participants intend. Determining how exactly this commitment or intention should be characterized is a delicate task. For the moment, I only make the modest claim that each individual in shared activity has this participatory commitment, and that underlying it is an intention *of some form or other*. I must turn to another feature of shared agency, something I call contralateral commitment, before I begin to make the case that the sort of intention underlying both participatory and contralateral commitments is *not* the sort we find in individual agency as it is generally understood.¹¹

2. Contralateral commitments

I turn now to an important feature of shared activity explored by Margaret Gilbert: the distinctive pattern of what I call contralateral commitments that hold *between* its participants.¹² An understanding of this phenomenon will suggest how an individual's participatory

my power to convince him not to come.' In each example, the speaker introduces some act to be performed by a group or organization of which she is a member, and expresses an individual intention conflicting with that of the group. There's nothing problematic about this *per se*. Groups often have dissenting members. Whether there is a problem here depends upon the extent to which the group in question will tolerate dissension and still include dissenters as members. Surely this will vary from group to group. Correspondingly, a member might identify himself with some group, but be vehemently at odds with some particular activity of the group.

I need therefore to distinguish between the shared activity of individuals (e.g. walking together, driving to Vegas together) and the activity of groups constituted independently of any particular activity they might perform. Setting aside the latter sorts of cases, we are left with cases where the rational relations (and the possibility of conflict) between the expression of a participant's personal intention on the one hand, and his expression of the intention associated with the shared activity on the other, demonstrate the existence of a participatory commitment had by that individual regarding the shared activity.

¹¹ Although I here attribute a commitment to each participant in shared activity, and claim that underlying this commitment is an intention that "attaches" to each individual, sections 5 and 6 will make clear that there are crucial differences between this kind of intention and the sort we find in individual agency.

¹² See Gilbert 1990, and 1989, Ch. 4, especially 162-3. I am influenced by Gilbert's insight here, though I don't quite agree with everything she says about these commitments (see below).

commitments, discussed in the previous section, are related to those of other participants in shared activity. Consider Gilbert's example of going for a walk. If Jack and Sue are walking together, and Jack begins to walk so fast that Sue cannot keep up, Sue is entitled to complain to Jack that they were walking together and that he is walking too fast. Sue's entitlement appears to be based on Jack being in some sense *committed to Sue* to maintain a pace that allows Sue to keep up. Gilbert speaks of obligations, but I prefer talking of commitments. Sue has reciprocal commitments to Jack. The commitments, of course, are not limited to keeping pace, but also concern a host of factors that are necessary for or facilitate walking together. For the sake of simplicity, I will focus on the particular commitment to keep pace.

Sue is in a *special* position to criticize Jack when he walks too fast. It's quite possible for there to be *prima facie* grounds for *anyone* to criticize Jack, much in the way that A's breaking of a promise to B, or A's assault on B, are grounds for some third individual C to criticize A. But it is important to distinguish the criticism that could be made by Sue from those that might be made by some third party not participating in the activity.¹³ Sue's criticism should be seen as a part of the shared endeavor. Criticism by non-participants cannot be seen in this light. Whether or not welcomed by the group as a whole, such third party criticism is not a part of the exercise of shared agency.¹⁴ So if Sue is in this special position to criticize Jack, this suggests that a distinctive form of commitment is a basis for her criticism, one that does not serve as a *prima facie* ground for non-participants to object to Jack. I will mark this special and *executive*¹⁵

¹³ Gilbert (1990) doesn't clearly indicate that a participant is in a *special* position to criticize another participant, but I suspect that she had something like this in mind.

¹⁴ Against such external criticism, a participant could in principle reply along the lines of, "Mind your own business". Compare the case of individual agency. Someone else might criticize your actions, but if they lack any authority over you (moral or otherwise), then their criticism lacks the sort of practical force that your own self-criticism has. A related point is made by Korsgaard (1993) concerning Nagel's idea of subjective or agent-relative reasons. Her view is that moral reasons are intersubjective. If they were merely subjective, then it wouldn't be one's business to intervene with a moral judgment.

¹⁵ The commitment is executive in that the criticism it would license would be a part of the activity in question. It is helpful to draw an analogy with individual cases. Think about when one is performing some task and says to oneself, *wait let's see - this is supposed to go there, not here*, and one then moves the object in question to the appropriate place. The criticism as made by the agent herself (at least as she's performing the task) has a practical and executive force that is lacking when it is put forth by someone else. The criticism is a part of what it is for her to carry out the task. It's true that the content of her own criticism is presumably available to others as a part of their criticism of what the agent is doing. Someone else looking at the agent might also judge that the object

commitment with the admittedly awkward locution of Jack's *commitment-to-Sue* to keep pace with her. I will also be describing these commitments as contralateral, to emphasize that they hold between or across participants in shared activity.

But why think that there is some special form of commitment between participants in shared activity? Why not simply explain contralateral commitments in terms of familiar moral principles that underlie our obligations to one another and apply when individuals act together? For example, mightn't Jack's commitment-to-Sue simply be a moral obligation that results from his having promised Sue that he would keep pace with her? Or, if there wasn't anything as formal as an exchange of promises, didn't Jack at least lead Sue to believe that he would not walk ahead of her? Then perhaps Jack's commitment-to-Sue is simply the moral obligation that stems from it being wrong to create in others false or misleading expectations concerning what one will do.

Characterizing Jack's commitment-to-Sue simply as a form of moral obligation does not, at least on the face of it, address the special and executive nature of contralateral commitment. Indeed, it is an important feature of many moral principles that they generate obligations that are *not* special. Moral obligations must be able to serve as a basis for third party critique and, if necessary, intervention. For example, when one individual assaults another, it's not simply their business. The duty not to harm others, or more specifically its violation on this occasion, can be grounds for someone not involved to get involved, to criticize and even to try to stop the aggressor.¹⁶ But morality would not justify the involvement of third parties if its obligations were only of the special variety. The specialness of the commitment or obligation, after all, lies precisely with something about them not being available to third parties.

This is not to say that there are no moral obligations that are not in *some* way special and executive. But it's not clear that moral obligations simply *as such* account for or explain the special character of these commitments. Therefore, if we're to see whether Jack's commitment-to-Sue turns out to be a moral obligation, we will need to turn away from appealing simply to the notion of moral obligation in the abstract. We must consider instead some specific moral principles to see whether they

belongs in one place rather than another and that she ought to move it. But their criticism is just a judgment about the agent's performance - perhaps a part of some evaluation of her, or at best some sort of advice regarding how she should proceed. It is not a part of what she is doing.

¹⁶ This is not to deny that on many occasions, there are reasons, some of them specifically moral, for avoiding involvement.

might generate, in addition to obligations that can ground third party involvement, special obligations that might account for Jack's contralateral commitment-to-Sue.

An exhaustive survey of possibilities along these lines will not be possible here. I propose instead to consider some principles adapted from those formulated by Scanlon (1990, 1998) for a different purpose, namely, accounting for the moral obligation to keep one's promises, but which look promising for our concerns. Although I do not dispute Scanlon's principles, I will argue that they do not suffice to *account* for contralateral commitments.¹⁷ These commitments exist over and above (or should I say under and below?) any sort of moral obligation generated by the principles proposed by Scanlon.

One such principle concerns the wrong of a form of manipulation:

Principle M*: In the absence of special justification, it is not permissible for one person, A, in order to get another person, B, to do some act, X, to lead B to expect that if he or she does X then A will do Y, when in fact A has no intention of doing Y if B does X.¹⁸

Whether an individual will participate in shared activity often depends on the expectations he or she has about what the other participants will do as their part of the activity. These expectations, in turn, tend to be a matter of what the others have led the individual to believe. So there is no doubt that Principle M* may *apply* to shared activity. But no element in this principle would seem to place B in a special position not had by some third party to object to A's actions. So M* will not generate *special* and *executive* commitments that account for Jack's commitment-to-Sue.

Notice that M* does not rule out the possibility of changing one's mind, even in the absence of special justification. All that it takes for M* to be satisfied is that one have the appropriate intention when one created the expectation. The situation is different in the case of promising, which as I mentioned was Scanlon's concern. The moral obligation here is such that once A has promised B that he'll X, A cannot just change his mind (absent special circumstances that would justify doing so). As Scanlon would put it, B has the "right to rely" on what A has promised. Since M*

¹⁷ Bratman (1997a) is surely correct in his claim that Scanlon's principles can apply to cases of shared activity, but I want to resist his suggestion that they account for contralateral commitments.

¹⁸ This is a simplified version of Scanlon's Principle M (1998, 298).

allows A to change his mind, it provides no such right. Scanlon concludes that a principle like M* does not underlie the moral obligation to do as one promises.

That M* allows a change of mind (even absent special justification) is also a reason to think that it is not sufficient for generating the contralateral commitments of shared agency. The worry now is not that M* fails to generate a commitment that is contralateral or special, but that it simply doesn't generate much of a commitment at all. Although I think that some sort of unilateral withdrawal from shared activity (absent special justification) is possible, M* makes withdrawal all too easy.¹⁹ Can M* be supplemented with another principle to generate a commitment that can be relied upon? Perhaps a principle to the effect that if A has led B to form an expectation regarding what A will do, A must take steps to prevent B from any loss stemming from A not living up to that expectation. But as Scanlon points out, such a principle of loss prevention is neutral between warning the other party that one will not be living up to expectations, compensating them for not doing so, or in the end living up to the expectations (1998, 300-1). The availability of the first two remedies entails that M* and the principle of loss prevention do not really amount to a commitment to do one's share; others cannot take as settled that one will do one's share, rather than one of the alternatives, in order to fulfill the obligation to them.

In M*'s place, Scanlon formulates the following principle F, which does introduce a special and executive element, and thus looks promising as a basis for contralateral commitments:

Principle F: If (1) A voluntarily and intentionally leads B to expect that A will do X (unless B consents to A's not doing so); (2) A knows that B wants to be assured of this; (3) A acts with the aim of providing this assurance, and has good reason to believe that he or she has done so; (4) B knows that A has the beliefs and intentions just described; (5) A intends for B to know this, and knows that B does know it; and (6) B knows that A has this knowledge and intent; then, in the absence of special justification, A must do X unless B consents to X's not being done. (Scanlon 1998, 304)

The element of consent serves Scanlon in capturing the idea that B has the "right to rely" on what A has promised. A cannot simply

¹⁹ More on this below.

change his mind after leading B to expect that A will X. Outside of special circumstances, A must now get B's permission not to X. What's interesting for our purposes is that the element of consent places B in a special position with respect to A's obligation. While anyone, including third parties, is in a position at least in principle to criticize A for not living up to his obligation, only B is in a position to release A from the obligation.²⁰ So does Principle F account for contralateral commitments? We should not be too hasty in answering in the affirmative.

Notice, first of all, that Principle F will not account for the commitments between participants when there is no intentional creation of expectations. Sometimes each participant is familiar with the other participants and thus already has expectations about what they will do in the circumstances. And sometimes, social circumstances are such that it would be reasonable to form expectations about what even unfamiliar others will do despite a lack of communication. Principle F does not apply in such circumstances, and therefore does not explain the contralateral commitments that nevertheless would hold between the participants.²¹

Recall that F was formulated to account for the moral obligation associated with promising. This points to a second worry about appealing to F to account for contralateral commitments. F's emphasis on assurance, while entirely appropriate as an account of the obligation to keep one's promises, is problematic as a basis for the contralateral commitments. Promising and its associated obligations are typically invoked when there is some concern about whether one individual can rely on what another says he'll do. Perhaps the promiser is not very reliable, or the matter is of particular importance for the promisee. The invocation of a promise is meant to assure that the promiser can be relied on in this situation, and that he'll accord it the importance that the promisee sees in the matter. This focus on assurance is reflected in the particularly strong condition for permissible change of mind: (absent special circumstances) I cannot without your consent change my mind about what I've promised you. But this overriding concern with assurance is misplaced when we turn from promising to shared agency. Shared activity is very often carried out by individuals who trust one

²⁰ A proposal based on principle F also has advantages over an account that explains contralateral commitment in terms of an exchange of promises. Although it is meant to be an account of the moral obligation to keep one's promises, F does not assume that there is any explicit exchange of promises, nor does it even assume that there is already in place a social practice or institution of promise-keeping. Presumably, shared agency is more basic than the social practice or institution of promise-keeping.

²¹ This is a limitation Bratman (1997a, 139-40) himself concedes.

another. There is often simply no need for the assurance that promises are meant to provide. Thus, familiars may act together even when no one particularly wants the assurance, and no one seeks to provide it – that is, when conditions (2) and (3) of F are not satisfied. So the contralateral commitments between the individuals in such cases will not stem from Principle F.²²

Another worry also stems from F being formulated with an eye toward accounting for promising. One’s participation in shared activity may be subject to revision in a way that F fails to capture. I might acquire some reason to withdraw from or modify my contribution, and I may at least in some circumstances do so unilaterally.²³ On the current proposal, (absent special justification) I must get the permission from all other participants before I can do this. This fails to acknowledge properly the (limited) element of discretion one has in revising one’s commitments in shared activity. The reasons one must have to revise one’s commitments in shared activity (or to withdraw from it altogether) needn’t be as strong as those necessary to be released from the obligation to keep a promise.²⁴

²² The trust in question needn’t be limited to the sort of trust that justifies A’s relying on what B has led A to believe. It’s also possible for A to trust B’s *judgment*. This may allow B to do something other than acting in accord with A’s specific expectations. If circumstances change, B may be trusted to modify the shared activity to take advantage of the new opportunities, and may do so even without A’s consent. This possibility wouldn’t be accommodated by an account of contralateral commitments that appealed solely to F.

²³ Similarly, you may insist, despite my protests, that I withdraw from the activity. This will also pose worries for an account based on Principle F.

²⁴ Of course, this criticism is also a worry for an account of contralateral commitments that appeals to an explicit exchange of promises. I disagree with Gilbert who thinks that one cannot unilaterally release oneself from contralateral commitments. I think that Gilbert’s characterization of the commitment is too strong, and invites analysis in terms of the exchange of promises (or Principle F), though she herself is not sympathetic to such an approach. On Gilbert’s view (1990, 1997), all of the other participants must concur in order for one to be released.

As an individual taking part in shared activity, there are two “points of view” imposing rational constraints to which you are subject: your own, and that which is not exclusively your own. There will then be two sets of commitments, and these might not be mutually satisfiable. Nevertheless, since the norm of consistency governing intentions requires the avoidance of this sort of conflict, you will need to revise the intentions corresponding to these commitments. The kind of authority one has to revise the shared intention (i.e. the commitment that comes with the shared activity) will not quite amount to the authority one has with respect to one’s individual commitments/intentions. Does this mean that you must inevitably adjust your individual commitments and intentions to make them conform to your commitments in shared activity? Not necessarily. It would be perverse to insist that one has no authority at all in the matter. Thus, if one discerns

Finally, the attempt to understand contralateral commitments in terms of the executive element of consent (as in Principle F) does not, at least by itself, account for the demand that B's contralateral commitments-to-A be compatible with A's own commitments and intentions regarding the shared activity. A and B might have agreed to some activity or goal G to be obtained in part by B doing certain tasks T. B has a contralateral commitment-to-A to do his share, specifically a commitment to do T. Suppose that B now thinks that what he should do is T' rather than T. A doesn't think that this is a good idea. He thinks that T' will not further the shared goal G, and that it might even undermine it. As A sees it, B's doing T' does not cohere with and may even be incompatible with A's G-related commitments and intentions. But B doesn't see this, and moreover, B is very stubborn. So A finally consents to B doing T' instead of T, figuring that when B runs into trouble, he'll finally realize his mistake, and will switch back to doing T.²⁵ According to Principle F, B has been released from the commitment to do what A takes to be necessary for G, and now has a commitment-to-A incompatible with A's own G-related intentions and commitments. If our view is that contralateral commitments are to be understood in terms of Principle F and the idea of consent, then we would have to deny that there is any problem or glitch in shared agency in this sort of case. This seems mistaken. When individuals act together, their relevant intentions and contralateral commitments are subject to interpersonal consistency and coherence. This demand is not met in the present example, where A's G-related intentions conflict with B's contralateral commitment-to-A

an emerging conflict between these sets of demands, one might be in a position to revise or modify the shared goal in such a way as to avoid a conflict between shared and individual commitments. Of course, how you revise the shared commitments is defeasible and subject to challenge by other participants. Nevertheless, this is quite different from the proposal based on F (or the exchange of promises) where one needs the permission of all the other participants before one revises the shared commitments or withdraws altogether from the shared activity. There is a spectrum of possible views about permissible revision or withdrawal. The spectrum begins at one end with a view like that based on M* which would allow any simple change of mind, and runs through e.g. a view that gives one a default entitlement to revise or withdraw, but where this is subject to challenge by other participants. Finally, at the far end of the spectrum we have a view that requires consent from all other participants. Accounts in terms of principle F or an exchange of promises pass over this spectrum and settle at the far end, with a view that imposes the strictest conditions on withdrawal or revision.

²⁵ For example, in painting a house together, you might insist now that you have to go purchase some special new brushes. I know that this sort of brush will not work as well as the those we already have, that it'll waste time for you to try them, and that the store in any case will be closed.

to do T'. And the problem stems from the fact that one may consent to another's doing something even when what the other seeks to do is incompatible with one's own intentions. Consent does not have built into it anything that would account for the interpersonal consistency and coherence that is demanded in shared activity.

We have been exploring the possibility that an account of contralateral commitments might be given in terms of some familiar moral principle. One principle, Scanlon's Principle F, does contain an executive element that might account for the special nature of contralateral commitments. However, there are a number of reasons that suggest that F does not embody the sort of executive element generally appropriate for shared agency. I propose now to focus directly on the task of capturing the executive element, without worrying whether it might be formulated in terms of a specific moral principle. To that end, I will turn to a consideration of "intention-based" accounts of shared agency.

But before I do this, let me briefly recap. Assuming that the discussion of section 1 is correct, Jack and Sue both have a participatory commitment to a common goal – that of walking together. We have developed from Gilbert's insight the further point that Jack has a *commitment-to-Sue* to pursue that commonly held goal, to do his share in their pursuit of this goal. Sue has the corresponding *commitment-to-Jack*. This distinctive pattern of special and executive commitments – what I've been calling contralateral commitments – are characteristic of shared agency.

Given that Jack and Sue are walking together, I maintain that in addition to his contralateral commitment to Sue, Jack also owes it *to himself* not to stray ahead of her. Thus, if Jack notices that he's walking ahead of Sue, he might criticize or rebuke himself for not doing what is necessary for the shared activity he and Sue are engaged in. Let us call *ipsilateral* the commitment each participant in shared activity has to him- or herself. We will see that accounting for ipsilateral commitments in shared agency is relatively straightforward, whereas contralateral commitments will be trickier to explain.

3. Initial intention-based accounts of contralateral commitment

Whether or not they are explicitly formulated to account for *participatory* commitments, a number of theories of shared agency will have

the resources to say something about those commitments.²⁶ That's because these theories are formulated in terms of the intentions attributed to each participant, and the concept of intention is tied to a modest form of commitment, as I've suggested in section 1. More problematic are attempts to give intention-based accounts of the *contralateral* commitments between participants, such as Jack's commitment-to-Sue to walk side by side with her. I will discuss the problems facing intention-based accounts of contralateral commitments in this section and the next. This discussion will also provide an opportunity to introduce some ideas concerning intending that will be relevant for the account to be defended later in this paper.

Assuming that our discussion in section 1 was on the right track, we have reason to believe that each participant in shared activity has a commitment to that activity. We are also working with the assumption that this commitment is a matter of having the right sort of intention. Thus, if several individuals throw a dinner party together, each has an intention which, at its most basic, might be naturally expressed as, "We will throw a dinner party". Failing to meet this commitment – for example, by not performing the necessary means – is a matter of violating a rational norm associated with intending.

The first intention-based account to be considered suggests that the agent's participatory commitment to shared activity suffices, in itself, to establish his or her contralateral commitment to the other participants. Thus, Jack's participatory commitment to the activity entails or perhaps just is his contralateral commitment-to-Sue.

It's not clear to me how this proposal can work. What does seem right is the idea that one's commitment to a goal will involve or perhaps amount to a *commitment- to-oneself*. To see this, consider, first of all, that in forming an intention, one sets a goal and settles some deliberative matter. One is thereby committed to the goal and subject to norms of intention. To the extent that one fails to live up to this goal and to satisfy these standards, one's practical rationality is open to a certain sort of criticism. Furthermore, these norms or standards are brought to bear by someone – oneself in this case – who had the authority in the first place to settle the matter and determine what to do. One has the right to commit oneself and to impose upon oneself these standards of intending. Failure to live up to this commitment would therefore be a failure to live up to one's own authority and standards. It is in this sense that one has not only a commitment to a goal (by which one's actions are evaluated and subject

²⁶ See references in note 5 above.

to rational criticism). One has, in addition, a commitment to oneself to act toward that goal. Thus, Jack had the authority to set this standard for himself, and he did so. If Jack does not act accordingly, Jack has not only failed to satisfy some standard. Jack has let himself down.²⁷ It is in this sense that when one intends to F, one has a *commitment-to-oneself* to F.²⁸

I think that it is very plausible to think of ipsilateral commitment along these lines. Ipsilateral commitment, recall, is the commitment each participant in shared activity has *to himself* to do his part in the activity. We have just seen that when one forms the intention, one does not merely take on a commitment to whatever it is that one has intended. This commitment to a goal is also a commitment-to-oneself to fulfill the goal, in the sense that the norms of intending are in force because the agent intends. So if each participant has the intention expressible as “we will A”, then the ipsilateral commitment each individual has to himself is thereby established.²⁹

The proposal under examination, recall, was that Jack's participatory commitment to the shared act or its goal suffices for Jack's commitment-to-Sue. The problem for this proposal is that it is unclear how one's own intention to pursue a goal amounts to a commitment to anyone *besides* oneself. Consider the hosts of the dinner party. Suppose each individually has an intention to serve as one of the hosts for the party. Suppose that one fails to act on his intention. Then he fails to live up to his commitment to this goal, and thereby (granting what we've just said) fails to live up to a commitment he has to himself. However, none of the practical norms of *individual* intention seem to have anything to say about whether this individual owes it to the other participants to make his contribution to the joint activity. After all, the supposition is that it's only *his* intention. And whether or not he lives up to it is not the business of the other participants. The relevant perspective from which to apply the practical norms of individual intending is the perspective of the intender/agent himself, and not the other participants. Hence, as far as the norms of intention are concerned, he owes it to himself – and not the

²⁷ Barbara Herman points out that this is related to the case of weakness of will. When I behave akratically, I let myself down.

²⁸ One's authority here qua intender is not arbitrary. It is governed by the norms of intention. That's why when one fails to act on an intention for no good reason, one cannot as the relevant authority simply rule that the intention is not in force.

²⁹ And this, I think, reflects why it is that a participant in shared activity is in a special position to criticize himself. His critique is from the agent's point of view. It is a part of the very exercise of agency that it critiques, whereas a third party critique – even if intended as supportive and conducive to the success of the act – is not an exercise of the agency behind the act.

others – to make his contribution. And that means that we haven't captured the contralateral commitments.

Another example might help. Suppose that Mark has the goal of getting \$10 from me. Suppose also that I form an intention to give Mark \$10. Finally, suppose that I don't promise \$10 to him, and don't tell him about my intention. (This will, I hope, eliminate the distorting effects of social practices and moral factors that are often a source of commitments between individuals.) Now, do I have a commitment to Mark? Well, my goal is to give him \$10. And we have seen that it is a commitment-to-myself. But it's not a commitment-to-Mark. Mark does not have any special entitlement to complain if I don't act on my intention to give him \$10. If I don't hand him the money, the only commitment I fail to live up to is that which I have to myself; I fail to live up to my own intentions. So just because you intend some goal or end that someone else intends, you haven't necessarily formed a commitment-to-that individual – even if what is intended is some activity involving the two of you. The case of one's intention regarding shared activity is similar. In and of itself, my individual intention regarding our activity, even in the context of a similar intention on your part, does not entail any commitment-to-you. I only manage to establish an ipsilateral commitment-to-myself.

Suppose we were to modify the previous case so that I intend not merely to give Mark \$10, but to give him \$10 *and thereby act in a way that accords with his desire or intention to get \$10 from me*. It is interesting to note that even in this sort of case no contralateral commitment-to-Mark is generated.³⁰ I'm merely committed to giving him \$10 and thereby conforming to his intention. I desire to give him \$10 and conform to his intention just because I think it's a good idea to do so. I want to spare Mark severe agitation and avoid a scene. It's not that I somehow recognize his intention and practical perspective as *authoritative* for me, as *decisively settling* what I am to do. Hence, practical norms of intending apply solely from the perspective of the agent himself. This is the case even if what is specified as the goal/content of the intention is that it conform to the goal that has been expressed in another's intention.

So like the original proposal it was meant to revise, the proposal that each individual in shared activity intends the activity “in accordance

³⁰ Actually, there might be some ambiguity in the expression ‘acting in accord with another's intention’. My remarks here are based upon what I take to be the most natural reading, viz., that my action will fit with or conform to the content A of the other's intention, but that the decision to A was mine, and that I'm really acting on my own intention to A. There may be another reading of the expression more in the spirit of the position I will defend later in this paper.

with” the others’ intentions fails to establish contralateral commitments between the participants.³¹ In intending what the other person intends (even if I so intend with the intention of fulfilling the content/goal specified in the other’s intention), I exercise my own practical authority, thereby establishing the commitment to myself, and not the other person. Thus my subsequent failure to act would not be evaluated in terms of the other’s intention, but my own. The other individual does not, on the basis of anything considered so far, have the authority to settle decisively what I am to do.³²

Of course, I could in the end act in a way that corresponds to what the other had in mind for me to do. But from what has been said so far, this is likely just due to the fact that I recognized what it is that the other sought for me to do, and *decided for myself* to do it. My actions do not somehow amount to a recognition that the reasoning or decision of the other as authoritative for me. (Had what the other intended been different, I could very well have rejected it as something for me to do.) To (re)intend another’s intention in this sense is to exercise one’s own practical authority; acting on such an intention is merely to live up to a commitment one has to oneself.

Of course, an individual intention might generate a commitment to others if the contents of the intention include concepts of the sort to establish such commitments. You might, for example, just intend to participate in the shared activity conceived as such. Then it might be that you establish some sort of contralateral commitment. But notice that any theoretical proposal that makes use of individual intention with this sort of

³¹ This is why Bratman’s (1992) condition – that each participant in shared activity has an intention to act “in accordance with” the intentions of the other participants – will not generate contralateral commitments. Though Bratman also adds that it is an intention to act “in accordance with and because of” the intentions of the others, this still will not suffice. The added causal condition does not imbue the other individuals with the normative status of (partial) authorities over the participant’s actions. It should be noted that Bratman did not propose this condition in order to account for contralateral commitments, but rather to rule out certain types of coercion that would otherwise have been counterexamples to his account of shared activity.

³² Bratman (1997b) tries to address criticisms made by Velleman (1997) regarding the issue of *discretion*, and in the end he seems to adopt a view similar in some respects to that of Velleman. Both think of discretion in a non-normative fashion, *not* as some kind of practical authority one has over what’s to be done, but as a power to cause some state of affairs to come about. While I think that perhaps Velleman and Bratman succeed in accounting for discretion understood as causal power, I think that they both fail to address how it is possible for one person to have (at least in part) discretion over another, where ‘discretion’ is understood in normative terms. In the individual case this discretion is the authority the agent has over his own actions.

content will take for granted the phenomenon we're investigating. It is hardly satisfying to try to understand shared agency and the sharing of intentions by invoking one's ability to form intentions to share agency with another.

So we are still without an account of the contralateral commitments. Intending seems to be a rather individual affair, and seems only to generate ipsilateral commitments. How might we make some progress in accounting for the contralateral commitments?

4. Another intention-based account: interdependent conditional intentions

One difference between shared activity and a mere aggregation of individual acts is that the intentions and actions of the participants in the former are dependent on one another in a manner not necessarily true of the latter. Perhaps contralateral commitments result from this interdependence. One sophisticated strategy along these lines makes use of the idea of interdependent conditional intentions. After emphasizing the importance of what I've been calling contralateral commitments, Gilbert describes how shared activity can come about:

...each person expresses a special form of *conditional commitment* such that (as is understood) only when *everyone* has done similarly is *anyone* committed. Thus all wills are bound simultaneously and interdependently.³³

Gilbert doesn't talk specifically of intentions, but whether or not she has intentions in mind will not matter for the points I want to make. In a similar vein, Velleman says of a case of two people considering a walk together,

I suggest that the statement "I will if you will"...means, "I hereby frame an effective intention that's conditional on your framing an effective intention as well" - that is, "I hereby *will* it,

³³ The idea of interdependent conditional intentions does not figure in Gilbert's recent work. Indeed, Gilbert has pointed out to me that even in the early work, she never intended the passage quoted as an analysis of joint or shared activity and the commitments therein. Rather, the passage gestures at a genetic account, a story concerning how these commitments come about. In recent work, Gilbert has emphasized the notion of joint commitment (1996, 330-331, and 1990 and 1999). Though she takes this to be conceptually basic or primitive, she does characterize it by saying that a party to a joint commitment cannot withdraw unilaterally from it. I have expressed my skepticism about this in note 24 above.

conditional on your willing likewise." And this statement just is the conditional willing that it describes itself as being...

Suppose that I say, "I'll go for a walk if you will," and you answer "Then I will." According to my analysis, each statement describes itself as an effective conditional willing, or intention, to take a walk;³⁴

I will take the intention-based commitment at the heart of the present proposal to be something like this: I am committed to activity A if or so long as the other participants are committed to A.³⁵ My commitment is thereby conditioned upon the commitments of the other individuals, and such a conditional commitment holds of each participant in shared activity.

This conditional commitment is understood to be the source or basis of the contralateral commitments. Thus, Gilbert holds that once the appropriate interdependent intentions have been expressed, the contralateral commitments are established:

The character of each one's commitment is then as follows: no one can release himself from the commitment; each is obligated to all the others for performance; each is (thus) entitled to performance from the rest.³⁶

³⁴ Velleman (1997, 45). To be fair, Velleman doesn't propose interdependent conditional intentions in order to account for contralateral commitments. Given his acknowledgement of Gilbert, this is a little surprising. The problem that animates Velleman is that of shared discretion. When an agent intends to F, this usually means that whether or not to F is up to the agent. In this sense, the agent has discretion over whether or not to F. F-ing is in his or her causal power. Velleman deploys interdependent conditional intentions to account for the sharing of discretion.

³⁵ Actually, Gilbert's proposal is that each intention is conditioned not only on the other individual's commitment, but also on its expression. Velleman's account involves intentions of the following form: I intend that we A, so long as you intend that we A. However, he sometimes speaks of the intention of form 'I intend to/will walk, if you will', which conditions on the act, rather than the intention. Ignoring these differences for now will not affect the points to be made below.

³⁶ Gilbert (1990, 8). Again, Gilbert is offering this as a genetic account, and not as an analysis of contralateral commitments. The criticism I go on to discuss will not apply directly to her understanding of contralateral commitments in terms of her notion of joint commitment. It does, however, apply to the story she gives in "Walking Together" about how this joint commitment comes about. Velleman seems to make use of interdependent conditional intentions as part of an analysis, rather than an

While I think it is on the right track to suggest that each participant's intention is dependent upon the intentions of his or her partners, an account of shared activity in terms of interdependent conditional intentions is ultimately unsatisfactory. First of all, it is not at all obvious how conditional intentions establish the contralateral commitments. If I were to form the conditional intention to go for a bike ride so long as it is sunny, it doesn't follow that when it is sunny, I somehow owe it to the weather – that I have a “commitment-to-the-weather” – to go for a ride. It's not even clear what this would mean. It doesn't seem to me that anything of significance would change were I to condition my bike ride on the behavior or intentions of other people. One's conditional intention to take surface streets if many people take the 405 does not generate in one any commitment to those many people on the 405 to take surface streets. Nothing intrinsic to the conditional intention itself generates any such commitment. So conditional intentions will not account for the contralateral commitments in shared agency.

There is a further important worry with the conditional intentions proposal of contralateral commitments: it fails even to establish the commitment each participant has to the shared goal, the participatory commitment which was the topic of section 1. If this charge is correct, then this would reinforce the conclusion that the conditional intention proposal will not account for contralateral commitments. For if, on this proposal, participants in shared activity are not even committed to a shared goal, then *a fortiori*, there is no commitment between the participants. Without any commitment to a goal, there can be no such contralateral commitments; the latter would be deprived of their subject matter.

To see why, on the conditional intention view, we lose our grip on the idea of each participant being committed to a goal, it is useful to keep in mind the case of an individual intention. Recall that to have an intention is to be settled on a goal. This means that the issue of whether

historical/genetic account. He says that participants in shared activity manage to share discretion,

in such a way that *their conditionally settling the issue separately adds up to their categorically settling it together...* [Gilbert] doesn't fully explain how individual exercises of conditional discretion *add up* to something that qualifies as a joint exercises of categorical discretion - how a collection of wills becomes one collective will. I want to clarify this feature of Gilbert's account...(Velleman 1997, 36; my emphasis)

or not to act on that goal is (defeasibly) settled. That's why it would be odd to have an agent who asserted something like:

- (i) I intend to A, but I'm deliberating about whether to A.

Part of the idea of intending some end is that you should no longer deliberate about whether to act toward that end. There is a (defeasible) practical norm of intending to the effect that when you have decided on some goal, you should move on to other issues (e.g. how to achieve it, or plan other activities so as not to conflict with it). The regular reconsideration of an intention violates this norm.

Not all forms of reconsideration entail that the agent actually re-opens deliberation on a matter. You might have an intention and yet consider whether to act on it, much in the way that you can wonder whether to trust your memory on some occasion. In the end, you might decide that your memory is reliable, and accept its deliverances. Similarly, you might realize that you have no reason to reconsider your prior intention to A, and decide to act on it without re-opening the issue of whether or not to A. Though this is all possible, the peculiarity of this situation is captured by the following:

- (ii) I intend to A, but I'm considering whether to act on this intention.

I believe that this, too, is a form of reconsideration that violates a practical norm of intending. There might be occasions where one should reconsider a decision, for example when one realizes that it may have been arrived at impetuously. However, someone who *globally* and *continuously* reconsiders his intentions in this way is not an intender; he simply has not really decided and settled on anything.

These examples bring out the idea that an intender qua intender can - indeed must (to a normatively significant extent) - take for granted that he is to act on his intentions. One's intention lingers, and needn't be re-considered in either of the manners described above. Typically, we don't need to *re-intend* what we've already intended in order to act on that prior intention. This is to say that once I've formed the intention to A, I don't typically re-open the issue of whether to A, nor do I deliberate about whether to re-open the matter.

In the case of conditional intentions one allows some external factors to finally settle the issue. For example, I might be settled on the goal of going for a bike ride so long as it's sunny. Conditional intentions

also involve being settled on a goal (e.g. of bike riding), *but only so long* as some condition is satisfied. If the condition is not satisfied, there is no commitment to the action – all bets are off. Nevertheless, the agent has to some important extent settled the issue of what he’s going to do. He’s committed himself to the action so long as some condition is in place, and he can go on to thinking about other issues.³⁷ He’s committed to some act F, on some condition C. That’s (normally) a substantive commitment, even if it is not a categorical commitment to F-ing. Thus it seems appropriate to describe the agent as in some significant sense having a goal, and intending, even if the full-blown commitment is not present.

But not all conditional intentions are quite so unproblematic. Consider the following:

(iii) I intend to A on condition that I intend to A.³⁸

Whatever is settled by such an intention, it is certainly not the issue of whether to A or not. Now consider the following:

(iv) I intend to A at noon on condition that I so intend (decide, will) at the time.

This intention does not settle the issue of whether you will A at noon. The important point here is that these sorts of intentions do not establish ANY kind of commitment to perform the relevant act. While my intention to bike if it is sunny commits me to a course of action if it's sunny, my intention to bike tomorrow so long as I intend to bike then involves no commitment on my part to any course of action. It remains completely up to me whether I will bike or not, and I violate no norms or commitments whichever option I choose. But if this sort of intention is pointless in virtue of not settling any issue, it is not really an intention. At least, no commitment characteristic of intention is formed when one

³⁷ Harman (e.g. 1986) emphasizes the importance in both theoretical and practical reasoning of settling certain issues so as to render manageable reasoning about other issues. Bratman (1987) speaks of intentions and plans in this way serving as filters for practical reasoning.

³⁸ It is possible to read this so that the intention in question is categorical or unconditional, where the condition noted isn't a part of the intention's content. On such a reading, (iii) is trivially true. However, the intended construal of (iii) is as a conditional intention wherein the condition specified figures as part of the content. On the intended reading, the content refers to the intention itself.

intends in this manner. My contention is that a problem like this afflicts the conditional intention proposal. In effect, this peculiar sort of conditional intention fails to satisfy a norm essential to or constitutive of intending.

Let's return to Gilbert's example of two individuals considering a walk together. Suppose that each of two walkers, Bob and Sue,³⁹ has a conditional intention to walk so long as the other intends likewise. That is, each forms a commitment to walk, conditional on the other's similar commitment to walk, thus:

- (iv) Bob is committed to walking, on condition that Sue is (similarly) committed.
- (v) Sue is committed to walking, on condition that Bob is (similarly) committed.

Along with some further standard conditions (e.g. common knowledge of these intentions), such conditionally interdependent intentions might be offered as an account of the pattern of commitments of shared intention. I will argue that this sort of proposal fails to generate the relevant commitment to a goal, let alone contralateral commitments.

Let's assume that Bob and Sue have the intentions in (iv) and (v). If Bob were suddenly to stop intending for no reason at all (or simply stop walking with Sue - e.g. start running way ahead), then we might say that he was violating a practical norm of intending, where the intention in question was his own. After all, his conditional intention says that so long as Sue is intending, Bob is committed to acting this way (or to continue intending). So it appears.

But on second glance, perhaps not. Suppose Bob suddenly changes his intention for no reason at all. You'd think that he thereby fails to live up to his commitment. But note then that Sue's intention, which is conditioned on Bob's, is no longer in effect; *she* is not committed because *Bob* is not committed. And since Bob's intention is conditioned on Sue's, then the norms that applied in virtue of his original intention are no longer in effect. That is, he is *not* committed to acting on his intention. But if no such intentions are in effect, there are no longer any commitments, shared or otherwise. So Bob would seem to have no contralateral commitment to Sue, so long as he simply stops intending.

³⁹ Things didn't work out between Sue and Jack, so she has moved on.

The reason there is no violation of the commitment arises from the peculiar interdependence of the conditional commitments. Whether or not your partner has the relevant commitment is *up to you*. Even if you have a commitment, if you suddenly cease intending, thereby failing to live up to your commitment, then your partner's intention is no longer in force. But if he's not so committed, then neither are you. So in failing to live up to your commitment, you're actually no longer committed, and in this sense, your failure is self-effacing.

But perhaps I've already granted too much to the conditional intention hypothesis. It's not even clear from the start that Bob has any commitment, and hence there's not even one for Bob to break. That's because his commitment is, in effect, conditioned on itself (by way of the conditioning on Sue's intention). But if Bob's intention is conditioned on itself, then (as we saw above in the discussion of (iii)) it is not really an intention, and in any case, lacks the commitment we need in order to account for shared agency.

Interdependent conditional intentions of the sort under discussion violate a practical norm constitutive of intention, and hence aren't really intentions. They fail to generate any of the sorts of commitment associated with shared agency (or, indeed, any commitments at all).

It might be replied that the proposal was not that one participant's commitment to A is conditioned on the others being so committed. Rather, it is conditioned on the others' *conditional* commitment. Turning to the walk shared by Bob and Sue, we have not (iv) and (v), but the following instead:

- (vi) Bob is committed to walking, on condition that (Sue has a conditional commitment to walking on condition that ...).
- (vii) Sue is committed to walking, on condition that (Bob has a conditional commitment to walking on condition that ...).

We saw that in (iv) we have Bob's commitment to walking being conditioned on Sue's commitment to walking. But since Sue's commitment is conditioned on Bob's (from (v)), Bob's commitment is in effect conditioned on itself, which is no commitment at all. The present proposal conditions Bob's commitment on a different commitment of Sue's, one that does not lead us directly back to his commitment. It is hoped that with this revised proposal we can avoid problematically conditioning an intention on itself.

But what, then, is this commitment of Sue's on this revised proposal? Well, it's not a commitment to walking, but a commitment to walking so long as Bob is committed to walking. That means that when we fill in the dots in (vi), we get:

(vi') Bob is committed to walking, on condition that (Sue is committed to walking on condition that Bob is).

Likewise, we have

(vii') Sue is committed to walking, on condition that (Bob is committed to walking on condition that Sue is).

In order for (vi') and (vii') to establish the requisite commitment of each individual to A, their respective conditions must hold. That is, we need the following for (vi'):⁴⁰

(a) Sue is committed to walking on condition that Bob is

and for (vii'):

(b) Bob is committed to walking on condition that Sue is

So, the revised view holds that when (vi') and (vii'), and (a) and (b) hold, we have in each agent the commitment to walking required for shared agency.

But notice that whether (a) and (b) hold depends upon whether Bob and Sue have formed the relevant conditional intentions. This is to say that (a) and (b) must be understood as practical commitments, and not simply conditional facts, independent of what the agent intends. This reflects the idea that when one shares agency with others, one thinks of one's partners as agents – and of their taking on of commitments as exercises of agency. When sharing agency, you do not rely upon the other agents as merely behaving or undergoing movements given certain

⁴⁰ If we interpret the conditionals (vi') and (vii') as *material* conditionals, then they might hold or be true even if the following conditions (a) and (b) are not satisfied. But (vi') and (vii'), and conditional intentions in general, should not be so interpreted. If I have only a conditional intention to go for a bike ride so long as it's sunny, then if the condition is satisfied, I'm committed to riding. But if it's not sunny, I can conclude that I'm not committed to going (so far as this intention is concerned). This latter conclusion cannot be drawn if we were to interpret the intention as a material conditional. For then either I'm committed to riding or it's not sunny. Here, when the condition goes unsatisfied (i.e. it's not sunny), I cannot conclude that I'm *not* committed to going for a bike ride. So it's wrong to interpret these conditional intentions in terms of material conditionals.

conditions. E.g. if (a) holds simply because Sue has a *compulsion* to walk so long as Bob does, and it is this upon which Bob is conditioning his intention to walk in (vi'), then they are not sharing agency. It is rather a form of manipulation. On the contrary, the conditional would have to hold because Sue has formed or *framed* the conditional intention.⁴¹ The question then arises as to whether you can frame the relevant conditional intention such as (a) or (b), without its ceasing to be an intention. And we have already seen from our earlier discussion of the initial conditional intention proposal that this is not the case (since the pair (a) and (b) correspond to the pair (iv) and (v)).

I'm suggesting that (a) violates the practical logic of intending and does not express a possible intention. Sue can't frame the conditional intention in (a) in the context of Bob's conditional intention in (b) (and their mutual knowledge). That's because Sue's intention is, in effect, conditioned on itself. We have already seen that this is an illegitimate form of intention, involving no commitment whatsoever.

But if (a) is an illegitimate form of intending, one that Sue cannot form, then the condition in (vi') does not hold. That means that Bob has no commitment to walking. A similar argument shows that the condition in (vii') doesn't hold, and hence that Sue has no commitment to walking. So the revised proposal that conditions the commitment on a conditional commitment will not work.

But how could interdependent conditional intentions fail to generate commitments? If you and I each have intentions of the form *I will A so long as you will (or intend to) A*, shouldn't we now form the corresponding categorical intention? Doesn't this show that such interdependent conditional intentions must embody a commitment to A-ing? The answer is no. When we have these conditional intentions, it makes at least as much sense NOT to form the categorical intention as it does to form it. If we have interdependent conditional intentions to shoot each other, should we each then form the categorical intention to shoot the other? One hopes not. Would it at least be *permissible* to form the categorical intention? This is doubtful in light of the shooting example. But even setting the example aside, the mere permissibility of forming the intention to A does not amount to any commitment to A-ing. For it is just as permissible not to form the intention, or even to form the intention *not* to A.⁴²

⁴¹ The expression is Velleman's (1997).

⁴² Nor does it help to add that each of us has a desire to A. For having such a desire is compatible with desiring not to A. Adding a desire, then, would hardly count as a commitment to A-ing.

Thus, the interdependent conditional intention proposal fails to capture the (participatory) commitment to the activity that made us want to appeal to the concept of intention in the first place. Notice, again, that this proposal will therefore not be able to handle contralateral commitments either, for on it participants have no commitment to a goal that can then be owed to any other individual in the sense of being a commitment-to that other individual.⁴³

5. Towards a solution: acting directly on another's intention

Our problem has been to apply the concept of intention to the interpersonal case so that we can capture the pattern of commitments characteristic of shared agency. So far two proposals along these lines have been considered, neither of them satisfactory. What has been particularly recalcitrant to analysis is the idea of contralateral commitment between participants in shared activity. This commitment seems to depend on an interdependence of intentions, but capturing this interdependence without undermining the participatory commitment to the goal is a non-trivial matter. I will take some steps toward addressing this problem by introducing an often-overlooked distinction regarding individual intentions, and then looking at an interpersonal analogue that might help with our problem.

I take it to be uncontroversial that it follows from

- (1) Agent acts on a prior intention to Φ

⁴³ Another proposal in terms of interdependent conditional intentions avoids the criticism developed in this section, but faces other serious problems. According to this proposal, the intention of one participant in the two person case would be something like:

Participant 1: I will (or am committed to) A so long as Participant 2 has a conditional intention with content like this.

The second participant would have the corresponding intention.

The problem with this proposal is that it fails to specify the content of these conditional intentions. In trying to articulate the conditional content of one intention, the proposal appeals to the content of the other's intention. But the content of the latter depends on that of the former. As a result, the content of neither is specified. As Velleman points out, "what I intend depends on what would count as a "like" intention on your part; and yet what counts as a "like" intention depends on what exactly I intend" (Velleman 1997, 45, n26).

The proposal might *seem* to succeed in specifying the intention contents, but this is only because we smuggle in ersatz content, pretend that the intentions are already there so to speak, and assume that there's a fact of the matter as to whether they're similar or not. All this before the contents have been specified, and hence before the intentions even exist. This is the criticism and diagnosis I would make of Velleman's attempted solution to this problem (1997, 1997a (Appendix)).

that

(2) at the time of action, Agent intends to Φ .

But we may distinguish between minimal and robust ways in which (2) can be true. Minimally, in this situation (where a prior intention has already been formed), (2) is no different from (1). There is no suggestion on this reading that the agent at the time of action is deliberating about Φ , or even that he considers *whether* to act on the intention. These issues have already been settled in the affirmative, and now the agent is simply acting on the intention he has already formed. In contrast, on the robust way that (2) might be satisfied, the issue of whether to Φ has not been settled for the agent. Perhaps the agent has taken up or resumed deliberating over whether to Φ , or whether to act on the prior intention to Φ . Or perhaps nothing so explicit is going on, and it's just that whatever authority the prior intention has for settling the matter of whether to \square is acquired well after it has been issued, deriving only from the agent *now, at the time of action* having the intention to F, or the intention to accord authority to that prior intention. In any case, what's important is that on the robust reading, the prior intention does not set a course of action for the agent. The issue is still open for the agent, and is only resolved for him at the time of action. On this robust reading, (2) amounts to *re-issuing the intention to Φ* - in effect, deciding anew or something close to it, as opposed to the simple *preservation of the prior intention to Φ* .

I do not think that (1) entails (2) understood as satisfied in the robust fashion just described. To insist that it does is, I believe, not to take seriously the idea that to intend involves being settled on a course of action. An intention that requires re-issuing does not really involve being settled or committed to a goal; it is not genuinely an intention. If after re-issuing the intention one finally does act, one is actually acting on the *new* intention - the re-intention - not the original intention. To insist that the agent must always re-issue her prior intention when she acts on it is to fail to appreciate the agent's own authority as an intender. As persisting agents who have some control over how we lead our lives, we plan for and commit ourselves to acts at some distance in the future. This requires that we have some level of authority and control over our acts that can be exercised over some temporal distance into the future. This control and

authority is constituted in part through our intentions for the future. But intentions will not afford us this control or authority if they must always be re-issued at the time of action.⁴⁴

Now consider the interpersonal case where someone, A, acts on an intention formed by someone else, B. For example, B intends for A to do some shopping. What I shall call Individualism holds that

(Individualism) The *only* way A can act on an intention formed by B is if A *re-issues* what B intended.

That's to say that if A is going to act on B's intention for A to go do some shopping, then A very well might have to consider reasons for shopping, and then decide in favor of shopping (or at least become settled on the matter by intending, where this needn't involve any conscious, explicit decision). Perhaps he might have to consider what could be said in favor of something more complicated, like doing some shopping *and thereby conforming to what B intends for him to do*. These reasons presumably figure in his deliberation, or would so figure were he to think about it some more. While the considerations that move A to go shopping might be the same as those that move B to intend for A to shop, they need not be. A could very well appeal to considerations that B would find unacceptable.

According to Individualism, B cannot settle the issue of whether A will go shopping. That's up to A. If A does re-issue the intention and if he then fails to act in accordance with the intention, then he's failing to live up to a commitment and the associated norms of intending that were invoked by *his own* authority – not B's – to intend and settle the issue. In this sense, A is only committed to himself. On this individualistic view, B has no authority over what A is to do, and A correspondingly has no commitment to B to act in accord with the intention.⁴⁵

⁴⁴ This sort of point is not limited to intentions formed for action at some distance in the future. Thus, when I act on a proximal intention (i.e. one to act here and now), or when I act with some intention (i.e. execute some intention in action), I should not continually re-issue the intention. The authority and commitment of the agent in intending is undermined when at the time of action, or throughout its course, the agent continually re-issues in the robust sense outlined above.

⁴⁵ The authority B lacks is that of the practical and rational sort associated with intending. Circumstances might be such that B possesses political or moral authority over A. But that's a different matter. (For the purposes of this paper, I set aside the possibility that political or moral authority is to be understood in terms of the kind of interpersonal intentions at issue here.)

This individualism entails that one cannot really act on another's intention, at least not directly. In re-issuing the intention, one is only acting on one's own intention; one has an ipsilateral commitment-to-oneself, and not a contralateral commitment-to-another.⁴⁶

I think that Individualism is mistaken in claiming that one can only act on another's intention by re-issuing it. When I form an intention to act in the future, the intention may simply be preserved in me to be acted upon later without being re-issued. And sometimes intentions may be formed by one person, and persist across an interpersonal divide (via communication) to be acted upon by another. Much in the manner in which one's own intentions create a (defeasible) goal and course of action for oneself, there are situations where someone else's intention might create a (defeasible) goal and course of action for oneself. So I endorse

(Practical Intimacy) It is possible for one individual to take up and to act on an intention formed by another without re-issuing the latter's intention.⁴⁷

Practical Intimacy suggests that one might *act directly* on someone else's intention, just as one might act directly on one's own prior intentions. I find the locution of 'acting directly' suggestive of the distinctive normative principle in play in shared agency wherein one is to act on some intention without re-issuing it.⁴⁸ However, it is important to rule out any possible misinterpretation that might be suggested by this locution.

In this regard, I want to make clear that the idea of acting directly on another's intention does not entail any sort of psychological "action at a distance". For example, when A acts directly on B's intention, I don't envision A somehow accomplishing this without having any knowledge or awareness whatsoever of B's intention. Presumably, B's intention needs to be *communicated* to A in some fashion, and A must thereby have some

⁴⁶ This was why forming an intention to act in accord with another's intention fails to generate a commitment to that other person, a point discussed earlier, in section 3.

⁴⁷ Thus there are circumstances or relationships where I am receptive to another's intention insofar as I need not exercise my own practical authority in order to act on that intention. As a result, I do not screen off the practical authority of the other individual. Such a scenario is, however, consistent with the prior exercise of my practical authority in establishing the appropriate relation with someone wherein I'm receptive to his or her intention, either on some particular occasion, or more generally.

⁴⁸ I have benefited here from conversation with David Kaplan.

knowledge or awareness of it.⁴⁹ A's acting directly on B's intention is thus conditioned on the intention being successfully transmitted from B and acquired by A. But this is just like the individual case wherein my acting directly on my prior decision/intention is conditioned on the successful transmission of the intention, usually via memory.

This communication condition does not suffice by itself to rule out all the worries about Practical Intimacy that might fall under the "action at a distance" heading. Consider a case where I act directly on Lisa's intention for me to go to the store. Suppose that the communication condition is satisfied in that she tells me to go to the store. Nevertheless, since I don't re-issue her intention, it might be thought that I don't have an intention to go to the store in any sense whatsoever. All I have is knowledge or belief concerning Lisa's intention. But then if I go shopping, I engage in some bit of intentional activity without it being true that I intend this activity or its end. The worry, then, is that if one can act directly on another's intention, there would be intentional activity with no relevant intention, or the intention is at some distance from my action, in the individual who authored it.

Many would regard this as a surprising and unfortunate result. It seems implausible to hold that intention and intentional action are so unrelated that one might intentionally F without it being the case that one intends F.⁵⁰ Moreover, if I am engaged in the intentional activity of going to the store, then there *must* be some psychological attitude of intending that plays the relevant role in my cognitive economy. That is, I am settled on this course of activity presumably because *I* intend it. Where

⁴⁹ This condition is meant to be pretty minimal. It does not imply that one necessarily knows whose intention one is acting upon. While certainly not the norm, it would appear that one might act on an intention that is not one's own, and yet not know who authored it. Perhaps one receives commands from a superior over a special phone line, without any identifying knowledge of who one's superior is. It might even be possible that one lacks any conscious awareness of acting on another's intention. An intention might be communicated to one without one's even knowing that it is another's intention. This might happen in some forms of advertising. It would be similar to subliminal suggestion, except that what's communicated is not a suggestion but something more decisive, corresponding to an intention. We do need to insist that the agent must in some manner *register* the intention. He can't simply act on the other's intention without even knowing in some fashion what that intention is.

⁵⁰ Of course, one might intentionally F without a *prior* intention to F. But that is not what is being envisioned here. Furthermore, in some cases, one might perform some act knowing full well that there is some side-effect S that will result. In such cases, one might say that one intentionally S's, but doesn't intend to S. But the case at hand is not supposed to be one involving side effects.

does this intention come from? It seems that I must have formed it, and that suggests that I have re-issued the intention.⁵¹

If acting directly on another's intention implies that there is intentional activity without intending, that would make it difficult to accept Practical Intimacy. Fortunately, the thesis need not have this consequence. I maintain that when an agent acts directly on another's intention (for him) to F, he *is* settled on and committed to a course of action. That is, we *should* be able to attribute to the agent the attitude of intending insofar as he intentionally F's and/or there is some attitude that plays the appropriate functional and normative role in the agent's planning and execution for bringing about F.⁵² I am not claiming that in this case the agent doesn't have any intention concerning the activity at all.

That being said, I want to insist that the way in which the agent comes to have this intention is different from the case where she acts on her own. The claim is that when I act directly on another's intention, I don't *re-issue* their intention. I am, rather, *preserving* and executing the intention of another. By this, I mean that I need not have deliberated over, weighed, or recognized reasons for doing this or that. If anyone had to have done so, it was this other individual, the intender. I didn't decide or settle the issue of what to do regarding this matter; the other individual did that. Her decision does not merely provide a strong *reason* for me to act as she intends for me to do, a reason that is supposed to enter into my deliberation about what to do. Rather, what we're envisioning is that I do not deliberate at all, or at least not enough to undermine this other individual's ability and authority to settle what it is that I'm to do.⁵³ If pressed on the reasons for the action, my ability to generate larger descriptions of the activity that exhibit its point or reason could very well peter out pretty quickly, and I might soon have to defer to the intender's descriptions of what's going on. Finally, she has the authority to correct

⁵¹ I owe this point to Tyler Burge, who argued that the position defended here should be formulated so as to avoid the sorts of problematic consequences just noted.

⁵² This underlies what I take to be the uncontroversial analogous thesis concerning individual intentions: when one acts on a prior intention to F, we would insist that at the time of action one intends to F.

⁵³ Of course, it's possible for me to make a contribution to the other's deliberation, for example by suggesting what my preferences are. Indeed, in paradigmatic instances of shared activity one could very well expect some such contribution. It is important, however, that one's contribution doesn't take the form of a necessary ratification or endorsement of the other's decision, without which the other's intention would not be in effect. For then the other will have failed to issue an intention – something that's supposed to settle what I'm to do.

what I'm doing, pointing out that I'm not doing what she intended for me to do. Thus, though I intend F, I lack the sort of authority regarding the intention I would have had had I formed it on my own. In sum, I act directly on the other's intention, without re-issuing it.⁵⁴

To ensure that Practical Intimacy will not be understood to entail a problematic psychological action at a distance, I will insist on the two auxiliary theses just considered: (i) the communication condition, which requires the transmission of the intention from intender to agent, and (ii) the thesis that the agent, in acting directly on another's intention to F, also intends to F (in the sense of preserving and executing the other's intention).

But how does one individual A come to be related to another B such that A will act directly on B's intention? This question might be taken genetically, as asking for the sorts of circumstances that *lead to* this

⁵⁴ In light of these remarks, it seems that the communication condition mentioned earlier should be understood as requiring not merely the transmission of some of content, but the transmission of the *attitude* of intending. I reject the idea that communication is necessarily limited to the transmission of contents. Such a view would imply that commands are fundamentally no different from mere suggestions, that they are presented to one as, in effect, something one is merely to entertain or consider doing, and not as something to be done. But this cannot be the right. Compare typical cases where I ask someone for the time. When someone answers me, her response is presented as true and as inducing in me all the relevant doxastic commitments. It is not presented as a mere proposition, something merely to be entertained and evaluated, but not to be believed until one has assessed the evidence for and against it. The answer to my question is thus presented to me to be believed straightaway. I think that what is going on in this sort of case is that the attitude itself - and not merely its content - is being communicated between persons.

Just as communication can involve the transmission of the attitude of belief over and above belief contents, so intentions and not merely their contents might be transmitted and preserved between individuals through communication. Though the communication of attitudes does not happen all the time (which is a good thing), it does happen, and this is a crucial element of shared activity. (The possibility that communication might preserve more than propositional content across an interpersonal divide is modeled on the distinction between evidential/substantive memory on the one hand, and preservative memory on the other. Memory preserves not only the proposition, but the attitude one has with respect to it. And this includes its justificatory status, and the sort of commitments involved. See Burge 1993.)

What if there is miscommunication? This certainly raises interesting questions. But similar questions are raised when someone mis-remembers or forgets his prior intention or what he was doing. The mere possibility of mis-remembering does not preclude my being able to act directly on my prior intentions. Likewise, the mere possibility of miscommunication does not preclude the possibility of acting directly on the intentions of others.

kind of relationship being established between A and B. So understood, a common way, I suppose, is for A simply to decide to recognize B as being knowledgeable about what he, A, should do, either generally or within a specific range of practical subject matter. Then, so long as B is willing to issue intentions for A to act, it seems that A will have opportunities to act directly on intentions formed by B.

But this sort of relationship between individuals is not always the result of some explicit decision. A could just fall into acting directly on what B intends for A to do. Children at certain ages very naturally act directly on each other's intentions without, it seems, explicitly deciding to enter into these sorts of relationships. Some may think it generally inadvisable simply to fall into this sort of relationship. But we seem to do it naturally, and often.⁵⁵

It might appear that the social phenomenon I'm gesturing toward is merely a regularity or tendency. This is not my intention. My concern is to account for a normative phenomenon - that of contralateral *commitment* - and I am not aiming to offer any reductive account of this in non-normative terms. So the question I've been asking - how does A come to be related to B so that A will act directly on B's intentions - might be put in more explicitly normative terms thus: How does B come to have the *authority* to settle what it is that A is to do?

I may allow that sometimes A's individual decision *leads to* a relationship with B such that the latter has the authority that underwrites A's acting directly on B's intentions. But this sort of genetic story does not entail that B's authority over A necessarily is to be *analyzed* individualistically, in terms of A's authority over him- or herself. The point is similar to the thought that altruism is not reduced to a form of self-interest just because it's possible that one might for selfish reasons become an altruist. Nevertheless, we may wonder whether the analysis is possible. That is, can B's authority to settle what A is to do be understood in terms of A's authority to form intentions to settle what he himself is to do?

A further but related normative question should be kept in mind. In the individual case, when I intend to F at t, I am settling what it is that

⁵⁵ The exact conditions under which individuals act directly on the intentions of others is a large topic that I do not try to address here. But I contend that it is a very common phenomenon, and something about which moral psychology, and social and other forms of empirical psychology will have something to say. I think that it's a safe bet that one will find this phenomenon associated with friendship, and more generally in social circumstances and organizations that involve any significant level of trust between individuals.

I am to do then. I am then *entitled* to the practical conclusion embodied in the intention. I am not merely settled in some causal or dispositional sense, for then my acting on the intention would turn into a mechanical affair and not count as full-blown agency. Thus, at the time of action, I am entitled to the intention in the sense that I am entitled to it as a practical judgment, a conclusion about what I am to do. Between the time I form the intention (presumably on sound practical grounds) and the time I act on it, the position of the intention in “justificational space” is maintained. My grounds for this practical conclusion are preserved, so that at the time of action, I do not have to re-think the matter, reconsider the alternatives, and re-produce my reasons. Now, something like this can hold in the interpersonal case as well. Thus, when A acts directly on B’s intention, A – at least in the ordinary case – will be entitled to B’s intention as a practical conclusion. So now, in addition to asking how B can come to have the authority to settle what it is that A is to do, we can ask the corresponding question: how does A come to be entitled to the practical conclusion embodied in B’s intention?⁵⁶

To sum up the story so far, the idea of acting directly on another’s intention preserves the authority of the intender, as well as the entitlement of the agent to the intender’s practical conclusion. But where does the authority or entitlement come from in the first place? This is a topic to be pursued in greater detail elsewhere; here I can only offer the following limited remarks.

On the most straightforward account, B’s authority over A is a matter of assent. A grants B authority, either in general, or within a certain range of practical subject matter. The granting of authority might take the form of an explicit assent or consent, or it might take a more implicit form. In the latter case, A needn’t explicitly decide to accord B authority, let alone outwardly assent to B as an authority. But the assent is implicit in what A thinks and does. Though he didn’t explicitly decide to grant B authority, there is implicit in his practical thought and action

⁵⁶ An ambitious thesis would be that we *generally* have an *a priori* though defeasible entitlement to such intentions of other individuals, and thus are generally able to act directly on the intentions of others. I do not take a stance on this thesis. For a defense of the analogous thesis in epistemology, see Burge (1993). Burge doesn’t look at *both* sides of the corresponding theoretical issue of testimony. That is, Burge only looks at the sort of *a priori* entitlement one has to believe the testimony of another. He does not consider what authority or right the other individual - the testifier - possesses to induce in the listener the theoretic/doxastic commitments of *a priori* belief or (*a priori* warrant for belief).

an intention to that effect.⁵⁷ So, whether explicitly or implicitly, B's authority on this view is granted to him by A. Then, when B intends for A to act in some manner, and B communicates that intention to A so that it is understood and recognized by him to lie within the relevant range of practical subject matter, A acts directly on the intention without deliberating and re-issuing the intention.

The assent-based story is a natural response to the normative issues stemming from the possibility of acting directly on another's intention. A's granting B the authority seems to be the most straightforward way for the latter to possess this authority to settle these issues for the former. But there may be some limitations to the scope of this account. There are occasions where A might act on B's intention without any prior opportunity to grant authority to B. Can one instead assent at the same time one is acting directly on the other's intention? This would imply that A grants authority to B *after* B has issued the intention, and *after* A knows of the issued intention. There are two worries with this. First, A cannot learn of B's intention and *then* grant authority, for B cannot issue this intention for A before having the authority to do so. Second, if A already knows what B has issued, it's hard to see how A can grant B the authority without the content of what's issued by B figuring in A's decision to grant authority. But if B's intention inevitably figures in A's granting of authority, then A is in effect making a decision about whether to act on it. B's intention therefore does *not* settle what A will do. A would not be acting directly on it. He would, it seems, be re-issuing B's intention, and thereby acting on his own authority, not that of B's.

A second reason to think that the assent-based model is limited becomes apparent when we recall that the flip side of the authority that B has over A is the *entitlement* that A has to B's intention as practical conclusion. One might wonder whether such an entitlement to a judgment about what to do can simply be a matter of assent. If the issue is merely whether or not B has authority over A, then A's simple assent to B's governing him would seem to address at least a large core of our concerns about this normative issue. Thus, if A accepts the idea of B telling him what to do, then the matter of authority plausibly may be regarded as settled. But A's assent will hardly suffice for his being *entitled* to the practical conclusions that are embodied in B's intentions for A to

⁵⁷ On this story, the assent involves A's intention to *grant* authority to B. The assent is thus the *source or basis* of B's authority. Contrast the sense in which one's assent to another is not so much a source or basis of the other's authority, but a recognition of the authority the other has independently of one's granting it.

act. The intention I get from you will not come with practical justification or warrant (and hence will not be something to which I'm entitled in a justificatory sense) if all that I've done is simply give you the okay to tell me what to do. What appears to be required in addition is that A's recognition of B's authority be warranted. So if the assent-based model is to capture all of the normative issues surrounding A's acting directly on B's intentions, the model will have to require A to make some judgment regarding B's competence and trustworthiness (or at least be entitled to such a judgment). Simple assent is not enough.

This complicates the assent-based model of acting directly on another's intention. The natural ways in which people immediately take up each other's intentions in social interaction make it implausible to think that there is explicit assent – let alone this further explicit judgment underlying entitlement to the intention as practical conclusion – in all instances where one individual acts on another's intention. Think of a situation where, at the site of an accident, a stranger tells another to call an ambulance. If the assent model is to handle such a case, it will have to make use of the idea of a background belief on the part of one individual concerning the competence of the other in order to justify the implicit granting of assent. I will not try to assess whether individuals generally make judgments regarding the competence of others and thereby have background beliefs to cover all instances where they act directly on another's intention. Whether the assent model should be accepted will depend on how well this is developed.⁵⁸ It will also depend on the availability of rival hypotheses.

⁵⁸ A further worry with the assent-based model stems from the fact that generally the authority one individual has over the intentions of another is not unlimited, but is instead restricted to a domain centered on some particular activity. The assent-based model would have to reflect this. For example, instead of assenting to you as an authority in general, I would assent to you as an authority to issue intentions concerning our day at the Museum of Modern Art. (Perhaps you're the expert on modern art, and it's your hometown, etc.) So I would act directly on your decisions regarding which galleries should have priority in being visited, and on where to eat for lunch. But when we get to the museum we find it ridiculously over-crowded. At this point, it would make a lot of sense to revise the overall plan of going to the museum. And given the position of authority you've recently been holding, you ought to be able to revise the overall plan, assuming you go about it in a reasonable manner. Notice, however, that on the assent-based story you are no longer in a position to do this. I've only granted you authority over intentions concerning our trip to the Museum of Modern Art, and we're no longer doing *that*. While not always the case, I submit that often the authority that underlies the issuing of an intention to X can extend to the issuing of an intention to Y if Y is a reasonable way to revise X in light of changing or unforeseen circumstances. But there is no way for the assent-based model to capture this possibility, for one cannot specify ahead of time – as the assent-based model must – all of the ways in which it might be

There *is* an alternative to the assent-based model of the authority and the coordinate entitlement involved in acting directly on another's intention. Consider the individual case. My prior intention, though defeasible, is nevertheless presented to me as to be acted upon straightaway, not as some *prima facie* consideration in favor of so acting. Thus my intentions settle the matter of what it is that I will do, and reflect some kind of authority I have over my future actions. But is there any reason to think that this authority is a matter of assent? Surely not. Certainly there is no explicit assent. It might be argued that there is an implicit form of assent involved here, though I'm not sure what that would amount to.⁵⁹ In any case, the simpler and more plausible view is that the authority one has over one's future actions is just a condition for being the sort of creatures that we are: creatures that persist over time, and engage in complex, long range planning and activity.

Similarly, there may be circumstances where other individuals present intentions to you to be acted on straightaway. If this indeed is the case, then it might be that there is a form of authority that one has over some of the actions of other individuals, and that this is not so much a matter of their assent, but of our natures as social beings routinely interacting with one another.

Of course, the sort of authority that B has over A is likely to be strictly limited, as would be the corresponding entitlement A has to B's practical conclusion. While one might as a default act on another's intention (or, more restrictedly, the intention of some *familiar* other), one might easily overturn such a default with a decision not to do so. And even if one doesn't make any such decision, the proffered intention is like one's own in being defeasible: if something comes up, I can override the other's intention for me to act and so not be saddled with any associated commitment. With these qualifications in place, we should be able to recognize a relationship that can hold between people such that one is liable to act directly on the other's intention.

Recall that simple assent is not enough to entitle one to another's practical conclusion; some sort of judgment of competence is required, at

reasonable to revise an intention in light of changing or unforeseen circumstances. There are some limits therefore to the assent-based model, at least as a comprehensive account of normative issues underlying acting directing on another's intention. I discuss this in greater detail in my 2003, where I consider a sophisticated form of individualism similar to the assent-based story discussed here.

⁵⁹ Note, again, that acting on a prior intention implies implicit assent to the intention. But that's not to say that the source of the authority of the prior intention *derives* from that assent. Rather, the assent is a reflection or recognition of the authority already in place.

least in the background. This was seen as a complicating factor for the assent model. It might be thought that similar complications beset the current proposal.⁶⁰ Notice, however, that the responsibility of the individual to grant authority in such a way that she will be entitled to the other's practical conclusions stems from the fact that the source of the other's authority lies with the individual's assent. But if the authority is not based on the individual's assent, then he or she will not have the same sort of responsibility for assessing competence. In particular, in order to be entitled to the other's practical conclusions, one need not formulate or justify a judgment of competence unless, of course, one has good reason to think the other to be incompetent, or else the issue of their competence is legitimately raised.

I believe that there is an analogy between how my own intentions present themselves to me to be acted upon, and how someone else's intentions for me to act can present themselves to me. In the individual case, the agent can settle what it is that he will do, and he then acts directly on his intention and is entitled to the practical conclusion. Underlying this is an authority over oneself and an entitlement to one's own prior practical conclusion. Analogously, one individual can settle what another will do, and the latter acts directly on this intention, and is entitled to it as a practical conclusion. On the assent model, the authority and entitlement is based on assent and judgments of competence. On the alternative model, the analogy is tighter. My authority over myself is not based on assent, nor is my entitlement to my own intentions based on a judgment of competence. Likewise, the intender has some sort of authority over the agent, and the agent has an entitlement to the intender's practical reasoning, and neither assent nor judgment about competence or trustworthiness need have anything to do with this.⁶¹

I have noted that the idea of acting directly on another's intention raises the normative issues concerning the authority to settle what someone else is to do, and the coordinate entitlement to that practical conclusion. I have also sketched a couple of approaches to answering these issues: the assent-based model, and the approach that sees this sort of relation as being a condition for being the sort of social agents that we are. I do not pretend to have given anything like a complete account of

⁶⁰ This worry was raised by one of the anonymous referees.

⁶¹ Of course, things are different when there is some reason to doubt the practical conclusions of the intender. Then again, this is also true in the individual case.

these normative issues.⁶² But I think I've said enough to indicate what I have in mind when one acts directly on another's intention.

6. Contralateral commitments: their source

The idea of acting on someone else's intention without re-issuing it – understood so as not to entail a problematic psychological action at a distance – suggests a promising way of capturing the contralateral commitments which have so far eluded satisfactory philosophical treatment.

Recall that an agent has some authority over how to act so that when she forms an intention to F, she has a commitment to herself. Now, suppose I take on someone else's intention – say, yours – without re-issuing it. This means that you had the practical authority to settle what it is that I'm to do (otherwise, the only way for me to act on your intention would be by re-issuing it).⁶³ In acting directly on your intention, moreover, your authority was not somehow superseded or suspended (as it might have been had I re-issued the intention). If I then fail to act in accord with the intention, I let you down. So in this sense, just as my own intention (prior or otherwise) commits me to some end and underlies a commitment to myself, so someone else's intention (in appropriate circumstances) commits me to an end and is the basis of my commitment *to that person* to act in accord with that intention.

In light of these remarks, I suggest the following. One individual, A, has a contralateral commitment to another, B, if and only if:

- (i) B has the authority to settle what it is that A will do with respect to the matter of whether or not to F, and A has the corresponding entitlement to B's intention as a practical

⁶² Notice, however, that similar normative issues can also be raised concerning individual intentions. It is generally taken for granted that an individual has the authority to settle what she is to do in the future. To insist that interpersonal authority requires some *further* justification or defense beyond the sort required of individual intentions is to buy into an individualistic theoretical bias. Instead, perhaps what is needed is a special argument to impose certain restrictions on interpersonal authority so as to establish a domain of individual or personal autonomy.

⁶³ Keep in mind that your authority is limited. You can settle what I do, but only within a certain range. If you form a crazy intention, this will trigger me to reconsider. But the default is that I do not.

It won't make sense for you to issue an intention for me to act upon unless you can expect to settle the matter for me. Barring some general entitlement to settle what another is to do, you would most likely appeal to background knowledge of our relationship or circumstances to justify the authority implicit in issuing this sort of intention to me.

- conclusion. Moreover, this authority is exercised, and B intends for A to F;
- (ii) this intention is communicated to A;
 - (iii) B's authority is not blocked, superseded, suspended, or nullified, so that A is in a position to act directly on the intention for him to F, without re-issuing it.⁶⁴

Take the example of the walk again. Our problem was to see how Bob could have an intention-based contralateral commitment to Sue. We can see now that what is needed is that Sue's intention be of the sort for Bob to act upon directly. Sue might intend for Bob to walk with her, or for them to walk together. Either way, Bob figures as (at least part of) what might be called the executive subject of the intention. The executive subject is to carry out the intention issued by what we might call the authoritative subject of the intention, in this case Sue. Bob takes up Sue's intention and acts on it as he would for intentions he has formed himself – that is, without re-issuing it. Of course, there is some attitude of intending that plays a role in Bob's cognitive economy. He is settled on walking with Sue, and so we speak of Bob intending to walk with Sue. But in an important sense, the intention in virtue of which Bob is thus settled is not fully his own. It is, rather, an intention that has been preserved, one that originated in the exercise of Sue's authority to settle what Bob is to do in this matter. Sue's intention has a practical and functional role in Bob's cognitive economy. At no point in the process of intention transmission does Bob re-issue Sue's intention to make it exclusively his own. Were he to fail to act in accord with this intention, he would be letting Sue down – just as he would have been letting himself down had he failed to act on his own intention. And so, just as I have an ipsilateral commitment-to-myself to act in accord with my intentions, Bob has a contralateral commitment-to-Sue to act in accord with the intention for them to walk together.

But the account so far seems more appropriate for cases where only one participant – in this case Sue – has the authority to command

⁶⁴ B's authority can be superseded by a higher authority. B's authority is suspended or nullified when for some reason B ceases to satisfy some necessary condition for being an authority (compare cases where an individual becomes temporarily or permanently incapable of making decisions for himself). B's authority is blocked when A's condition or new circumstances rule out any resolution of practical matters that fall within the scope of B's authority. B's authority may also be blocked when changing circumstances defeat B's practical conclusion/intention at a time when he is unavailable to revise it. This list is not exhaustive.

the rest and thereby settle how they will act. Our story seems to describe a more hierarchical social arrangement than we would have expected in shared agency. Perhaps this is not how we imagined Bob and Sue to be related. More to the point, it seems that nothing in our account so far entails that Sue has any corresponding commitment to Bob to do her share of the activity. We have been focusing on Bob's commitment to Sue probably because he has been misbehaving in walking far ahead of her. But Sue has, for example, a commitment to Bob not to lag so far behind when she is quite capable of walking faster. Where does her contralateral commitment-to-Bob come from?

On our account so far, Sue presumably acts directly on the very intention that Bob acts upon: the intention for them to walk, which she herself has issued. But this only accounts for Sue's ipsilateral commitment to herself to do her share of the activity. We have yet to account for her commitment to Bob.

One suggestion is to have Sue act directly on an intention issued by Bob. The worry with this proposal is that Bob's intention would then be distinct from that issued by Sue, and hence their respective commitments would be independent of one another. That does not seem right for shared agency. Whatever the problems with the conditional intentions proposal discussed earlier, it was right in seeking to reflect the interdependence of these commitments.

It turns out, however, that there is no difficulty in accounting for Sue's reciprocal commitments to Bob. When Bob acts directly on Sue's intention, we have the basis not only for his contralateral commitment to her, but also her commitment to him. We've seen that his commitment to her is based on her practical *authority* to settle what it is that he will do, an authority which is a precondition for his being able to act directly on her intention. But recall that a further precondition is that he have an *entitlement* to her intention as a practical conclusion about what to do. I contend that this entitlement is the foundation for Sue's symmetrical contralateral commitment to Bob. Let me explain.

I speak of Bob *acting* directly on Sue's intention. He isn't merely caused to move in some way, but is exercising intentional agency. And this means that he must have some reason.⁶⁵ But this doesn't mean that he must have reasons in quite the way he normally does when he is acting on his own intention. The reasons that figure in Sue's decision, and which justify and rationalize what Bob does, need not be available to Bob when he acts directly on her intention. He is nevertheless entitled to those

⁶⁵ Davidson (1963).

reasons, and to the practical conclusion embodied in her intention, and it is in virtue of this entitlement that he *acts* on her intention.

Given Bob's entitlement, Sue has special responsibilities to Bob. If she acts in a way that undermines the practical conclusions of hers to which Bob is entitled, then Sue would be letting him down. In our case, her conclusion was for them to walk together. A condition for the success of this intention is that she not lag so far behind as to make it impossible for them to walk together. Were she to lag and ensure the failure of the intention, she would undermine Bob's entitlement to her conclusion – something for which she's responsible. She would thus be letting him down. Analogously, suppose that today I form an intention to □ tomorrow. I thereby settle what I'm to do; when it comes time to act, I can take the matter as settled, and am entitled to yesterday's conclusion. But what if, around when I issue this intention to □□□□□ also lay the groundwork for ensuring that I not succeed at □-ing. (Perhaps this is because of my forgetfulness or sloppiness, or perhaps more perversely it is a matter of self-deception.) Then I would be letting myself down by failing to be responsible for something to which I should later be entitled.⁶⁶ It is in a sense analogous to this that Sue has a commitment-to-Bob to do her share of the activity.

So Bob has a commitment to Sue, and vice versa. But there may be some residual concerns about asymmetry. It might be suggested that Bob's commitment to Sue has a somewhat different basis than Sue's commitment to Bob. The former stems from Sue's authority, and the latter from Bob's entitlement. Moreover, if they are sharing agency, shouldn't he have as much authority as she?

Before I try to allay these worries, I would like to point out that it's not at all obvious that *all* instances of shared agency actually exhibit the sort of symmetry posited by these complaints. The sharing of agency amongst individuals should not exclude the possibility that one or some of the participants take a lead, or exhibit agency more robustly, or have a sort of authority not had by the rest who nevertheless remain genuine participants in the activity. So the proposal being defended here should not be faulted if it allows that shared agency may exhibit certain forms of asymmetry, so long as it can also handle cases exhibiting greater symmetry. Indeed, it is a strength of the account that can handle the asymmetric cases.

What, then, can the account say about the more symmetrical cases? First, notice that Sue's authority is balanced by Bob's

⁶⁶ This is then the inverse, so to speak, of the ipsilateral commitment one has to oneself to act on one's prior intentions.

corresponding entitlement to Sue's practical judgment, and that the authority and entitlement are not independent of one another. So there is a unity that underlies what is a relatively superficial asymmetry. Second, and perhaps more significantly, in many instances of shared agency, authority is had by *each* of the participants, and it was just a matter of chance that it was exercised by one and not the other. If our example is such a case, then Bob just as well as Sue could have issued the intention. It can be allowed that amongst the variety of forms of shared agency, a common (perhaps even paradigmatic) form features *mutual* authority.

But might there be a problem with the notion of mutual authority?⁶⁷ If Sue can settle what she and Bob will do, does this not exclude Bob from having this authority? Not necessarily. Sue's authority consists in the fact that she can issue an intention and count on Bob to act directly on it. Bob's authority likewise consists in his being able to issue a similar intention for Sue. Suppose Sue exercises her authority. It doesn't follow that she wouldn't have acted directly on an intention issued by Bob. Thus Sue's authority and its exercise are compatible with Bob having a similar authority.

Nor is there a problem if both were to issue incompatible intentions. Compare the case of conflicting intentions in the individual case. It's quite possible for me on Monday to form an intention to attend a meeting on campus on Friday, and for me on Tuesday to decide to spend all day Friday at the beach. Come Thursday, I realize that my intentions for what to do the next day conflict with one another. This possibility of conflict does not entail that on one or both of Monday and Tuesday I *lack* the authority to issue intentions regarding what I'll do on Friday. In situations like this, one seeks some sort of rational resolution between the intentions, or the rational revision of one or both of them. One does not dismiss one of the intentions as illegitimate, as having been issued without the requisite authority. This would not be a way to seek consistency amongst one's intentions, but to assert – bizarrely – that there really wasn't any inconsistency in the first place because one of the intentions wasn't really an intention, or wasn't one's own intention. Returning to the interpersonal case, suppose that Bob finds that Sue has issued a conflicting intention. What is then called for is some sort of resolution involving the rational revision of intentions. In particular, it's not simply open to Bob to dismiss Sue's intention as illegitimate, as having been issued by someone who has no authority over what he is to do in this area, and hence as subject to circumvention or undermining. So long as

⁶⁷ These concerns are related to those raised by Velleman (1997) under the heading of 'shared discretion'.

this sort of rational recourse is available, and it is not simply open to one to circumvent or undermine a conflicting intention issued by another authority, then the fact that conflicting intentions might be issued shouldn't preclude the possibility of mutual authority.⁶⁸ Thus, there does not seem to be any reason to stop us from making use of the notion of mutual authority to describe the more symmetric instances of shared activity.

There remains a further asymmetry: Sue *exercised* her authority, and Bob did not. But such an asymmetry need not be symptomatic of any important difference in the authority possessed by each participant. Nor need it indicate any difference in their more general social status. Thus, suppose you and I are thinking about what to do together for the afternoon. We've considered shopping, the beach, and the museum. The malls are over-crowded, and we've already had a bit too much sun from a hike in the morning. But there won't be crowds at the Museum of Contemporary Art, even though there has just opened a new exhibit we have both expressed a keen interest in seeing. Neither of us can think of any other options. Now, in such a context, it could be just a matter of chance who it is that finally goes ahead and exercises the authority that each evidently has. Although there may well be a fact of the matter as to who takes this first step, I don't think that my account can be faulted for attributing such an innocuous asymmetry to cases of shared activity of individuals of similar status and authority.⁶⁹

Concluding remarks

When individuals engage in shared activity, each has a participatory commitment to that activity. I understand this commitment to be intention-based, so that in the case of Bob and Sue, each has an intention that might be expressed as "We are walking together". This intention and its preconditions also figure in our account of the contralateral commitments. It so happens that in this case, Sue originally issued this intention, and Bob acts directly on it. A precondition for this is that Sue has the authority to settle what it is that Bob will do, and Bob is entitled to her intention as a practical conclusion. This authority underlies Bob's contralateral commitment to Sue, and this entitlement

⁶⁸ These issues receive more extended discussion in my 2003.

⁶⁹ Individuals regularly acting together would probably do well not to let there be too much asymmetry in the exercise of authority. Otherwise, passive partners could become resentful of active partners as bossy and overly assertive know-it-alls, while the active become irritated at all the planning and decision-making the passive have left for them to do.

correspondingly underlies Sue's contralateral commitment to Bob. Cases exhibiting more symmetry will involve mutual authority.

There are a host of issues yet to be addressed. More needs to be said about just when it is that individuals are related so that the appropriate relations of authority and entitlement are in place. What kind of trust is necessary for the practical intimacy that allows individuals to act directly on each other's intentions? There may be the suspicion that the sort of trust needed is rarely if ever available. It must be conceded that there are very many cases where it is not advisable to act on an intention formed by someone else. Some individuals may not be in a good position to know what your interests are, and others simply should not be trusted to have your interests in mind. Social and cultural circumstances might become so unfortunate that the wise default position would be to refuse all proffered intentions. However, it would be overly pessimistic to think that circumstances *always* preclude the possibility of trust between people who are less than fully acquainted with one another. Often others can be trusted, or the stakes are simply not so high so as to warrant suspicion. There might be cases where someone else could do a better job at deliberating with your interests in mind than you yourself could. And even if in some circumstances, because of a lack of knowledge, someone is not in a good position to intend for you, she can get into a better position simply by asking you what you'd like to do.

Shared agency requires some disposition to act directly on the intentions formed by others. If we are social creatures who share agency, then we do take up the intentions of others in this manner. We might imagine a spectrum of psychological tendencies to act directly on the intentions formed by others, and different types of social circumstances might make some tendencies more reasonable to have than others. At one extreme there is a gullible fool, and at the other there is someone who is so utterly suspicious that they cannot connect with others in any normal way. Those who would insist that we never directly act on the intentions formed by others are very implausibly taking a psychological extreme as the norm.

In any case, if this concern based on a lack of trust has any force, consider how a similar objection might be formulated against the idea of acting on one's own prior intentions. Can one trust one's memory of what one has intended? Can one now rely on the soundness of decisions one has made in the past? Do such questions suggest that we should never act directly on intentions we've formed in the past? No. These questions only suggest that the norm of acting directly on one's past intentions is defeasible. Sometimes you shouldn't trust your memory of

what you intended, and sometimes you shouldn't trust your past self, which was young, or foolish, or drunk, etc. But hopefully one is not always too young, too foolish, or too drunk to make decisions about one's future.

I hope to have said enough about acting directly on another's intention to have shown that this idea must be taken seriously and that it makes for an elegant solution to the problem of accounting for the contralateral commitments of shared activity. Moreover, the theory is, I hope, at least intriguing in that its view of intentions and practical reasoning is substantively different from the prevailing individualism. Its implications need to be explored further.

I complained at the outset that even on a sophisticated form of individualism featuring interdependent conditional intentions, the suggested picture of social interaction is one of delicate negotiation between wary individuals. While diplomacy is called for all too often in human affairs, it should not serve as a theoretical model for shared agency in general. As I have said, most walks are not negotiated endings of hostilities. A majority of walks seem to come about rather easily. How? Often it simply amounts to someone saying "Let's go for a walk", and (barring conflicting commitments) we all act directly on the intention and go for a walk.⁷⁰

Bibliography

- Anscombe, G.E.M. 1963. *Intention*, 2nd ed. Ithaca, NY: Cornell University Press.
- 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
- 1997a. Shared Intention and Mutual Obligation. Originally published as "Intention Partagée et Obligation Mutuelle", in *Les limites de la rationalité*, Vol. 1, edited by Jean-Pierre Dupuy and Pierre Livet, 246-66. Paris: Editions La Découverte. Reprinted in Bratman 1999.
- 1997b. 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*.

⁷⁰ Apologies to David Velleman.

- 1999. *Faces of Intention*. Cambridge: Cambridge University Press.
- Burge, Tyler. 1993. Content Preservation. *Philosophical Review* 102:457-88.
- Davidson, Donald. 1964. Actions, Reasons, and Causes. *Journal of Philosophy* 60: 685-700.
- 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.
- 1993. Is an Agreement an Exchange of Promises? *Journal of Philosophy* 90:627-49. Reprinted in Gilbert 1996.
- 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.
- Grice, H. P. 1971. Intention and Uncertainty. *Proceedings of the British Academy* 57:263-79.
- Harman, Gilbert. 1976. Practical Reason. *Review of Metaphysics* 79:431-63.
- 1986. *Change in View*. Cambridge: MIT Press.
- Kavka, Gregory. 1983. The Toxin Puzzle. *Analysis* 43:33-6.
- Korsgaard, Christine. 1993. The Reasons We Can Share. *Social Philosophy and Policy* 10(1): 24-51.
- Kutz, Christopher. 2000. Acting Together. *Philosophy and Phenomenological Research* 61:1-31.
- Mele, Alfred. 1992. *Springs of Action*. New York: Oxford University Press.
- Raz, Joseph. 1975. *Practical Reason and Norms*. London: Hutchinson.
- Roth, Abraham Sesshu. 2003. Practical Intersubjectivity. In *Socializing Metaphysics: the Nature of Social Reality*, edited by Frederick Schmitt, 65-91. Lanham, MD: Rowman and Littlefield.
- Robins, Michael. 1984 *Promising, Intending and Moral Autonomy*.
Cambridge: Cambridge University Press.
- Scanlon, Thomas. 1990. Promises and Practices. *Philosophy and Public Affairs* 19:199-226.

- 1998. *What We Owe to Each Other*. Cambridge: Harvard University Press.
- 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.
- Sellars, Wilfrid. 1963. Imperatives, Intentions, and the Logic of 'Ought.' In *Morality and the Language of Conduct*, edited by G. Nakhnikian and H.-N. Castañeda. Detroit: Wayne State University Press.
- 1968. *Science and Metaphysics*. London: Routledge & Kegan Paul.
- Scheffler, Samuel. 1997. Relationships and Responsibilities. *Philosophy and Public Affairs* 26: 189-209.
- Tuomela, Raimo. 1991. We Will Do It: An Analysis of Group-Intentions. *Philosophy and Phenomenological Research* 51: 249-77.
- Tuomela, Raimo and Kaarlo Miller. 1988. We-Intentions. *Philosophical Studies* 53: 367-89.
- Velleman, J. David. 1997. How to Share an Intention. *Philosophy and Phenomenological Research* 57:29-50.
- 1997a . Deciding How to Decide. In *Ethics and Practical Reasoning*, edited by Garrett Cullity and Berys Gaut, 29-52. Oxford: Oxford University Press.
- Wilson, George. 1989. *The Intentionality of Human Action*. Stanford: Stanford University Press.