

Shared Intention

Author(s): Michael E. Bratman

Source: *Ethics*, Vol. 104, No. 1 (Oct., 1993), pp. 97-113

Published by: The University of Chicago Press

Stable URL: <https://www.jstor.org/stable/2381695>

Accessed: 20-10-2018 15:19 UTC

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



JSTOR

The University of Chicago Press is collaborating with JSTOR to digitize, preserve and extend access to *Ethics*

Shared Intention*

Michael E. Bratman

In *Choice: The Essential Element in Human Action* Alan Donagan argued for the importance of “will” to our shared understanding of intelligent action.¹ By “will” Donagan meant a complex of capacities for forming, changing, retaining, and sometimes abandoning our choices and intentions. (Choice is, for Donagan, a “determinate variety of intending.”)² Our capacity to intend is to be distinguished both from our capacity to believe and from our capacity to be moved by desires. And Donagan thought that intentions involve what, following Austin, he called “‘as it were’ plans.”³

I am broadly in agreement with these main themes in Donagan’s book, and I will pretty much take them for granted in what follows.⁴ I will suppose that intention is a distinctive attitude, not to be reduced to ordinary desires and beliefs; that intentions are central to our shared understanding of ourselves as intelligent agents; and that “the study

* Thanks to Margaret Gilbert and Raimo Tuomela, thoughtful commentators on presentations of earlier, shorter versions of this article. Thanks also to Philip Clark, Rachel Cohon, Fred Dretske, David Hilbert, Henry Richardson, and Debra Satz for their useful philosophical advice. Barbara Herman and David Velleman provided rich and probing comments when this article was presented at the September 1992 Memorial Conference in Honor of Alan Donagan, held at the University of Chicago. Some of the issues they raised are discussed further in my “Shared Intention and Mutual Obligation” (presented at the Pacific Division American Philosophical Association, San Francisco, March 1993). Work on this article was supported in part by the Center for the Study of Language and Information, made possible in part through an award from the System Development Foundation.

1. Alan Donagan, *Choice: The Essential Element in Human Action* (London: Routledge & Kegan Paul, 1987).

2. *Ibid.*, p. 97.

3. *Ibid.*, p. 96.

4. I developed ideas that are in some respects similar to Donagan’s themes in my *Intention, Plans, and Practical Reason* (Cambridge, Mass.: Harvard University Press, 1987). Of course, there are various differences in our views. Donagan discusses one of these—concerning the consistency demands to which intentions are subject—in *Choice*, pp. 98–105. My detailed treatment of choice differs in certain ways from Donagan’s (see *Intention, Plans and Practical Reason*, chap. 10). And there are other differences as well. But these differences are not relevant here.

Ethics 104 (October 1993): 97–113

© 1993 by The University of Chicago. All rights reserved. 0014-1704/94/0401-8731\$01.00

of intention" is in part the "study of planning."⁵ My hope is that these common elements in our views about intention can serve as a basis for reflection on the phenomenon of shared intention.

I

That we do sometimes have intentions that are in an important sense shared seems clear. We commonly report or express such shared intentions by speaking of what *we* intend or of what *we* are going to do or are doing. Speaking for you and myself I might say that we intend to paint the house together, to sing a duet together; and I might say that we are going to New York together. In each case I report or express a shared intention.

Sometimes we speak of the intentions of structured social groups: the Philosophy Department, for example, intends to strengthen its undergraduate program. But some shared intentions are not embedded in such institutional structures. These will be my main concern here: I will focus on cases of shared intention that involve only a pair of agents and do not depend on such institutional structures and authority relations. Supposing, for example, that you and I have a shared intention to paint the house together, I want to know in what that shared intention consists.⁶

On the one hand, it is clearly not enough for a shared intention to paint the house together that each intends to paint the house. Such coincident intentions do not even insure that each knows of the other's intention or that each is appropriately committed to the joint activity itself. On the other hand, a shared intention is not an attitude in the mind of some superagent consisting literally of some fusion of the two agents. There is no single mind which is the fusion of your mind and mine.

Now, one way in which you and I may arrive at a shared intention is to make an appropriate, explicit promise to each other. But such promises do not ensure a shared intention, for one or both parties may be insincere and have no intention to fulfill the promise. Nor are explicit promises necessary for shared intentions. Consider Hume's example of two people in a row boat who row together "tho' they have

5. Donagan, *Choice*, p. 95.

6. There is a recent literature in artificial intelligence that focuses on similar issues. See, e.g., Philip R. Cohen and Hector J. Levesque, "Teamwork," *Nous* 25 (1991): 487–512; Barbara J. Grosz and Candace L. Sidner, "Plans for Discourse"; and Jerry Hobbs, "Artificial Intelligence and Collective Intentionality: Comments on Searle and on Grosz and Sidner." The latter two essays are in *Intentions in Communication*, ed. Philip R. Cohen, Jerry Morgan, and Martha E. Pollack (Cambridge, Mass.: MIT Press, 1990), pp. 417–44 and pp. 445–59, respectively.

never given promises to each other.”⁷ Such rowers may well have a shared intention to row the boat together.

To understand shared intention, then, we should not appeal to an attitude in the mind of some superagent; nor should we assume that shared intentions are always grounded in prior promises. My conjecture is that we should, instead, understand shared intention, in the basic case, as a state of affairs consisting primarily of appropriate attitudes of each individual participant and their interrelations.⁸

How do we determine in what this complex of attitudes consists? Begin with a related query: What do shared intentions do, what jobs do they have in our lives? I think we can identify three main answers to this query.

First, our shared intention to paint together will help coordinate my activities with yours (and yours with mine) in ways that track the goal of our painting the house. Someone will scrape before, not after, the new paint is applied by someone. Second, our shared intention will coordinate our actions in part by ensuring that my planning about my role in the house-painting is coordinated with your relevant planning, and vice versa. If I plan to get the paint but not the brushes I will likely check whether you plan to get the brushes. Third, our shared intention will tend to provide a background framework that structures relevant bargaining. Though we share the intention to paint together we might have conflicting preferences about who scrapes and who paints, or about what color paint to use. Such conflicts call for bargaining in some form—not bargaining about whether to paint together but, rather, bargaining about how we are to paint together.

Our shared intention, then, performs at least three interrelated jobs: it helps coordinate our intentional actions; it helps coordinate our planning; and it can structure relevant bargaining. And it does all this in ways that track the goal of our painting the house together. Thus does our shared intention help to organize and to unify our intentional agency in ways to some extent analogous to the ways in which the intentions of an individual organize and unify her individual agency over time. An account of what shared intention is should explain how it does all this.

So what we want to know is this: Are there attitudes of each of the individual agents—attitudes that have appropriate contents and

7. David Hume, *A Treatise of Human Nature*, ed. L. A. Selby-Bigge (Oxford: Oxford University Press), p. 490. See David Lewis's remarks about this example in his *Convention: A Philosophical Study* (Cambridge, Mass.: Harvard University Press, 1969), p. 44, e.g.

8. Let me explain why I say only “primarily.” I claim below that shared intentions involve “common knowledge.” I do not try here to say what common knowledge is. But it may be that it involves some external situation in the environment of the agents that functions as what Lewis calls a “basis for common knowledge” (p. 56).

are interrelated in appropriate ways—such that the complex consisting of such attitudes would, if functioning properly, do the jobs of shared intention? Can we describe an appropriate complex from whose proper functioning would emerge the coordinated action and planning, and the relevant framework for bargaining, characteristic of shared intention? If so, we would have reason to identify shared intention with this complex.

II

Such an approach to shared intention will need to draw on an understanding of the intentions of individuals, with special attention to the roles of such intentions in coordination. Here I briefly sketch an approach to such matters that I have developed elsewhere.⁹

Suppose I intend now to practice the tenor part tomorrow at noon. If all goes well my activity between now and then will include all necessary preliminary steps—for example, getting the music if I don't already have it—and it won't include activity incompatible with my practicing then—for example, screaming too much at an athletic event the night before. And when tomorrow noon arrives I will be in a position to practice; I will not be, say, attending a movie. This normally happens, if it does happen, because of my intention. My intention to practice my part tomorrow coordinates my activity between now and then in a way that supports my practicing at noon.

How does my intention play this coordinating role? In part, by shaping my planning between now and later. My intention to practice is an element of a partial plan. As time goes by I need to fill in this plan appropriately; otherwise it will suffer from means-end incoherence. So my intention poses relatively specific problems of means and preliminary steps for my planning. I am faced, for example, with a problem about how to get a copy of the tenor part by noon. In contrast, my plan poses no special problem about how to get a copy of *The Iliad*, even if I would much like one. Further, my intention constrains my plans in ways necessary to ensure that my plans remain internally consistent and consistent with my beliefs: for example, it precludes going to a movie tomorrow at noon. In these ways my intention helps insure that my activities between now and tomorrow are coordinated with each other in ways that support my practicing then.

For all this to work my intention will need to have a further property. Prior intentions are revocable. If things change in relevant ways it may behoove me to change my plan. Still, prior intentions will need to have a certain stability.¹⁰ If we were constantly reconsidering

9. Primarily in my *Intention, Plans, and Practical Reason*. See also my "What Is Intention?" in Cohen, Morgan, and Pollack, eds., pp. 15–31.

10. See my discussion in *Intention, Plans, and Practical Reason*, and in "Planning and the Stability of Intention," *Minds and Machines* 2 (1992): 1–16.

our prior plans they would be of little use. The nonreconsideration of one's prior intentions will typically be the default.

Intentions, then, are normally stable elements of partial plans. These plans are subject to demands for coherence and consistency, demands which help structure further planning. Such planning is not the only mechanism that coordinates an individual's purposive activity over time. A tiger hunting her prey may exhibit wonderfully coordinated activity without being capable of such planning. But for creatures like us—as Donagan says, “creatures . . . of will”—planning is an important coordinating mechanism.¹¹

III

I need now to discuss two more preliminary issues. First: my strategy is to see our shared intention to *J* as consisting primarily of attitudes of each of us and their interrelations. At least some of these attitudes will specifically concern our joining action of *J*-ing; after all, our shared intention to *J* supports coordination specifically in the pursuit of our *J*-ing. But much talk of joint action already builds in the very idea of shared intention. For us to try to solve a problem together, for example, we need an appropriate shared intention. We would risk criticizable circularity if our analysis of shared intention itself appealed to joint-act-types that involved the very idea of shared intention.¹² So we will want to limit our analyses to joint-act-types that are, as I will say, neutral with respect to shared intention. For example, we will want to use a notion of painting the house together that does not itself require that the agents have a shared intention.¹³ I assume that we will have available appropriate conceptions of joint activity that are neutral with respect to shared intention; or anyway, my discussion is limited to such cases.

A second problem: the attitudes of the individual participants that are constitutive of a shared intention will include intentions of those participants. But what I intend to do is to perform actions of my own: I cannot intend to perform the joint action *J*. So how will the conception of the joint action get into the intentions of the individuals?

Distinguish two strategies. First, we can appeal to my intention to play my part in our *J*-ing, where this entails that our *J*-ing, while not something I strictly speaking intend, is something I want.¹⁴ Second,

11. The quote from Donagan is from *Choice*, p. 137.

12. Donagan discusses an analogous problem for individual intentional action in *Choice*, pp. 87–88.

13. Think of a case in which we paint it during the same time period but we are each ignorant of the other's activity.

14. An appeal to my intention to play my part in our *J*-ing is similar to the approach of Raimo Tuomela and Kaarlo Miller to what they call “we-intention” (see “We-Intentions,” *Philosophical Studies* 53 [1988]: 367–89, esp. pp. 375–76).

we can try to exploit the fact that we speak not only of intentions *to*, but also of intentions *that*—for example, my intention that Scott clean up his room. Accordingly, we can speak of my intention that we *J*.¹⁵

Consider the second strategy. The idea here is not to introduce some fundamentally new and distinctive attitude. The attitude we are appealing to is intention—an attitude already needed in an account of individual intelligent agency. But we are allowing this attitude to include in its content the joint activity—our *J*-ing.¹⁶ Such appeals to my intention that we *J* will seem reasonably natural given an emphasis on the roles of intentions in plans. This is because my conception of our *J*-ing can function in my plans in ways similar to my conception of my own *A*-ing: in each case I face problems of means and preliminary steps; and in each case I need to constrain the rest of my plans in the light of demands for consistency. And susceptibility to these demands for coherence and consistency is a characteristic sign of intention.

It might be objected that talk of an intention that we *J* conflicts with the plausible idea that one must see what one intends as to some extent within one's influence or control. That is why I can intend to raise my arm but not that the sun shine tomorrow. But, in fact, this need be no objection to the second strategy; for that strategy can build an appropriate influence condition into its understanding of my intending that we *J*. It can say, roughly, that for me to intend that we *J* I need to see your playing your role in our *J*-ing as in some way affected by me.

So the second strategy coheres with the planning conception of intention and can acknowledge a plausible influence condition. In what follows I will pursue this second strategy: my account of our shared intention to *J* will appeal to your and my intention that we *J*. I will not try to settle the question of exactly what version of the influence condition we should accept, for none of my main points depends on this issue. Nor will I try to argue that the first strategy must fail. My claim here is only that the second strategy is fruitful.

IV

I want to say what it is for us to intend something primarily in terms of (a) intentions and other attitudes of each and (b) the relations of

15. A strategy similar to one once urged on me by Philip Cohen. In "Objects of Intention," (*Philosophical Studies*, in press) Bruce Vermazen defends appeals to intentions that are not intentions to act.

16. This contrasts with John Searle's conception of "we-intending" in his "Collective Intentions and Actions," in Cohen, Morgan, and Pollack, eds., pp. 401–15. A we-intention, for Searle, is a distinctive attitude of an individual—an irreducible addition to the kinds of attitudes of which we are capable. On the tack I am taking, my intention that we *J* and my intention to play my part in our *J*-ing are both intentions—they are both instances of the same attitude; but they are intentions that differ in their contents.

these attitudes to each other.¹⁷ This account should explain how it is that shared intentions support the goal-directed coordination of shared activity, in part by way of coordinated planning and relevant bargaining. Limiting myself to joint-act-types that are neutral with respect to shared intention, I proceed by considering a series of views.

VIEW 1: We intend to *J* if and only if I intend that we *J* and you intend that we *J*.

View 1 does ensure that the participants in a shared intention to *J* each are, in a way, committed to their *J*-ing. But View 1 is nevertheless too weak. After all, each of us can intend that we *J* without even knowing of the other's intention that we *J*.¹⁸ Yet at least that much cognitive linkage is involved in shared intention. Indeed, it seems reasonable to suppose that in shared intention the fact that each has the relevant attitudes is itself out in the open, is public. This suggests that we turn to:

VIEW 2: We intend to *J* if and only if
 1. I intend that we *J* and you intend that we *J*, and
 3.¹⁹ I is common knowledge²⁰ between us.²¹

Now consider an example: you and I each intend that we go to New York together; and this is common knowledge. However, I intend that we go together as a result of my kidnapping you, throwing you in my car, and forcing you to join me. The expression of my intention,

17. Note that my target is *our* shared intention. My direct target is not what Tuomela calls a "we-intention"; for a we-intention is an intention of an individual that concerns a group's activity (see Raimo Tuomela, "We Will Do It: An Analysis of Group-Intentions," *Philosophy and Phenomenological Research* 51 [1991]: 249–77). Nor is my target what John Searle calls a "collective intention" in his "Collective Intentions and Actions." A collective intention, as Searle understands it, is an intention of an individual concerning a collective's activity. Indeed, both Tuomela and Searle want to allow that there can be a we-intention/collective intention even if there is in fact only one individual—one who falsely believes others are involved (see Searle, "Collective Intentions and Actions," pp. 406–7; and Tuomela, "We Will Do It," p. 254). In contrast, it takes at least two not only to tango but even for there to be a shared intention to tango.

18. This is true even if, to intend that we *J*, I must believe that your relevant activity depends on mine.

19. This numbering will help keep matters clearer as we proceed.

20. There is a large literature on the idea of common knowledge. See, e.g., Lewis. I use here an unanalyzed notion of common knowledge.

21. View 2 is in the spirit of Raimo Tuomela's analysis of "intentional joint goal" (see his "What Are Goals and Joint Goals?" *Theory and Decision* 28 [1990]: 1–20, esp. p. 10). View 2 is also close to what Margaret Gilbert calls a "strong shared personal goal analysis" of the psychological background of what she calls "acting together" (see "Walking Together: A Paradigmatic Social Phenomenon," *Midwest Studies* 15 [1990]: 1–14, esp. p. 3). Gilbert rejects such an analysis: she argues that it does not guarantee appropriate obligations and entitlements. My reasons for rejecting View 2 are quite different. I turn to Gilbert's concerns later.

we might say, is the Mafia sense of “we’re going to New York together.” In intending to coerce you in this way I intend to bypass your intentional agency. And that seems to rule out a shared intention to go to New York: my intention will surely not support coordinated planning about how we are going to get to New York. Granted, if I succeed in what I intend, our activity will in a way be unified: we will indeed go together to New York. But since the way our activity is tied together bypasses your relevant intentions, this is not the kind of unified agency characteristic of shared intention.²²

This suggests that in shared intention I not only intend that we *J*; I also intend that we *J* in part because of your relevant intention. I intend that our performance of the joint activity be in part explained by your intention that we perform the joint activity; I intend that you participate as an intentional agent in a joint activity that, as I know, you too intend. However, once we bring into the content of an intention of mine the efficacy of your intention, it is a short step to including as well the efficacy of *my own* intention. In a case of shared intention I see each of the participants, including me, as participating, intentional agents. If this obliges me to include the efficacy of your intention in the content of my relevant intention, then it seems plausible to suppose that it also obliges me to include the efficacy of my own intention. After all, I see each of us as participants in the shared intention and the shared activity. Why would what I intend include a requirement that your intention that we *J* be effective, and yet not include an analogous requirement concerning my own intention that we *J*?

These considerations, taken together, argue for:

VIEW 3: We intend to *J* if and only if

1. (a) I intend that we *J* and (b) you intend that we *J*
2. I intend that we *J* because of 1a and 1b; you intend that we *J* because of 1a and 1b
3. 1 and 2 are common knowledge between us.

In shared intention the constitutive intentions of the individuals are interlocking, for each agent has an intention in favor of the efficacy of an intention of the other. And the intentions of each involve a kind of reflexivity, for each has an intention concerning the efficacy of an intention of her own.

Now, Donagan has argued that the choice characteristic of individual intentional action is a choice that one act in a way explained by that very choice: “The choices that explain actions are explanatorily

22. This example, and the one to follow after View 3, are also discussed in my “Shared Cooperative Activity,” *Philosophical Review* 101 (1992): 327–41. See esp. pp. 332–33 (where I have more to say in defense of conditions to be added below in Views 3 and 4) and pp. 334–35 (where I have more to say about the kind of coercion involved in the Mafia example).

self-referential.”²³ The idea that shared intention involves reflexive intentions of the individuals is in a way similar in spirit to this claim of Donagan’s. Nevertheless, my claim about shared intention is compatible with the rejection of the need for self-referentiality in the case of individual intentional action. In a case of shared intention each agent sees herself as one of a pair of participants. Given that she intends that the relevant intention of the other be effective, and given that she recognizes that she and the other each have an intention in favor of the joint activity, there is pressure on her also to intend that *her* intention be effective. But this pressure arises from the social context of the shared intention and need not be present in the case of individual, nonshared intentional activity. So there is room for the conjecture that it is only when we get to shared intention that each agent is obliged to include in what she intends a reference to the role of her own intentions.

To return to the main thread, note that View 3 does not require that you and I either have or aim at having a shared conception of how we are to *J*. Suppose you and I each intend that we paint the house together in part because of each of our intentions. However, I intend that we paint it red all over, and you intend that we paint it blue all over. All this is common knowledge; and neither of us is willing to compromise.²⁴ On View 3 we have a shared intention to paint the house. But this seems wrong, for neither of us is committed to the interpersonal coordination of our relevant subplans.

Granted, for me to intend that we paint the house, despite my knowledge of our differences, I need to think there is some real possibility that we will nevertheless paint it. But perhaps I think this because I think I can trick you about the color of the paint in your can. We might then satisfy 1–3 of View 3; and yet we would still not have a shared intention. For our intention to be shared neither of us can intend that the other’s relevant subplans be subverted. A shared intention should function to unify our *intentional* agency at least to this extent; otherwise it would not support appropriately coordinated planning.

So we need to go beyond View 3. But we also need to be careful not to go too far. First, it would be too strong to require that the subplans of our intentions in *1a* and *1b* completely match, for there

23. Donagan, *Choice*, p. 88. Others who have defended similar views about the self-referential causes of intentional action include Gilbert Harman, “Practical Reasoning,” *Review of Metaphysics* 29 (1976): 431–63, and *Change in View* (Cambridge, Mass.: MIT Press, 1986); John Searle, *Intentionality* (Cambridge: Cambridge University Press, 1983); and David Velleman, *Practical Reflection* (Princeton, N.J.: Princeton University Press, 1989). For a trenchant critique of such views, see Alfred Mele, “Are Intentions Self-Referential?” *Philosophical Studies* 52 (1987): 309–29.

24. Rachel Cohon helped me get this example into shape.

can be features of your subplan that I do not even know or care about, and vice versa. Perhaps your subplan includes painting in overalls or buying the brushes at a certain store. While I need to know you will show up with the brushes, I may well neither know nor care how you are dressed or where you get the brushes. So our subplans may well not completely match. Still, it seems that we will each want them in the end to *mesh*: our individual subplans concerning our *J*-ing *mesh* just in case there is some way we could *J* that would not violate either of our subplans but would, rather, involve the successful execution of those subplans. If I intend that we paint solely with red paint and you intend that we paint solely with blue, our subplans do not mesh. But if you intend to get the paint at Greg's Hardware, and I simply do not know or care about where you get the paint, then our subplans, while they do not completely match, may still mesh. And it is meshing subplans that are our concern in shared intention.

There is a second way in which we must be careful not to go too far. For you and I to have a shared intention to *J* we need not *already* have arrived at subplans that mesh. Much of our relevant planning may occur after we have arrived at our shared intention. All that is plausibly required is that we each intend that we *J* by way of meshing subplans. This leads us to:

VIEW 4: We intend to *J* if and only if

1. (a) I intend that we *J* and (b) you intend that we *J*
2. I intend that we *J* in accordance with and because of 1a, 1b, and meshing subplans of 1a and 1b; you intend that we *J* in accordance with and because of 1a, 1b, and meshing subplans of 1a and 1b.
3. 1 and 2 are common knowledge between us.

On View 4, then, I need neither know nor seek to know of all your subplans for us to have a shared intention; nor need we already have arrived at complete, meshing subplans. What is required is that I *intend* that we *J* by way of meshing subplans. I can so intend even though there are as yet no specific, meshing subplans such that I intend that we *J* by way of them. You and I may not yet have filled in each of our subplans, or we may have filled them in in ways which do not yet mesh. We may have conflicting preferences concerning subplans and be involved in negotiations about how to fill in our plans even while we have already started to *J*.

It is worth reflecting on this last point. Our shared intention can serve as a relatively fixed background against which relevant bargaining can take place. Suppose you and I jointly intend to paint the house together but we have yet to agree on the colors or on the division of roles. Given our conflicting preferences we may engage in various forms of bargaining. Difficulties in such bargaining may, of course,

lead either of us to reconsider the intention that we paint together. But so long as we continue so to intend, our bargaining will concern not whether to paint together but how. Our bargaining will be framed by our shared intention.

Recognition of such potential bargaining raises a question. Suppose that you and I satisfy conditions 1–3 of View 4 with respect to our going to New York together but that there are large differences between us in relevant bargaining power. Perhaps it is a very important matter for you but only a welcome break from work for me. Suppose I plan to use this difference to bargain hard for meshing subplans that are very much to my liking. Perhaps I plan to put a lot of pressure on you to pay for both tickets. According to View 4 we could still have a shared intention. Is that an acceptable result?

I believe that it is; though, of course, too much stubbornness might result in the dissolution of our shared intention. Granted, at some point the exploitation of large differences in bargaining power becomes coercive. When it does our activity of going to New York together (if that is what we manage to do) will not be a fully cooperative activity. But it may still be one that is jointly intentional; and we may still have a shared intention so to act.²⁵ There still may be appropriate kinds of coordination in our planning and action.

A virtue of View 4 is that it allows for shared intention even when the agents have different reasons for participating. We can intend to sing the duet together even though my reason is the love of the music and yours is, instead, the chance to impress the audience.

View 4 does have a drawback: it does not yet provide for a shared intention to play a competitive game together. You and I might have a shared intention to play chess together and yet neither of us intend that our subplans mesh all the way down. After all, I intend to try to scuttle your plans for checkmating me. I think such cases will force modest modifications in View 4; but I will not try to get this straight here. Instead, I want to explore further whether, cases of competitive games to one side, View 4 provides for appropriate explanations of the coordinated planning and action, and associated bargaining, characteristic of shared intention.

V

Begin by reflecting on three basic points. First, shared intention, as I understand it, is not an attitude in any mind. It is not an attitude in the mind of some fused agent, for there is no such mind; and it is not an attitude in the mind or minds of either or both participants. Rather, it is a state of affairs that consists primarily in attitudes (none of which

25. For suggestions of other conditions on cooperative activity that are not insured by the successful execution of a shared intention, see my “Shared Cooperative Activity.”

are themselves the shared intention) of the participants and interrelations between those attitudes.

Second, to say in what shared intention consists I have sought to combine two main elements: (1) a general treatment of the intentions of individuals and (2) an account of the special contents of the intentions of the individual participants in a shared intention. Intentions of individuals are normally stable elements in larger, partial plans of those individuals. These plans are subject to demands for means-end coherence and consistency. Because of these demands, intentions tend to pose problems for further practical reasoning and to constrain solutions to those problems. Given these features of the intentions of individuals, and given the special contents identified in View 4, I want to explain how that in which a shared intention consists supports coordinated planning and action, and appropriate bargaining, in pursuit of the joint activity.

Shared intention consists primarily of a web of attitudes of the individual participants. These attitudes of the individuals are subject to various rational pressures. In particular, the intentions of the participants are subject to demands for consistency and coherence. The specific impact of these demands will depend, of course, on the contents of these intentions. And in shared intention the relevant intentions of the individual participants have the special contents we have been discussing. So—and this is the third point—what we want to show is that intentions of individuals with these special contents should lead to planning, bargaining, and action of those individuals which, taken together, constitute appropriately coordinated planning and unified shared activity. The unified action and coordinated planning characteristic of shared intention is to be explained primarily by appeal to the functioning of the attitudes which are constituents of the shared intention.

Let us see how steps in the direction of View 4 contribute to such an explanation. Begin with View 2. Condition 1 of View 2 requires that each intends that we *J*. So the demand for means-end coherence of the plans of each insures rational pressure on each participant to pursue means to the joint *J*-ing. It also follows, given the demand for consistency of each agent's plans, that there is rational pressure on each to eschew courses of action believed by her to be incompatible with the joint *J*-ing.

So far so good. But what we learn from the Mafia case is that this does not insure that there is rational pressure on each participant to aim at coordination with the other's successful execution of *her* intention. Yet the pursuit of coordination with the other's successful execution of her relevant intention is essential to the kind of coordinated planning characteristic of shared intention.

This brings us to View 3. The conditions of View 3 insure rational pressure on each participant to seek means not only to the joint *J*-ing but also to the joint *J*-ing by way of the other's intention. Now, I frequently form my intentions in the light of my expectations about your intentions and actions, including expectations about how my intentions will influence yours. Since my expectations about how my intentions will influence yours may depend on my expectations about how you expect my intentions to be influenced by yours, this can get quite complex. But in this, as Schelling says, "spiral of reciprocal expectations," we still each see the other's intentions merely as data for our deliberations, albeit as data that are potentially affected by our own decisions.²⁶ In contrast, agents who satisfy the conditions cited in View 3 do not see each other's relevant intention merely as a datum, for each intends that the joint activity go in part by way of the efficacy of the other's intention. Each is rationally committed to pursuing means, and eschewing obstacles, to the complex goal of their *J*-ing by way of the other agent's relevant intention. Each aims at the efficacy of the intention of the other.

In requiring that the participants' intentions interlock in this way, View 3 gives up on the idea, implicit in View 2, that the crucial linkage between the attitudes of those who share an intention is merely cognitive.²⁷ Appropriate common knowledge, or the like, is not a sufficient link for shared intention. Each agent needs also to embrace as her own end the efficacy of the other's relevant intention.

However, the conditions of View 3 still do not insure that each agent aims at there being meshing subplans. The conditions of View 3 do insure that each agent seeks a consistent individual plan in support of a joint *J*-ing in which each agent's intention that they *J* is efficacious. But these conditions do not insure that each agent intends that the subplans of both, taken together, be jointly consistent: that is the lesson of the painting case. But shared intention should bring with it rational pressure in the direction of subplans of both participants that are, taken together, jointly consistent. By requiring that the participants intend that they *J* by way of meshing subplans, View 4 insures such rational pressure.

Finally, View 4 makes it clear why shared intentions will sometimes frame relevant bargaining. On View 4 each agent aims at a performance of the joint *J*-ing that goes by way of each participant's relevant intention and its meshing subplans. So even if the participants

26. Thomas Schelling, *The Strategy of Conflict* (Cambridge, Mass.: Harvard University Press, 1960), p. 87.

27. Gilbert in "Walking Together," and Searle in "Collective Intentions" also reject related ideas, though for different reasons.

have differing preferences about how they are to *J*, neither participant will be in a position to pursue such preferences in ways that bypass the other's intentions/subplans. This makes it likely that in such cases the demand on each agent that her plans be means-end coherent will lead to rational pressure in the direction of bargaining that is framed by the shared intention.

Suppose, then, that the intentions of individual participants have the contents and interrelations cited in View 4; and suppose that these intentions—like intentions generally—are subject to demands for consistency and means-end coherence. These rational pressures on these intentions of those individuals will issue in pressure in the direction of coordinated planning and action, and appropriate bargaining, directed at the joint action of *J*-ing. And that is what I wanted to show.

VI

Margaret Gilbert has argued that in an important sense of “acting together” each participant has associated nonconditional obligations to act and nonconditional entitlements to rebuke the other for failures to act.²⁸ On View 4, if you and I have a shared intention to *J* then you ought to perform your role *if* you continue to intend that we *J*. But View 4 by itself seems to offer no guarantee that by virtue of our having a shared intention you have a nonconditional obligation to perform. Does this suggest that something is missing in View 4?

Recall that intentions are subject to a demand for stability. One reason for this is that the reconsideration of an intention already formed can itself have significant costs; a second is that an agent who too easily reconsiders her prior intentions will be a less reliable partner in social coordination. This latter, social pressure toward stability is particularly relevant to the stability of intentions constitutive of a shared intention. So our approach to shared intention can account for rational pressure on a participating agent not too easily to abandon her relevant intentions.

Note further that if each agent's relevant intentions are fairly stable it will normally be reasonable for each to rely on the other to stick with the joint project. The stability of the constituent intentions thereby supports each in planning on the contributions of the other, just as we would want in coordinated planning.

When I too easily abandon my intention that we take a walk together I am, then, being unreasonable. But it does not follow that in abandoning my intention I am violating a nonconditional obligation to you, a nonconditional obligation grounded in our shared intention. To be sure, shared intentions are frequently accompanied by such

28. For example, pp. 5–6 of “Walking Together.” This summarizes aspects of her much longer discussion in *On Social Facts* (London: Routledge, 1989).

obligations. In arriving at a shared intention we frequently make promises or reach agreements which generate corresponding nonconditional obligations. Further, once we begin executing a shared intention implicit promises frequently arise—promises that generate nonconditional obligations. Still, such a promise or agreement does not seem to be, strictly speaking, necessary for a shared intention.

Imagine two singers who each highly value their duet-singing but nevertheless have a clear understanding between them that neither is making any binding promise to or agreement with the other concerning their singing. Each publicly states that she reserves the right to change her mind. These two could still share an intention to sing a duet together.²⁹ They could still engage in coordinated planning aimed at their singing the duet and in which each relies on the participation of the other. Granted, the normal case of shared intention will not be like this. In a normal case there will likely be some promise or agreement; and that will further contribute to the confidence of each that she can plan on the participation of the other. Nevertheless, such a promise or agreement does not seem essential to shared intention. And when there is no such promise or agreement, or some other obligation-generating process, the shared intention may not impose a nonconditional obligation to stick with the joint action.

Consider two different responses to this. First, one might try to insist that the mere satisfaction of the conditions of View 4, in the absence of some further obligation-generating agreement, does not ensure shared intention.³⁰ So our singers do not in fact have a shared intention.

At this point perhaps the dispute is merely verbal and we should simply speak of shared intention in a weaker and in a stronger sense. The weaker sense is captured, pretty much, by View 4. The stronger sense involves yet a further condition, that there be a binding agreement.³¹ I have argued that shared intention in the supposed weaker sense supports coordinated planning and action, and relevant bar-

29. Lewis makes a similar point (p. 34).

30. This is roughly in the spirit of some of Gilbert's remarks as commentator on an earlier and shorter version of this article at the Central Division of the American Philosophical Association (APA), Louisville, Ky. (April 1992). (Gilbert put the point in terms of a special notion of "joint commitment," indicating that "it may be reasonable enough to think of [joint commitment] as an 'implicit agreement'.") This was also Raimo Tuomela's tack in his replies as commentator on an earlier shorter version of this article (presented at the meetings of the Society for Philosophy and Psychology, Montreal, June 1992).

31. In her comments at the Central Division of the APA, Gilbert suggested (as Paul Weirich brought out in the discussion period) that such a binding agreement, and the resulting obligations and entitlements, would itself be sufficient for a shared intention. But that seems to me wrong, since binding agreements do not guarantee intentions on the part of the individual agents to act accordingly. That is why I understand the

gaining, aimed at the joint activity and that it is typically but not necessarily accompanied by relevant nonconditional obligations. That seems to me a reason to see the phenomenon captured by View 4 as at the heart of the matter. At the least we have seen that there is an important kind of shared intention that does not essentially involve such obligations. Such shared intention is primarily a psychological—rather than primarily a normative—phenomenon. The step to nonconditional obligations and entitlements is a step beyond this more basic phenomenon.

Consider a second response to my defense of View 4. One might urge that a shared intention in the sense of View 4 could only come about by way of a process of a sort that generates corresponding nonconditional obligations. Perhaps the process is not, strictly speaking, one of agreement or the exchange of promises; it may just be a more general kind of mutual assurance. But this process will nevertheless be sufficient to support corresponding obligations.

My reply to this is twofold. First, the main claim—that shared intention must *always* come about by way of an obligation-generating process—does not seem to me very plausible: the case of the cautious singers who disavow obligation seems a fairly clear counterexample.³² But, second, even if I were wrong about this, this need not be an objection to View 4. We could still allow that View 4 says what shared intention *is*, while noting that the creation of a shared intention brings with it certain normative consequences. We could still agree with View 4 that shared intention consists primarily of a web of individual psychological states and their interrelations. It would just turn out that the creation of this psychological web has normative consequences.

VII

This approach to shared intention is broadly individualistic in spirit.³³ Granted, much recent work in the philosophy of mind has argued that our ordinary ways of specifying the contents of the attitudes draw on features outside of the individual whose attitudes are in question.

stronger sense of shared intention, if such there be, to include the conditions cited in View 4 as well as a further condition specifying an appropriate normative relation between the participants.

32. I believe that certain cases of coerced shared intention would also provide counterexamples to this overly general claim. Other potential counterexamples may come from cases of shared intention in which the common knowledge is grounded in the background knowledge of the participants and is not the result of assurances each gives the other (a point David Velleman helped me see—though he did this while trying to convince me that such cases posed problems for View 4). I discuss these matters further in my “Shared Intention and Mutual Obligation.”

33. Assuming that the common knowledge condition can be understood along individualistic lines.

Such external features may include the causal context of the use of names or natural kind terms,³⁴ as well as relevant linguistic practices of the community in which the individual is located.³⁵ The individualism of my approach to shared intention can grant these insights about what determines the content of an individual's attitudes. The claim is not that we can specify these contents in ways that do not appeal to elements outside the individual whose attitudes are in question. The claim, rather, is that shared intention consists primarily of attitudes of individuals and their interrelations. The coordinated planning and action, and framework for bargaining, characteristic of shared intention emerge from the proper functioning of these attitudes of the individual participants.

34. Hilary Putnam, "The Meaning of 'Meaning'," in his *Mind, Language and Reality* (Cambridge: Cambridge University Press, 1975), vol. 2.

35. Tyler Burge, "Individualism and the Mental," *Midwest Studies in Philosophy* 4 (1979): 73–121.