

*Collective Intentional Behavior
from the Standpoint of Semantics**

KIRK LUDWIG
University of Florida

The mutual dependence of men is so great in all societies that scarce any human action is entirely complete in itself, or is performed without some reference to the actions of others, which are requisite to make it answer fully the intention of the agent.

An Enquiry Concerning Human Understanding

David Hume

I

Some things we do by ourselves. I typically shave and brush my teeth without any contribution from another. In this case, going it alone improves efficiency. But there are things that I could not do alone. I could not carry on a conversation, sing a duet, play a symphony or a game of chess, poison the environment, have my hair cut, or vote in a national election, just by myself. Some of these things I could not do alone because I contingently lack

*I would like to thank John Biro, Michael Bratman, Sara Chant, David Copp, Aaron Edidin, Charles Hermes, Kevin Savage, Raimo Tuomela, audiences at the Florida Philosophical Association, 11 November 2005, a colloquium at the University of Kansas, 2 March 2005, the “Action and Agency” conference at the University of Florida, Gainesville, 19 February 2005, a colloquium at the University of Fribourg, 4 November 2004, the “Collective Intentionality IV” conference in Siena, Italy, October 14, 2004, members of my fall 2003 and spring 2005 philosophy of action courses at the University of Florida, and Sara Chant’s spring 2006 philosophy of action seminar for helpful comments and discussion.

© 2007, Copyright the Authors
Journal compilation © 2007, Blackwell Publishing, Inc.

the power to do them. No one person has the power to poison the environment, so it must be done collectively if at all. Some of them I could not do alone because they are by their nature joint or collective behavior (I will use the two terms interchangeably), requiring the participation of others. I could not carry on a conversation alone because conversing by its nature requires at least two people coordinating their responses in a talk exchange. Some things I could do by myself or with others, such as push a car or build a boat. Among the things that I do with others, some of them are intentional, that is, things that we do together intentionally. We don't poison the environment intentionally. This is rather a by-product of individual intentional behavior, uncoordinated with the behavior of others. If I sing a duet with you, in contrast, it is something that I must be doing intentionally, and it entails, moreover, that *we* are doing something intentionally and that *we* intend to do something together. Some things I do with others may be intentional or not. You may push the same car I am pushing so that we push it together, though we do not push it together intentionally. On another occasion, we may decide to push it together, and then push it together intentionally.

What is the difference between individual action, my brushing my teeth, and collective action, our carrying on a conversation, on the one hand, and individual intention, my intending to brush my teeth, and collective intention, our intending to carry on a conversation, on the other? What is the difference between intentional and unintentional collective action, our singing a duet together, and our poisoning the environment together? What is the difference between the intentions of individuals engaged in collective intentional action, when they intend to do something together ('we-intentions'), and those of individuals engaged in collective action that is not intentional, when we do something together but have no such intentions (when we have only 'I-intentions')? In particular, to understand collective intentional behavior: (1) must we admit the existence of agents over and above the individuals who are members of groups that act; (2) if so, must we attribute intentions and plans to such agents; and (3) in any case, must we accept that the intentions of agents who participate in collective intentional behavior, we-intentions, cannot be exhaustively analyzed in terms of the concepts which are used in understanding individual intentional action, that is, must we introduce concepts of forms of agency or intention specific to group action?

In the present paper, I approach these questions from the point of view of the semantics of action sentences. I think this holds out the hope of clarifying a number of these vexed issues in the philosophy of collective intentional action. I sketch the approach here and draw out some of its consequences, without attempting a full defense or comparison with other views in the literature.¹ In particular, in this paper, my aims will be circumscribed by the class of sentences that will be my main concern, namely, action sentences with plural subject terms. I will not take up here the analysis of action sentences whose subject terms, though singular, refer to groups or institutions

(e.g., 'The Army of the Potomac', 'The Catholic Church', 'Congress', 'The French', 'General Motors', 'Great Britain', etc.). The account I outline here can be formally extended to this latter class of sentences, and I believe that, in the light of the analysis I give of plural action sentences, it would be surprising if our understanding of this latter class of sentences required a radically different semantic account. But the extension would require additional argumentation and refinement both because of the difference in grammatical number and because the referents are in many of these cases not just groups but institutions, which introduces additional complexities.

I will argue for an analysis of plural action sentences according to which their truth does not require the existence of plural agents or intenders and an analysis of we-intentions which does not require concepts which are not already used in describing and understanding I-intentions. *Prima facie*, plural agents are puzzling. They must, it seems, have beliefs, desires and intentions. But prior to collective action we have only individuals with their beliefs, desires, and intentions. How could individual agents doing things together give rise to the existence of a super agent with its own beliefs, desires and intentions distinct from those of any of the individuals who are members of the group which acts? Much of the pull to admit plural agents into our ontology comes from our commitment to the truth of action sentences whose subject terms refer to groups. The analysis presented here, if correct, bears on questions (1) and (2) above by showing that at least with respect to an important class of sentences attributing collective action, those with plural subject terms, there is no ground to think that commitment to their truth commits us to the existence of group agents. If, as I hope, this analysis can be extended to action sentences with grammatically singular terms that refer to groups and institutions, we will be able to remove the most compelling reason to think that there are group agents.

In section II, my strategy will be to begin with an account of the logical form of sentences about individual intentions and actions, specifically the event analysis introduced originally by Davidson, and then to use that as the pattern for analyzing sentences about collective action and intention. The key to seeing how to project correctly the pattern from singular to plural action sentences, I will argue, lies in properly taking into account the ambiguity between the distributive and collective reading of plural action sentences. In section III, I draw out some consequences of the analysis. In section IV, I provide a more complete analysis of we-intentions, which I truncate in section II in order to concentrate on the analysis of attributions of collective action. In section V, I turn to a partial analysis of the concept of agency which underlies the account I give, in order to fill in the account of collective action, and to raise and respond to some puzzles about group action. In section VI, I raise, by way of puzzle cases, several related questions about the adequacy of the account I offer which focus on the requirements that all members of a group must be agents of what the group does and that all

must intend all (and only) group members be agents of it for it to be done by them intentionally. I argue that the account does not need modification in the light of these cases. Since my account implies, contrary to some others, that there is no common belief or knowledge requirement on joint intentional action, I conclude briefly in section VII by arguing, by example, that this is in fact correct. Though common belief or knowledge is a common feature of collective intentional behavior, it is not a necessary condition for it.

II

I begin with a comparison of [1]-[3], which deal with individual intention and action, with [4]-[6], which deal with collective intention and action—at least on one reading. My interest is in what minimally is required for the truth of sentences such as [4]-[6] on their collective reading.

[1] I intend to build a boat	[4] We intend to build a boat
[2] I built a boat	[5] We built a boat
[3] I built a boat intentionally	[6] We built a boat intentionally

Structurally, the only difference between these seems to be that in the first three the subject term is singular, while in the last three the subject term is plural. An analysis of action sentences and sentences attributing intentions for the case of a singular subject term then should provide the basic pattern or patterns for sentences with a plural subject term.

This makes it seem as if we must accept as a matter of the logical form of [4]-[6] that we are committed to group agents. As the only difference between [1]-[3] and [4]-[6] is that in the former three the grammatical subject is singular while in the latter it is plural, and as in the former the referent of the subject term, an individual, is the agent, so it seems in the latter the agent must be the referent of the subject term, namely, a group.² As we will see, however, this conclusion would be too quick. A quite different picture is suggested once we look more closely at the logical form of singular action sentences and notice that [4]-[6] are ambiguous between a distributive and collective reading.

I will follow a modified Davidsonian line (Davidson, 2001c) in analyzing [1]-[3] as in [1']-[3']. First, I represent the complement of [1] as a *that*-clause in logical form, i.e., I treat 'to build a boat' as coding for a full sentence in which the subject term is suppressed because it is recoverable from the subject of the main clause. I emphasize though, that this is to be thought of as an intermediate stage in the full analysis of [1], and I therefore represent the complement as 'that-[p]'. Further analysis is postponed to section IV. I follow Davidson's event analysis of action sentences,³ but include an additional

quantifier over times to accommodate tense as an independent semantic feature of the verb (Lepore & Ludwig, 2003). In the following ‘ e ’ is a variable that ranges over events and ‘ t ’ is a variable that ranges over time intervals, of which an instant is a limiting case. ‘ t^* ’ is an indexical term that picks out the time of utterance of the sentence or clause in which it is embedded. ‘ $x < y$ ’ means ‘ x is earlier than y .’

[1'] I intend that-[I build a boat]

[2'] $(\exists e)(\exists t: t < t^*)(\text{agent}_x(e, I) \ \& \ (\text{no } y: y \neq I)(\text{agent}_x(e, y)) \ \& \ \text{building}(e, t) \ \& \ \text{patient}(e, \text{a boat}))$

[3'] $(\exists e)(\exists t: t < t^*)(\text{agent}_x(e, I) \ \& \ (\text{no } y: y \neq I)(\text{agent}_x(e, y)) \ \& \ \text{building}(e, t) \ \& \ \text{patient}(e, \text{a boat}) \ \& \ \text{intend}(I, t) \ \text{that-}[\text{agent}_x(e, I) \ \& \ (\text{no } y: y \neq I)(\text{agent}_x(e, y)) \ \& \ \text{building}(e, t) \ \& \ \text{patient}(e, \text{a boat})])$

In the analysis, I separate out the thematic or case roles of agent and patient,⁴ which will be of particular use when we turn to analyzing collective intentional behavior. In any action, there is an agent, and so the role of agent is common to all action verbs; but there are also different relations we can bear to events which make us agents of them, which are determinates of the determinable agency. Particular action verbs may require a particular form of agency. For example, MacBeth caused Banquo's death because he hired the assassins who killed him, but he did not himself kill Banquo because he was not a direct agent of his death. For this reason, I add to the predicate for the thematic role of agent in the analysis of singular action sentences a subscript ‘ χ ’ which represents a parameter set by the action verb in question that determines the particular form of agency (if any) it requires. The utility of this will be clearer in section V. In contrast to the usual representation, I also add a clause to secure the uniqueness of the agent of the event implied by singular action sentences.⁵ If I built the front half of a house and Jack built the back half, neither of us gets to say he built the house, though we were both agents of its building and we did build it together. One might say: I played a game of tennis. But we recognize this has to be filled out with ‘with ...’ to be true, and this is in effect to say that we played together, that is, ‘I played tennis with x ’ is equivalent to ‘ x and I played tennis together’. Finally, I analyze the adverb ‘intentionally’ as contributing a propositional attitude verb, ‘intends’, in which a variable for the event in the complement sentence is bound by the quantifier introduced by the action verb (this treatment departs from Davidson's suggestion in some significant ways).

I leave open the possibility of further, even syntactic, analysis of the ‘ $\text{agent}_x(e, x)$ ’ relation (Pietroski, 2000; Schein, 2002). In particular, I leave open the possibility that it involves quantification over another event of which x is a primitive agent, and which is related to e by a mechanism specified by the action verb in question, so that upon fuller analysis we have

' $(\exists f)(\text{primitive-agent}(f, x) \ \& \ R_x(f, e))$ '. This will leave open the possibility of identifying an agent's actions with those events of which he is primitively an agent, rather than the events the agent is related to thematically as agent in the above representations. For present purposes, though, we can, I believe, put this aside. If the extension I suggest below to the plural case is correct, further syntactic decomposition will not undermine any of the conclusions I wish to draw about collective intentional behavior.

The sentence that appears in the complement of [1'] would be analyzed as indicated in [2']; however, as noted, I will suggest below (section IV) that [1'] does not represent the fully analyzed form of [1]. A fuller analysis of [1] will also be reflected in the last conjunct of the embedded conjunction in [3']. But for the present we can leave this additional complexity to one side. In [2'] 'building(e, t)' represents a component of the meaning of the English verb 'build' which picks out the type of the event without reference either to its agent (if any) or its patient (or product) if any. For very many English verbs there may be no predicate in the language that performs this function. I press into service the present participle, which may be regarded as a term of the metalanguage, the language into which the object language sentence is rendered to make explicit its logical form, in this case, an extension of English.⁶ I leave open for the moment what it is to be an agent of an event, though I will return to this question in section V.

In [3'] the initial event quantifier binds variables inside the scope of 'intends,' that is, the intention is represented as one directly about the action, not merely for an action of the type. It is not enough that one intend, e.g., to help someone, and then help him, in order to have helped him intentionally (Harman, 1976). And it is not enough that the intention directed toward the type of action play a causal role in the action of that type coming about. The prior intention must give rise to an intention simultaneous with the action and directly about it, as of that type, say helping, for the action done intentionally as of that type, to be an intentional helping. I follow Davidson and Anscombe in holding that an action is intentional or unintentional "under a description" (Anscombe, 1957, 1979; Davidson, 2001b; Hornsby, 1980). [3'] gives a precise form to this claim. The very same action may be intentional under one description, and unintentional under another. Whether it is intentional under one description, 'built a boat', depends on whether the last conjunct of the embedded conjunction in [3'] is made true by the event. It is unintentional under the description otherwise. Thus, the analysis of [7] is [7].

[7] I built a boat unintentionally

[7] $(\exists e)(\exists t: t < t^*)(\text{agent}_x(e, I) \ \& \ (\text{no } y: y \neq I)(\text{agent}_x(e, y)) \ \& \ \text{building}(e, t) \ \& \ \text{patient}(e, \text{a boat}) \ \& \ \sim(\text{intend}(I, t) \ \text{that-}[\text{agent}_x(e, I) \ \& \ (\text{no } y: y \neq I)(\text{agent}_x(e, y)) \ \& \ \text{building}(e, t) \ \& \ \text{patient}(e, \text{a boat})])$)

I leave aside here the controversy about whether the so-called simple view of the connection between doing something intentionally and intending to do it, according to which one *A*'s intentionally only if one intends to *A*, is correct (Adams, 1986; Bratman, 1984; Chan, 1999; Harman, 1976; Ludwig, 1992; McCann, 1991), and also any controversies about whether there are individual actions which are non-intentional as distinct from intentional and unintentional (Chan, 1995; O'Shaughnessy, 1973). I endorse the simple view (Adams, 1986; Ludwig, 1992; McCann, 1991), and am inclined to say that our core conception of an individual action is of an event that is intentional under some description, and I assume these things in the following.⁷ However, even if this is mistaken, the modifications required to accommodate those views would not affect any of the basic lessons to be gleaned from extending the analysis above to sentences about collective actions and intentions. Nor will anything in the following depend on taking a side in the debate over whether a causal analysis of action is correct.⁸ My attention will be focused on the aftermath of basic or primitive action rather than its antecedents.

Before we apply this analysis to [4]-[6], it is important to note that, as mentioned above, each of them has two readings. If I am right, this is the key to understanding the form of sentences about collective action. To see this, consider first [8]-[10], for which the two readings will be easy to distinguish.

[8] We intended to go to college.

[9] We went to college.

[10] We went to college intentionally.

Suppose that when I utter [8]-[10], I intend to be speaking about myself and my brother, and that what I want to convey is that we decided that we would go to (the same) college together [8], and then that we did so together [9], and that we did so intentionally [10]. In this case, there is something we do together, not two things we do independently. I will call this the collective reading of [8]-[10].

Now, in contrast, suppose I utter each of [8]-[10] intending to be speaking about myself and my brother, but that we attended college in different states, and at different times. Clearly there are readings of [8]-[10] on which they would be true in these circumstances. No collective action is in question in this case. There are two actions, and two agents, one for each. What makes [9] true in this case is that *I* went to college and *my brother* went to college. There is clearly no single going to college that is in question that has as its agent all of us, or some composite of us, a plural subject of some kind, whatever that would come to. This is the *distributive* reading of each of [8]-[10]. We obviously cannot recover from the use of 'we' who exactly are the members of the group. Thus, in representing the distributive reading, we must give

[9], for example, the reading represented in [9'] (mutatis mutandis for [8] and [10]).

[9'] (For all x : x is one of us)(x went to college)

(Read: for all x such that x is one of us, x went to college.) Thus, 'we' must be interpreted as in effect introducing a restricted quantifier, though there is no suggestion here of eliminating the plural referring term, which is displaced to the nominal restriction. For notational convenience, I will henceforth write 'for all x : x is one of us' as 'we _{x} '. I will call 'we _{x} ' a "step down quantifier" because it is a mechanism for moving from talk of the group to its members.

Thus, there are both collective and distributive readings of [8]-[10]. Similarly, there are both collective and distributive readings of [4]-[6]. What is the difference between the collective readings and the distributive readings of plural action sentences?

Let us focus on [9] for the moment. One suggestion is that 'we' is ambiguous between an interpretation as a restricted quantifier as in [9'] and a term that refers to a plural entity. The collective reading would then be the one on which 'we' is interpreted as referring to a group, which is then treated as the agent of the event. There is an alternative, however, which does not require treating the contribution of the noun phrase differently on the two readings, that is, as contributing a restricted quantifier in the one case, and simple plural referring term in the other. Other things being equal, it would be preferable not to have to treat its contribution differently in the two cases. We do not have to do so because it is possible to get the two readings on the basis of our independently motivated analysis of ' x went to college' and the treatment recommended for 'we' by the distributive reading. Specifically, if we analyze 'we' as a restricted quantifier, there is a scope ambiguity in [9]. In [9'], the event quantifier introduced by the verb is inside the scope of 'we _{x} '. The other reading may be obtained by giving the event quantifier wide scope with respect to 'we _{x} '. For this is to say that there is one event of which each of us is an agent. That is the collective reading.⁹ Thus, the ambiguity in [9] can be accommodated solely by appeal to a scope ambiguity introduced by the interaction of the event quantifier introduced by the action verb and the "step down" quantifier introduced by 'we'. There is no need to postulate an ambiguity in contribution of 'we' to obtain both readings.¹⁰

The observation carries over to [8] and [10]. In [8], we get the distributive reading by treating the initial quantifier introduced by 'we' as binding the argument place in the complement clause in the subject position of ' x built a boat'. We get the collective reading by treating the complement's subject position as occupied by 'we' in turn, which is treated as a step down quantifier, and interpreting it on the pattern of [9]. [10] patterns after [9] with an extra conjunct within the scope of the event quantifier which patterns after [8].

To develop this more precisely, let us return to [4]-[6] on their collective readings. We carry out the analysis in two stages. In the first stage, in analyzing [4]-[6] on the model of [1']-[3'] we will replace 'I' with 'we.' We must then also replace 'no y : $y \neq$ ' by replacing 'y' with a plural variable. I will use upper case italicized letters as plural variables. Thus, we will replace 'no y : $y \neq$ I' with 'no Y : $Y \neq$ we.' This gives us [4'a]-[6'a].

[4'a] We intend that-[we build a boat]

[5'a] $(\exists e)(\exists t: t < t^*)(\text{agent}_x(e, \text{we}) \ \& \ (\text{no } Y: Y \neq \text{we})(\text{agent}_x(e, Y)) \ \& \ \text{building}(e, t) \ \& \ \text{patient}(e, \text{a boat}))$

[6'a] $(\exists e)(\exists t: t < t^*)(\text{agent}_x(e, \text{we}) \ \& \ (\text{no } Y: Y \neq \text{we})(\text{agent}_x(e, Y)) \ \& \ \text{building}(e, t) \ \& \ \text{patient}(e, \text{a boat}) \ \& \ \text{intend}(\text{we}, t) \ \text{that}-[\text{agent}_x(e, \text{we}) \ \& \ (\text{no } Y: Y \neq \text{we})(\text{agent}_x(e, Y)) \ \& \ \text{building}(e, t) \ \& \ \text{patient}(e, \text{a boat})])$

In the second stage, we introduce our notation 'we_x', and indicate its proper scope on the collective reading. This requires stepping down also from the plural variable 'Y' to singular variables. 'No x : φ ' is equivalent to 'All x : $\sim\varphi$.' When we step down from a universal quantifier with a plural variable, we want the quantifier over individuals inside the plural quantifier. 'All X : $\sim\varphi$ ' becomes 'All X : all z of X : $\sim\varphi$ ', that is to say, 'All X : no z of X : φ ', or 'No z of any X : φ .' Abbreviate this as 'No z - X .' Intuitively, 'we_x' is inside the scope of 'No z - X .' Thus, we have [4'b]-[6'b] in which we indicate more explicitly the appropriate scope of each appearance of 'we_x.'

[4'b] (We_x)(x intends that-[we build a boat])

[5'b] $(\exists e)(\exists t: t < t^*)(\text{we}_x(\text{agent}_x(e, x) \ \& \ (\text{no } z\text{-}Y:(\text{we}_y)(z \neq y))(\text{agent}_x(e, z)) \ \& \ \text{building}(e, t) \ \& \ \text{patient}(e, \text{a boat})))$

[6'b] $(\exists e)(\exists t: t < t^*)(\text{we}_x(\text{agent}_x(e, x) \ \& \ (\text{no } z\text{-}Y:(\text{we}_y)(z \neq y))(\text{agent}_x(e, z)) \ \& \ \text{building}(e, t) \ \& \ \text{patient}(e, \text{a boat}) \ \& \ \text{intends}(x, t) \ \text{that}-[(\text{we}_x)(\text{agent}_x(e, x) \ \& \ (\text{no } z\text{-}Y:(\text{we}_y)(z \neq y))(\text{agent}_x(e, z)) \ \& \ \text{building}(e, t) \ \& \ \text{patient}(e, \text{a boat}))])$

I leave the complement of [4'b] for the moment unanalyzed. Clearly there is the same ambiguity in the complement of [4'b] as in [5] between a collective and distributive reading, one of which, the collective reading, will pattern after [5'b], though there will be some differences as well, since the tense construction is different. For the collective reading, we will want the complement to pattern after [5'b]. In [5'a] the uniqueness condition carries over from the analysis of singular action to group action, which is unpacked in [5'b]. In [6'b] the outermost quantifier over members of the group binds two arguments, one for agents of the event, and one for intendings of the event that it be a joint action. Within the scope of 'intends' the restricted quantifier is introduced because members of a group that intentionally does something together need not intend of each member that that member be an agent of

the event. For example, the electorate can vote for a new president without any of them knowing who all the voters are. This provides a straightforward analysis of these sentences on their collective readings that is the result of applying the analysis of singular action sentences to collective action sentences, and the lesson about the proper treatment of ‘we’ which we derived from the distributive readings of [8]-[10].¹¹

III

I now highlight four features of this analysis. The first is that it exhibits the crucial distinction between the collective and distributive readings of plural action sentences as a matter of logical form. Specifically, it is a matter of whether the quantifier over events introduced by the verb takes wide scope over the plural subject term. The ambiguity of sentences such as [2] and [9] can be removed by adding the adverb ‘together.’¹² Thus, ‘together’ functions as a scope indication device. Arguably the joint reading of sentences without ‘together’ involves treating them as elliptical for the longer sentence obtained by including it explicitly. Thus, the difference between joint and distributive readings of plural action sentences is a matter of whether they imply that there is a single event of which there are multiple agents. This also seems to be intuitively correct. Think of what is involved in our poisoning the environment together. This is a single (complex) event of which there are multiple (mostly unwitting) agents.

A second connected point is that, if this analysis is correct, the appearance that there are any agents other than individuals is shown to rest on a mistake about the logic of the relation of plural referring terms to predicates.¹³ We saw that on the distributive reading of [9], in the scenario discussed, there could not be a plural agent of the event because there was no single action involving both of the agents. The actions involved are clearly, in the situation described, actions of individuals, namely, of each of the individuals among those intended by the speaker in his use of ‘we.’ Thus, we need to invoke a connection between ‘we’ and individuals in order to say what the sentence expresses. We need a way of stepping down from the plurality to each of the individuals who are members of it, in order to say of each of them that he went to college. The phrases ‘each of us’ or ‘each of them’ are made for this purpose. Thus, we represent ‘we’ as in effect a quantified noun phrase. This represents ‘went to college’ and other predicates expressing actions as of a uniform type, i.e., as being satisfied only by individuals. This in turn is just what is needed to make sense of *two* readings of sentences such as [4]-[6], for it shows us how ‘we’ can enter into scope relations with other elements of the sentence, and so how the event quantifier introduced by the verb can take wide scope with respect to it, generating the collective reading without collective agents. Most people are baffled by the idea that there are genuinely collective agents. It is only the misleading surface form of plural

action sentences that pulls us toward what on the face of it is an absurd idea.

A third connected point is that the only intentions we need in order to understand “collective action” and “group intentions” are the intentions of individuals. There are no *group* intentions *per se*. The need to countenance these disappears with the need to countenance collective agents. If there are no collective agents, there are no collective intenders. In ‘We intend to build a boat’, the main clause subject ‘we’ is interpreted as ‘(For all x : x is one of us)’ on both the distributive and collective readings. The difference between the two readings lies in whether the argument place in the complement is bound by the quantifier introduced by the initial plural term, or whether ‘We’ is taken as the subject of the complement sentence in logical form and then given narrow scope with respect to the event quantifier introduced by the main verb in the complement. We are committed only to intentions of individuals in saying what we intend when planning joint actions.

Against this it might be objected that, though ‘intend’ is a state verb rather than an event verb, Davidson’s analysis of adverbial modification should be extended from event verbs to state verbs to handle adverbial modification of state verbs. For example, to explain the contribution of ‘firmly’ in a sentence such as ‘John firmly intends to do it’, it might be suggested we must take ‘intends’ to introduce a quantifier over states and treat ‘firmly’ as contributing a predicate of the state variable so introduced. It is controversial whether state verbs introduce a quantifier over states, though there is some evidence for this.¹⁴ Let us suppose that they do for the sake of argument. If we take state verbs to introduce a quantifier over states, this would in principle, it seems, allow a non-distributive reading. This would not be a reading on which there was a collective intender, some super agent with states of its own. It would, rather, be a reading on which several individuals would in some sense partake in or support the same state, a state of intending.

However, even if it is semantically and syntactically available, we do not in fact accept such a reading. This is shown by the oddity of trying to force it by adding ‘together’ as a modifier of the main verb of the superordinate clause. [11] is clearly acceptable, but in this case ‘together’ is read as modifying the main verb in the subordinate clause. If we place ‘together’ in front of the verb in the superordinate clause as in [12], it sounds decidedly odd, and one might almost say ungrammatical.¹⁵

[11] We intend to build a boat together.

[12] ?We together intend to build a boat.

To the extent to which we can interpret [12] intelligibly, it is by treating ‘together’ as a displaced modifier of the subordinate clause verb. This is intuitively what we should expect. A state of intending cannot be shared by

distinct agents, anymore than states of consciousness can. Each such state is a state of an agent, and the only agents on the scene in our intending to build the boat together are each of us. It is a conceptual truth that intentions are not shared between agents, and it shows up here in the judgment of the unacceptability of [12] when ‘together’ is interpreted as modifying ‘intend.’ This is reflected in [6’b] in the quantifier over members of the group binding the variable ‘ x ’ in ‘agent $_{\chi}(e, x)$ ’ and in ‘intends(x, t).’ There are no collective intendings, and there is no collective intending.

The fourth feature has to do with the vexed question of what is special about the intentions of agents who are engaged in intentional joint action, so-called we-intentions (J. Searle, 1990; Tuomela, 1991a, 1991b, 2005; Tuomela & Miller, 1988).

In one sense, this is transparent on the present analysis. What is special about the intentions which agents of joint intentional action have has to do with their content. They are intentions whose contents involve every member of the group being an agent of the action. That is, they are intentions of individuals to the effect that, *inter alia*, every member of a group of which he is a member be an agent in the action. This explains why in Bratman’s case of my intending that we go to New York by way of my kidnapping you (the Mafia sense of “we’re going to New York”) I do not have a we-intention (Bratman, 1992, p. 333). For in this case, I do not intend that there be any event of which we are both agents. I intend there be an event of which I am the agent which results in us going to New York. ‘we go to New York’ is given the distributive reading in the content of my intention; and, indeed, in your case we do not interpret ‘go’ as an action verb at all but only as a verb of movement, as in ‘I go around the sun once a year.’¹⁶ *Prima facie*, then, the question whether we-intentions are in some sense irreducible to individual intentions and actions rests on the question whether the notion of joint agency that is invoked is reducible to the notions involving only individual intention and agency.¹⁷

The concept of a joint action as such is just that of an event of which there are multiple agents. To see this, consider an example of an unintentional joint action. When we poison the environment together, none of us intend to do so. So, none of us intend that we poison the environment, on either the non-distributive or distributive reading of ‘we’ in the complement. Yet, it is something that we do together. So we poison the environment together, but we do not poison the environment intentionally. If we consider the analysis of ‘We poisoned the environment intentionally’ in [13], we can see that we get the right result.

[13] $(\exists e)(\exists t: t < t^*)(we_x)(agent_{\chi}(e, x) \& (no\ z-Y)((we_y)(z \neq y))(agent_{\chi}(e, z)) \& poisoning(e, t) \& patient(e, the\ environment) \& intends(x, t) \text{ that-}[(we_x)agent_{\chi}(e, x) \& (no\ z-Y)(we_y)(z \neq y))(agent_{\chi}(e, z)) \& poisoning(e, t) \& patient(e, the\ environment)])$

There is no event such that each of us intends that each of us is an agent of it and it is a poisoning of the environment. We poison the environment together unintentionally.

The concept of a joint action, then, of an action of which there are multiple agents, does not involve the concept of any intentions other than individual intentions, and, in particular, it does not involve the concept of a we-intention. This shows that the concept of a joint action can be analyzed completely in terms of individual intentions (I-intentions, in contrast to we-intentions), individual actions, and their consequences.

Does this show that the concept of a we-intention is itself completely analyzable in terms of concepts that we invoke in understanding individual actions? I think we are not yet in a position to draw this conclusion. What these reflections show is only that so far as the concept of our doing something goes, we do not need to invoke any concepts beyond those required to understand individual intention and action. And so insofar as we focus on what is contributed to the content of a we-intention by 'we will build a boat' we do not need to invoke any notions which are not available in understanding individual intention and action. But as I said above, we have not fully explicated the content clause of 'We intend that-[we build a boat].' For we need to relate this to what each individual who has such an attitude intends in order to explain how this is related to what he does. And it is open, prior to doing this, that we may find we have to invoke some notions beside those available already in understanding individual intention and action. A full vindication of the claim that we-intentions can be understood wholly in terms of concepts already available in understanding individual intentions and actions awaits a fuller analysis of the content of we-intentions, which I will undertake in the next section.

The analysis thus far, however, handles the example discussed by Searle of a group of business men who are educated at a business school at which they are taught Adam Smith's theory of "the invisible hand" (Searle, 1990, pp. 404–405).¹⁸ They are each convinced that they will do best in helping humanity by pursuing their own selfish ends, that is, all accept that if "... he intends only his own gain," then he will be "in this, as in many other cases, led by an invisible hand to promote an end which was no part of his intention" (Smith, 1976, p. 423). They each form the intention to pursue their "selfish interests without reference to anybody else," and so each "intends to do his part toward helping humanity" (Searle, 1990, p. 404). Searle adds that they believe that the preconditions for success obtain and that they will all do their parts, and they all believe that the others believe that they all believe this. I will discuss the question of the status of the belief condition in section VII. For now I note only that this is not a feature of my analysis. The question is whether this example raises a difficulty for anything I have claimed. It should be noted first that there is a joint action that they perform, assuming that in fact they are correct in their beliefs, namely, they do in fact

help humanity by each pursuing his own selfish ends. And, on the distributive reading, they intend to help humanity. Intuitively, Searle claims, they do not perform a collective intentional action, and do not have we-intentions. This claim, however, is entailed by the analysis I have given, for Searle has not described each of the businessmen as intending that *they* help humanity. In fact, the intentions they have are described as intentions to pursue their own selfish ends *without reference to anybody else*. This is then to deny that they each intend that *they do something together*. And thus, on my analysis, they fail to have we-intentions. My analysis, then, explains why Searle's example fails to be an example of collective intentional action.

IV

I turn now to the question of how the we-intention of each individual who participates in a joint intentional action expresses the individual's role in the action. The intention is expressed above in a sentence of the form: I intend that-[we do *A*]. But for this to play a role in an individual's doing his part in doing *A*, this either must not fully express the content of the intention, or it must lead to an additional intention that expresses what his role in doing *A* is.

As a preamble to deciding how we should understand this, let us consider what we would say about the case of an I-intention whose complement does not express directly an action of the agent's. Our strategy, as above, will be to apply the lesson for the individual case to the collective case. Suppose that Jim intends that Sue sit next to Bill at the party.¹⁹ What is required for Jim to satisfy or carry out his intention? Is it enough that Sue in fact sit next to Bill at the party? Suppose that Jim finds at the end of the evening that he had never got around to doing anything to bring it about, but learns that in fact Sue *did* sit next to Bill. Jim's desire that Sue sit next to Bill is certainly satisfied. But his intention looks not to have been carried out. Why not? The answer is that Jim did not in fact do anything that contributed to Sue's sitting next to Bill. He did not execute his intention.

The content of an intention or other propositional attitude is given by a statement of the conditions under which it is or would be satisfied (Searle, 1983, ch. 3).²⁰ In the present case, that means that the content of Jim's intention that Sue sit next to Bill is not given by the proposition that Sue sit next to Bill. Rather, it must advert to Jim's bringing it about that Sue sit next to Bill. But even this is not enough, because an intention is not satisfied if its end is achieved accidentally relative to the plan embodied in the intention. Jim's intention that Sue sit next to Bill will be satisfied just in case Jim does something with that intention, which brings it about, in accordance with a plan he has, that Sue sits next to Bill at the party (Bishop, 1990; Ginet, 1990; Mele & Moser, 1994). Here, we do not have to specify the details of the effective plan, for Jim may not have settled on how to achieve his aim when he

forms his intention. His intention is to bring it about in accordance with a plan or plans he has or will develop. We will read 'in accordance with a plan he has' in this way. When we read this into the content of his intention, we get [13]. The general rule is given by [R].

[13] Jim intends that he bring it about in accordance with a plan he has as a result of so intending that Sue sits next to Bill.

[R] Where '*p*' is replaced by a sentence that does not have *a* as its subject a full rendering of '*a* intends that *p*' would be represented by '*a* intends that *a* bring it about in accordance with a plan he has as a result of so intending that *p*.'

Notice that Jim need not do anything that is causally sufficient by itself for Sue to sit next to Bill at the party in order to carry out his intention. He need not do more than, for example, put together a couple of chairs near where he notices that they have stopped to talk, with the expectation that in the circumstances they will take advantage of it and do their parts, not knowing it is part of his plan. It may even be enough that he invite them, if he has a reasonable expectation that if they are presented with the opportunity, they will sit next to each other. Thus, Jim's intending that Sue sit next to Bill does not require him to have it in his control what they will do or intend, but typically only that he have reasonable expectations that they will intend and act appropriately given what he does.

Now we can apply this to the case of intending that we do something. Take as an example our intention to build a boat. We analyze this as 'We intend that-[we build a boat]', and since we treat the first 'we' as distributive, if I am in the group, this means that I must have the intention that we build a boat. As a first pass, we can represent 'I intend that-[we build a boat]' as:

[14] I intend that I bring it about in accordance with a plan I have as a result of so-intending that we build a boat.

Is this adequate? One worry that might arise about this is that 'in accordance with a plan I have' is not specific enough. But we should not require that the intention contain a detailed plan about how one will contribute to bringing about the intended result. We may decide to do something together, for example, reform political campaign financing, without yet having any detailed plans about how to go about doing it. The intention we have to do this together will lead us to make further plans for achieving that end.

This does not extinguish the worry. Surely, it might be said, we must at least suppose that someone who intends that we do something thinks of himself as doing *his part* in what he intends that *we do*, even if his part has not yet been determined. And this idea that he intends to do his part is not captured by his intending to do it in accordance with a plan or plans he has

or will develop. But this objection overlooks the fact that the intention is an intention that *we* do something, and that among those picked out by ‘we’ is the agent whose intention is in question. Thus, in intending to bring about in accordance with a plan he has as a result of so intending that we, i.e., he and every other agent who is one of us, do something, he must be committed to bring it about that he does his part, whatever form it takes, in the joint action. And thus we have not left out this aspect of the we-intention in giving its content the canonical form we have above.

However, this arguably still leaves something out, namely, that the agent should not only intend to do his part, but that he should intend to do his part in a joint plan the agents of the action have, however loosely specified it might be. This is perhaps brought out by the following thought experiment.²¹ Suppose we have two colorblind painters, *A* and *B*. *A* and *B* both intend that they paint the house red. However, *A* thinks that *B* wants to paint the house green and that *B* thinks that that is their plan. He knows that *B* is colorblind, and relies on labels to choose his paint. His plan is to put red paint in a can marked ‘green’ for himself, and to empty *B*’s can marked ‘green’ and replace the paint with red paint. So he intends to do his part in accordance with the plan he has in their painting the house red. *B* also intends them to paint the house red, but he thinks *A* thinks they plan to paint it green. He plans that they should paint it red by his engaging with respect to *A* in the same plan of deceit that *A* plans to practice with respect to *B*. *A* and *B* both satisfy the requirement given above on having a we-intention, and so on this account, *A* and *B* could each truly say (on the *collective* reading): We intend to paint the house red. But something seems to have gone amiss. There’s a sense in which they have the same plan! But also a sense in which it is not a *joint* plan. For neither *intends* that they be operating in accordance with a shared plan, but on the contrary each intends by his plan to be putting something over on the other. And so it seems intuitively that they do not intend to paint the house red, on the collective reading. Each would take himself to be saying something false if he said ‘We intend to paint the house red’, intending the collective reading. But this sense that each has that he would be saying something false would not be removed if each were to learn later that the other intended a similar deceit. The same words could be used to say something true (the distributive reading is true), but what was originally expressed would not be true. If this is right, then what seems to be missing is the idea that they intend to be operating in accordance with a shared plan for doing the thing in question.

Can this be repaired in line with the method of projecting logical form from the singular to the plural case? Let us return for a moment to [R]. We applied this above to ‘I intend that-[we build a boat].’ However, this is in effect an instantiation to me of ‘We intend that-[we build a boat].’ If we apply [R] to ‘We intend that-[we build a boat]’ instead of ‘I intend that-[we build a boat]’ we get [15].

[15] We intend that we bring it about in accordance with a plan we have as a result of so intending that we build a boat.

Suppose that I am a member of the group. Then, reading the initial ‘we’ distributively, we get:

[16] I intend that we bring it about in accordance with a plan we have as a result of so intending that we build a boat.

Notice at this point that we must use [R] again because the complement here does not have ‘I’ as the subject. This then gives us:

[17] I intend that [I bring it about in accordance with a plan I have as a result of so intending that [we bring it about in accordance with a plan we have as a result of so intending that [we build a boat]]].

Now we must decide how to treat each appearance of ‘we’ in the complement. The last one, of course, receives the collective reading. It looks as if we should give the same reading to the first ‘we’ since none of us intend that each of us should bring it about, but rather that each of us should be an agent in bringing it about. That leaves the ‘we’ in ‘in accordance with a plan we have as a result of so intending.’ Essentially, this reads ‘in accordance with a plan we have as a result of *our* so intending.’ We cannot make sense of a shared state of intending, or so I have argued. So for ‘we’ and ‘our’ we need the distributive reading. Thus what is required is that each of us have a plan to bring it about that we bring it about in accordance with a plan each of us has by our each so intending that we do something. And this will rule out the case of coinciding mutual deceit considered above, because in that case neither *A* nor *B* intends that he should bring it about that they bring it about in accordance with the same plan that they paint the house red.

Now let us try to spell this out more precisely. For the sake of perspicuity, let me ignore tense, and abbreviate

$$(we_x)(agent_x(f, x) \ \& \ (no \ z-Y:(we_y)(z \neq y))(agent_x(f, z)))$$

as

$$the-agent_x(f, we_x)$$

I will concentrate on spelling out

[we bring it about in accordance with a plan we have as a result of so intending that [we build a boat]].

‘in accordance with a plan we have’ and ‘as a result of so intending’ both modify ‘bring it about’, which introduces an event variable, and, hence, they are to be taken as contributing predicates of the event variable. For present purposes I will treat ‘bring it about’ as contributing in logical form ‘bring it about(e , that p)’, but not consider any further analysis of this, which I think will not be necessary for our purposes. ‘We build a boat’ receives the analysis provided in [5’b]. So here we may leave it unanalyzed. In the light of this, we can represent ‘We intend to build a boat’ when it is read as expressing a joint intention as:

[18] $(\text{We}_x)(x \text{ intends that } [x \text{ brings it about in accordance with a plan } x \text{ has as a result of } x \text{ so intending that } [(\exists f)(\text{brings-about}(f, \text{ that we build a boat}) \& \text{the-agent}_x(f, \text{ we}_x) \& (\exists p)(\text{plan}(p) \text{ and } (\text{we}_x)\text{have}(p, x) \& \text{according-to}(f, p) \& (\text{we}_x)(x \text{ so-intends}))])]$

Compare this with, for example, A ’s and B ’s intentions in the example of the color-blind painters:

[19] $A/B \text{ intends that } [A/B \text{ brings it about in accordance with a plan } A/B \text{ has as a result of } A/B \text{ so intending that } [(\exists f)(\text{build}(f) \& \text{the-agent}_x(f, \text{ we}_x) \& \text{patient}(f, \text{ a boat})]$

Or perhaps, if we think each intends that they bring it about that they build a boat:

[20] $A/B \text{ intends that } [A/B \text{ brings it about in accordance with a plan } A/B \text{ has as a result of } A/B \text{ so intending that } [(\exists f)(\text{brings-about}(f, \text{ that we build a boat}) \& \text{the-agent}_x(f, \text{ we}_x)]$

Clearly, on our new analysis, neither has a genuine we-intention, because they lack the intention to bring it about that they build a boat in accordance with the same plan. In light of this, our mistake originally was not to apply [R] directly to ‘We intend to build a boat’ but rather to a consequence of it.

This immediately raises the question of what having the same plan comes to. Two people may not be settled on all the details of what is to be done, of course. But that is okay because having the same plan does not require having a very specific plan with all the details of action worked out. It can leave many things, and even most, for later development. Indeed, we are, as noted earlier, reading ‘in accordance with a plan x has’ as ‘in accordance with a plan x has or will develop.’ Thus, all that is required above is that to have a we-intention one intend that members of the group will do the thing in question in accordance with a shared plan. Part of what it is for their intentions to be carried out involves their arriving at and executing a shared plan. But might it not be that joint intenders actually have distinct and even

non-compatible conceptions of how they are to proceed, at least at certain points, while they still jointly intend to do something? Should we deny that they intend that they do the thing in question? The key to seeing why this is not likely to be a problem is that the requirement is only that there be *a* plan that they share for carrying out what they intend, not that *every* plan any has for carrying out what they intend be shared. For any genuine intentional joint action, there will be one at a sufficiently *general* level of specification that they share, for example, that each should coordinate as necessary with the others in the light of information about how things are going to achieve their intended end. Might this not let in too much? It won't let in the case of mutual deceit above, for in that case neither *A* nor *B* intends that there be *any* shared plan. So far then, we seem to have caught the right balance. And we have done so without appealing to any conceptual resources other than those required to understand individual action. For the plan they share can be specified in terms of individual coordinate action.

A worry that a number of writers on collective intentional action have had about intentions that we do things being taken to be about the intender's actions is that it suggests that the agent intends to make others intend to do things, and this is certainly not the case in every collective intentional action.²² However, that I intend that we do something does not entail that it is any part of my plan to manipulate the intentions of others. As noted above in the case of my intending that Bill and Sue sit next to each other, my doing something to bring it about does not require me to intend to manipulate what they intend or do or imply that I have any direct control over that. I may even intend that they do so freely and without any coercion. I intend rationally in typical cases provided that I have reasonable expectations that what I can do will aid in leading to the result I want. The same applies to intending that we do something.

Joint intentional action of the sort I have characterized should be distinguished from what Michael Bratman has called shared cooperative activity (Bratman, 1992). The latter requires the former, but not vice versa. Shared cooperative activity requires mutual responsiveness and a commitment to mutual support. Joint intentional action does not. Bratman's example of two people singing a duet together but being unwilling to help out if the other falters is an example of joint intentional action, for they each intend that they sing a duet together in the sense I have above characterized, and they sing it together intentionally, though it is not shared cooperative action in Bratman's sense, since neither is committed to supporting the other. Further, the case of prepackaged cooperative behavior is an example of joint intentional action which does not involve mutual responsiveness or a commitment to mutual support. For example, members of a team of saboteurs whose plan is to destroy all communications into a city by each of them acting in an independent but coordinated way to sabotage different elements of the communications infrastructure would carry out their destruction of the

communications infrastructure together, but would not have a commitment to mutual support or to be mutually responsive. Commitment to either of these might be thought, for example, to lead to too great a risk of detection, so that the plan calls specifically for none of them to give a thought to responding to the others or coming to their aid.

V

On the account above, we do something together when there is some event of which we are all agents. A deeper understanding of what is involved in collective behavior or action requires then an investigation into what is involved in there being multiple agents of the same event, and so into the concept of being an agent of an event. Since on the present account, the notion of agency at work in collective intentional action is the same as in individual intentional action, this is in effect to ask in what ways each of us can be an agent of an event, and which of these ways allow for more than one of us to be an agent of it. I do not aim to provide here a full account of this but to provide instead enough of a sketch to fill in our framework and to raise and answer to some questions that naturally arise about when certain agents are participants in a collective action. This discussion of the different ways in which one can be an agent of an event will also be put to work in the next section in responding to challenges to my account.

When I brush my teeth, I am *the* agent of the brushing, but I am also *an* agent of it. It suffices for me to be an agent of it (whether intentionally or not) that, intuitively, my doing something contributed to bringing it about. But we can do something that contributes to an event's coming about in a number of different ways. Thus, 'doing something that contributes to an event's coming about' is to be treated as a determinable, of which there are a variety of different determinate forms. An action verb may require a determinate variety of agency. In the account in section II, this is represented by the subscript ' χ ' in 'agent $_{\chi}(e, x)$.' There are, I think, at least four important general ways in which we do something that contributes to an event's coming about, each of which may admit of further more determinate forms.

(1) First, we will wish to allow as a limiting case that I do something that contributes to bringing something about if it is a primitive action of mine, one I do but not by doing anything else. No further analysis of primitive action is required for present purposes. We are agents of other events in virtue of the events of which we are primitively agents bearing some appropriate relation to them. We may therefore characterize other sufficient conditions for being an agent of an event recursively, in the following form (understood to be closed under universal quantification over the free variables): agent(e, x) if $(\exists f)(\text{agent}(f, x) \ \& \ R(e, f))$.

(2) Secondly, I can do something that contributes to bringing something about if I do something that is a cause of its occurrence. When I build a boat,

there is the event of the boat's construction, which is distinct from my actions, though caused by them. I am an agent of its construction in virtue of doing things that causally contribute to it. I contribute to its coming about in this sort of case if something I do would be cited in a complete causal explanation of its occurrence.²³ By a *complete* causal explanation I mean one that does not leave out any positively causally relevant factors.²⁴ This will allow "negative actions" to count as contributing something to an event's coming about, if we allow a decision not to do something to count as an action.²⁵ If each of us, able to intervene and save a child from drowning stands by indifferently and does not intervene deliberately, then we let a child drown. Our deliberately not intervening (distributive reading) is positively causally relevant. (In this case, 'we let a child drown' is true on both the distributive and collective readings.) Only a tenuous causal connection is required for us all to be agents in an event. If everyone on the planet should happen to move counter to the earth's rotation at the same time, they would together slow the passage of day by ever so little.

(3) Third, I can do something that contributes to something's occurring because it includes that as a part. To cite a well-known example due to A. I. Melden (Melden, 1961), in clenching my fist, I contract my forearm muscles, but my clenching my fist does not cause the muscles to contract but rather includes their contraction.

(4) Fourth, I can do something that contributes to something's coming about because what I do is partly constitutive of it. If my moving my hand in a certain social setting constitutes dismissing a case, then my moving my hand brought about the dismissal of the case, but by constituting it rather than causing it. (In Alvin Goldman's terms, this is an instance of the "conventional generation" of an action (Goldman, 1976, pp. 24–5), though I intend the constitution relation to be broader than just "conventionally generated" actions—see the example of applause below.) When you and I play a game of chess, our individual moves do not cause the game to be played, but rather constitute the play of the game. The game occurred because there were these individual moves made with the intention of playing the game, etc. Our shaking hands is constituted by our individual actions rather than caused by them. When the audience applauds after a performance, the individual actions don't cause the applause but contribute to it by partly constituting it. The relation of individual actions to collective actions is often like this.

We then have the following partial characterization of when someone is an agent of an event:

- agent(e, x) if e is a primitive action of x
- or $(\exists f)(\text{agent}(f, x) \ \& \ f \text{ is a cause of } e)$
- or $(\exists f)(\text{agent}(f, x) \ \& \ e \text{ is a part of } f)$
- or $(\exists f)(\text{agent}(f, x) \ \& \ f \text{ in whole or in part constitutes } e)$.

If these capture all the basic ways in which someone can be an agent of an event, then they provide jointly a necessary and sufficient condition.

I assume that if an event is a primitive action of *x*, it cannot be a primitive action of any other agent. If this is correct, then there are multiple agents of an event, and so collective behavior, only if for one or more of the agents the event is not a primitive action of theirs, but one they are an agent of because a primitive action of theirs stands in an appropriate relation to it. I have avoided making any commitment to what actions ultimately are beyond a commitment to a coarse-grained individuation of them. If, as Davidson has argued, all actions of individuals are their primitive actions, however, and to be a primitive action is to stand in some 'direct' causal relation to an agent or the agent's intention, it would follow, first, that there are no primitive actions of collectives, and, second, that strictly speaking there are no actions of collectives.²⁶ Rather, when we speak of collective action and behavior, this is always to be understood as referring to individual actions and behavior in which the primitive actions of more than one agent lead, via one of the ways in which we can be agents of an event, to an event's coming about.

I turn now to some complications that appear to arise when we consider collective behavior in the case of causal contributions to something's coming about. If I tell a group of people to build a pyramid, and they do so on my command, we didn't do it together, but rather they did, though an action of mine would be cited in a complete causal explanation of its coming about. Similarly, if something I do merely inspires a group of people to get together to do something intentionally, or something I do merely influences conditions that are relevant to someone or some group coming to have intentions that lead to their jointly intentionally doing something, it seems that I am not a member of the group that does that.

A first suggestion for handling this might be to require that one's causal influence not be mediated by another's agency. As noted in section II, I do not kill someone if I hire it done, even though I cause his death. So, it might be said, we do not do something like build a pyramid if the role some of us play is to have others do it. I think there is something right in this thought, but that the requirement suggested is too strong as it stands. Often enough when we do things together intentionally, our influence is essentially, and sometimes wholly, mediated by the actions of others. The quarterback on a football team may call the plays, and so tell his teammates what they are supposed to do once the ball is snapped. Here the quarterback's causal influence is both direct and mediated by the agency of others. But even further, one's part may be limited to directing others. A general may do no fighting himself, but only tell others where to go and what to do, but he is part of the army that wins, or loses, the battle, and he plays his appropriate part. His influence on the outcome is wholly mediated by the actions of others. Similarly, the conductor of an orchestra is part of the group that performs, say, Schuman's

Third Symphony, though his contribution is mediated wholly by the actions of the musicians.

What is the difference between the case in which I tell some people to build a pyramid and intuitively they do it without my help, and the case in which a general tells his subordinates what to do, and they intuitively all are participants in fighting the battle? We can first observe that with respect to doing something as a group intentionally there are requirements in addition to causally contributing to an event's coming about. In distinction from mere collective behavior (collective behavior which is not intentional under any description), there is the additional requirement that all members of the group in question intend that every member of the group in question is an agent of the action under a particular description. In the case of my telling some people to build a pyramid, we can see that when they do it, they do not conceive of me as a part of the group that is doing it. When they each intend that they build a pyramid, I am not included in the group that they intend to build the pyramid. In this case, I also fail to think of myself as a member of the group building the pyramid. That is, I do not intend that we build a pyramid, but that they build a pyramid. So I am not part of any group that builds the pyramid intentionally. In contrast, in the case of the general commanding his troops, he clearly is a part of the group that all the members of the collective fighting the battle have in mind. Even if they do not all have him in mind in particular, all of them have a conception of the group they are a part of, which includes him. They would acknowledge him as a member of the group on being given the relevant facts about his position in the organization of which they are a part. However, if my role in the pyramid's coming to be built had involved a blueprint and detailed instructions about how to go about doing it, then I and the others would have thought of me as a participant (the architect of the pyramid), and in this case we would also say intuitively that *we* built the pyramid. We do not, therefore, have to modify our account of what it is to be an agent of an action in order to understand the case involving my telling others to build a pyramid and their doing it intentionally, when I was not a member of that group. This is handled by the analysis of the contribution of 'intentionally' to action sentences.

But this doesn't solve the problem if we just focus on building the pyramid as opposed to building it intentionally. For there is a group larger than the group that builds the pyramid intentionally whose actions had a causal influence on its coming to be built, and it includes at least me, if I order it done, for I am then a *sine qua non* of the pyramid's being built. So did we build the pyramid together, albeit not intentionally? This can't be right. But what about the question whether we caused it to be the case that the pyramid was built? In this case, it does not seem wrong to say that we together caused it to be the case that the pyramid was built. This suggests there is something genuinely wrong in the case of saying that we, I and they, built the pyramid,

that has to do with the concept of building, for otherwise the facts are the same. This thus seems parallel to the contrast between saying that I killed someone when I hired it to be done (false), and saying that I caused his death by hiring someone to kill him (true).

In the case of individual action, we can explain this by taking ‘kill’ to require a special type of agency, direct agency, which requires more than just being the cause of the death of the person concerned, but also having one’s causal influence not be mediated primarily or wholly by another agent. However, we cannot deny me a role in building the pyramid on the grounds that the contribution of my agency is mediated, even wholly, by others; for we have seen that this is often the case in collective behavior. What is wanted here is an articulation of the concept of having a *role* in building something, which is not wholly a matter, at least, of having a causal influence on something’s coming to be built. The concept of having a role is not specially tied to the concept of collective intentional behavior. We all have our roles in unintentional collective behavior, such as poisoning the environment, as well. A sufficient condition for one’s having a role in doing something is that one’s agency is direct. The difficulty is to say what is required if one’s agency is not direct. I will make a brief suggestion here about what is required to have a role in a building if one’s agency is not direct. However, whatever the correct answer is to this difficulty, it is important to note that the extra requirement imposed on being an agent of a building imposed by ‘build’ (and similarly for other action verbs) is a matter of setting the parameter χ in the thematic role for agent in our account of the logical form of action sentences. This represents more determinate varieties of the four basic sorts of agency listed above. Accommodating additional restrictions will therefore not imply any modification to the account of the logical form of plural action sentences offered above. In terms of the suggestion that ‘agent $_{\chi}(e, x)$ ’ may be represented as ‘ $(\exists f)(\text{Primitive-agent}(f, x) \ \& \ R_{\chi}(f, e))$ ’, the additional restriction affects the relation expressed by the action verb between events which are primitive actions of an agent and further events of which he is an agent. My suggestion is that one plays a role in a building (or other similar group activity) when one’s agency is not direct only when the activity is under some description a collective intentional action and one is assigned a coordinating or directive role in the joint plan relative to which it is a collective intentional action. This handles the case of my ordering that a pyramid be built, for I do not contribute directly and play no coordinative or directive role in the joint plan of those who intentionally build it. It likewise handles the case of the general and the quarterback, for there are directive and coordinating roles in the joint plan for the battle and the game which they fill.

VI

I turn now to five connected questions about my account. (Q1) The first is whether *all* members of a group must be agents of an *A*-ing in order for

the *group* to be said to have *A*-ed. (Q2) The second is whether *all* members of the group that *A*'s need intend that members of the group are agents of the *A*-ing in order for them to have *A*-ed intentionally. (Q3) The third is whether members of the group need intend that *every* member of the group be an agent of an *A*-ing when the group *A*'s together for them to have done it intentionally. (Q4) The fourth is whether a group which *A*'s under the impression that they are doing it intentionally as part of a larger group may yet be said to have done it intentionally. (Q5) The fifth is whether in cases of overdetermination we ought to allow that groups whose contributions are minimally sufficient for the event can be said to have acted to bring it about. According to my account, the answers to each of (Q1)-(Q3) is 'yes', and the answer to each of (Q4) and (Q5) is 'no.' But there are cases which may seem to call these answers into question.

I begin with some examples which challenge the affirmative answers to questions (Q1)-(Q3). In each case, I choose examples to illustrate a class. I hope in responding to the particular examples also to provide the pattern of response for other members of the class.

(1) The first concerns members of social groups who seem to play no direct role in what the group does, but instead, for example, play the role of potential contributors to the effort, as needed. For example, consider a football team playing and winning a game. The members who did not leave the bench are also part of the team that won, but it may appear that they contribute nothing to the team's winning the game. This directly challenges the affirmative answer to (Q1).

(2) The second concerns groups which are involved in competition, cooperation, or conflict by proxy, i.e., when the competition, cooperation, or conflict is carried out by a subset of the members of the larger group. I want to look in particular at three sorts of example falling under this general heading, which I will treat differently. (a) The first may be linked to the previous example. When a college football team wins a game, often students (not on the team) at the college will proudly assert that they won, meaning to speak roughly of the members of the extended college community. (b) An archetype for the second is The Battle of Champions fought between the Argives and Spartans in 545 B.C. Three hundred hoplites were selected as champions for each. They fought each other (until two Argives and one Spartan remained alive, the latter of whom later committed suicide for the shame of having alone in his company survived) as a ritual substitute for a general battle between the opposing armies, as reported by Herotodus:

Now at this very time the Spartans themselves were feuding with the Argives over the country called Thyrea; for this was a part of the Argive territory which the Lacedaemonians had cut off and occupied . . . The Argives came out to save their territory from being cut off, then after debate the two armies agreed that three hundred of each side should fight, and whichever party won would possess the land. (I.82.8)

The question is whether this counts as a battle between the armies, or the Argives and Spartans, though only 300 fought on each side, and, if so, how that can be compatible with the requirement that all members of a group be agents of what the group does. (c) A third case, which appears similar to the second, but, I will argue, is in fact importantly distinct, is the administration announcing a new proposal for reducing the deficit by cutting taxes. The actual announcement, it seems, is made by the administration's spokesman, yet it is the administration which is counted as making the announcement. Cases of sort (2a)-(2c) appear to challenge the affirmative answer to question (Q1). (2c) may be used also to challenge the affirmative answer to (Q2), on the grounds that, first, not every member of the administration may intend that the administration announce new tax cuts, either because they are unaware of it, or because of policy disagreements. And (2c) may be used to challenge the affirmative answer to (Q3) on the grounds that those who intend that the announcement be made may be said to intend that just the spokesman be an agent of the announcement.

(3) A third sort of case involves large organizations or groups which are said to do things, and to do them intentionally, when it is clear that many members of the organization are not participating in the action, at least in any direct way, and are not, like the members of a football team on the bench, in readiness to contribute, and may even not be aware that the organization is undertaking the thing in question. For example, when General Motors introduces a new mid-size sedan many employees may be, and remain, unaware of it. This is a challenge to affirmative answers to both questions (Q1) and (Q2).

(4) The fourth sort of example concerns groups doing things or making decisions in which conflicting opinions about what the group should do are aggregated in some way to reach a decision to act in a way that does not agree with what some members of the group think should be done. For example, when a tenure committee makes a decision in favor of tenure by a three to two vote, the committee recommends tenure, though it is clear that not all members of the committee were in favor of it. Similarly in a national election for a new president, not all the voters need have voted for the winner, and not all the eligible voters, let alone all of the citizens, need vote. These cases raise a challenge to affirmative answers to each of (Q1)-(Q3).

My responses to these challenges will employ two strategies, sometimes in combination. The first is to argue that despite appearances each member of the group in question is in fact contributing to the relevant activity and has intentions appropriate to it. The second is to argue that in one or another respect the things we say ordinarily often involve some idealization, exaggeration, loose talk, or confusion.

(1) In the case of the football team, those members who sit on the bench do not participate on the field. However, intuitively they are part of the team that wins. Some may participate by sending in signals. But, in Milton's phrase,

others “serve who only stand and waite.” What is their contribution? How are they agents of the win? I think the answer lies in the fact that the members of the team who sit on the bench are *doing their part on the team in the play of the game*, namely, being available to play as needed. What they do is part of what the team did in playing the game, that is, bring backup or reserve players, to be called on as and when needed on the field of play. These actions constitute part of the team’s play of the game, what the team did in playing the game. So, when the team wins because of the causal contributions of the players on the field, the members whose actions constituted in part the play of the game by the team are members of the team that won. In contrast, if a member of the team is at home with the flu, even if he suits up to watch the game on television, we will not say that he was a member of the group that won the game: he does nothing that counts as part of the team’s play of the game. If he is officially on the team, then speaking of the team per se (given not by the roster of players for the game but *in toto*) as winning is to speak loosely.

(2) What of the student’s claim that *we* won the game, speaking of the larger community associated with the college or university in which he includes himself, when the football team wins? This is best explained as the result of identifying with an entity which one admires, that is, by someone’s appropriating to himself the accomplishments of another with whom he would like to be identified. In other words, the student has slipped into a confused state of mind. This sort of appropriation is commonplace, usually harmless, arguably socially beneficial in many cases, and therefore apt to seem innocent, but it is nonetheless illegitimate. We can take as a model for this the pride we take in the accomplishments of another whom we conceive to be close to us or with whom we identify, but in whose accomplishments we have played no part. I can take pride justifiably only in things that I have contributed to. In some cases we can legitimately take pride in the accomplishments of others. I can take pride justifiably in the accomplishments of my children if, for example, I conceive of my contributions to their upbringing as relevant to them. But people often take pride in things that they have no legitimate connection with, and have made no contribution to. Someone may feel proud of his country or his president, if he speaks out spontaneously and with feeling against bigotry and racial hatred, at home and abroad. This arises from identifying oneself with the individual or group in question through some conceived connection. One then feels his or their act to be symbolically one’s own act. But strictly speaking it is not something one can rationally take pride in because it does not in fact reflect an accomplishment of one’s own. I suggest that a similar mechanism is at work in the claim to be part of a group that has done something one regards as admirable because of some connection one has with the group. One claims to be a part of the group which performs the admirable act, but strictly speaking one is not. Similarly, we speak loosely or confusedly if we say that, e.g., the United States won the

most gold medals at the winter Olympics, because strictly speaking it was the team of athletes they fielded that did so. Someone may speak for us or act for us in a legal proceeding, but no one can win a gold medal for us except in the sense in of doing it to honor or please us.

(3) Let us turn now to the case of The Battle of Champions. This much is clear: The armies, or the Argives and Spartans, or relevant subgroups, decide that they will settle who is to receive the land in dispute on the basis of the outcome of a combat between 300 hoplites selected from each side. However, this does not entail that the armies themselves engage in a battle. In fact, the description of how the dispute is to be settled shows that the combat is conceived of as not taking place between the two armies, or sides more generally. Rather a battle by subgroups of each of the larger groups is a part of a decision procedure which they have adopted to settle a dispute. They adopt this procedure *to avoid* a general battle between the armies.²⁷

(4) When the administration proposes a new tax cut, only the spokesman speaks. In this case, the administration counts as announcing something because it has deputed the task to one or a few individuals who are members (or instruments) of the administration, its spokesmen. That is one way for a group to act, by acting through an individual's agency, who represents the group as a whole, in some action or subset of his actions. In David Copp's apt phrase the administration has commissioned their spokesman to act for them (Copp, 1980, p. 591). The commissioned agent may or may not be a part of the group that commissions an act, though in this example the spokesman is part of the administration. In this case, all members of the group are agents of the event, through having granted responsibility to an individual to act in the group's name, or accepting as a condition of membership in the group that some other member appropriately placed has the authority to grant to someone responsibility to act on the group's behalf. Here each member of the group does something that enables someone to act in their name, some by giving permission, others by making decisions and giving directions, others by carrying them out, and so each makes a contribution to what the group does by doing something that is constitutively connected with what the representative does being counted as the group's doing something. There will obviously be considerable complexity involved in an account of what it is to commission an individual to do something in the group's name, when what the individual does can be disavowed, and so on.

But is it true that all members of the administration do something that contributes, even if it is as minimal as accepting as a condition of membership in the group that others may direct someone to speak on behalf of the group? May not some formally join by accepting a job without knowing much about the organizational structure, and protest later that the spokesman does not speak for him? I think the answer to both questions is plausibly 'yes', and that when people speak of the administration announcing something, we either idealize or we intend to be speaking specifically of a relevant policy making

group and its instruments, and that the use of 'administration' is an instance of metonymy.

Granting that this explains why all members of the group can be counted as agents of the announcement, can we say that the group, here the administration, made the announcement intentionally, on the account I have given? There are two worries here. One is that not all members of the administration intended the outcome, either because they were unaware of it, or opposed it. The other is that members of the administration did not intend that every member of the administration be an agent of the event.

The latter worry we have responded to in explaining how others besides the spokesman can be agents of the announcement, and the objection is based on the thought that since only the spokesman is an agent of the event, members of the administration could only rationally intend that he be an agent of the event.

The first worry divides into two cases. In the first, the concern is that not every member of the administration is aware of the announcement, and so cannot be said to intend that they make it. I will respond to this worry in the response to the next sort of case (5). In the second, the concern is that some members of the relevant policy making groups may have opposed the decision, and so could not be said to intend that the administration make the announcement. The issue this raises will be dealt with in the discussion of the case of voting, in which a group does something, presumably intentionally, though not all members of the group were in favor of what the group does (6).

There is a final connected worry about the relation between the actions of someone acting on behalf of a group and so acting for the group in a way which licenses saying the group did something on the basis of an individual's doing something, and the intentions of the members of the group, that arises when the agent is authorized to act independently of consulting the authorizing group. The group could not then intend anything about the particular actions undertaken by the representing agent on the group's behalf, and which, if other constraints are met, will be attributed to the group. Should we say that in this case the group does the things in question intentionally, though not all the members of the group intended that they do it? I think not. There is a parallel for individual action which helps us to see this case in the right light. If I authorize someone to act on my behalf in a legal matter, for example, by giving her my power of attorney, I would not usually be in a position to intend even that the particular actions undertaken on my behalf take place. However, I will be taken to have, e.g., sold shares in a company, closed an account, and so on. In the case of someone representing me as an individual, although we do allow that I did the things in question, if I do not pay attention to how my power of attorney is exercised, we do not take me to have done those particular things intentionally. I intend that my representative act on my behalf, and so that she perform various acts as my

representative. I do not intend the particular things that she does. I suggest we should say the same thing for the case of a representative acting for, but exercising judgment independent of, an authorizing group. Unless the group has directed a particular action under a particular description to be performed by the representative, the group cannot be said to have performed the action *intentionally* under that description, though they take responsibility for it (within limits).

(5) Often we speak of large groups doing things of which it is clear not all members of the group are even aware. When General Motors introduces a new mid-size sedan, many of its approximately 340,000 employees may not even be aware of it. A complication in this case is the use of the institution name 'General Motors.' It is not clear that it should be construed simply as a name for a group that includes all its employees. For example, when we say that General Motors employs over 340,000 people, we do not mean that the employees of General Motors employ over 340,000 people. The project of understanding institution names and how they are used in action sentences is important, but it is too large a task to undertake here, and I will therefore set it aside. However, we can ask a related question which may shed some light on the issue, for we can imagine someone claiming, 'We introduced a new mid-size sedan last month', meaning to speak of himself and his coworkers at General Motors. In this case, if the speaker intends genuinely to be referring to all those in the employ of General Motors, it has a definite air of exaggeration about it, and what we understand to be the case is rather that a subgroup assigned that project introduced a new sedan last month. And what if some members of that project did not know what they were working on, and so did not have the relevant *we-intention*? Then they were not members of any group that did it intentionally. What we must do to understand more fully the structure of the group activity is what different subgroups within the larger subgroup did knowingly and intentionally. The actual story is likely to be quite complex. That we resort to some loose talk or idealization is unsurprising, since there will be no short and accurate account of the group behavior that results in the introduction of a new mid-size sedan.

(6) In the case of a committee that decides a tenure case favorably by voting, it seems correct to say that they recommended tenure intentionally. But if there were dissenting votes on the committee, surely not all of its members intended that they recommend tenure. How can this be reconciled with the present account? In this case, it is important to notice that none of the members of the committee, not just those who voted against tenure, intended flatly that they recommend tenure. What they all intended was that they recommend for or against tenure on the basis of a majority vote of the members of the committee. The members of the committee have accepted a decision procedure for settling on a joint recommendation, and they intend to be making a recommendation in accordance with that procedure. Thus, they all had conditional intentions: that they recommend tenure if the

majority votes in favor of doing so, and that they recommend against tenure if the majority votes against it. When the majority votes in favor of tenure, then the condition is met, and the entire committee forms the intention that they recommend tenure. Thus, the committee members recommend tenure intentionally though not all of them were in favor of it, in the sense that some of them voted against it, and may still feel that tenure should not be awarded.

The case of electing a new president is the committee case writ large. The voters who choose someone other than the winner are participating in the election, and so in the procedure for choosing the new president. What they do is partly constitutive of an election's taking place. They do their part by voting for one of the candidates, intending that the candidate who wins a majority (of the electoral votes) become president in accordance with the legal and constitutional framework for transfers of power from one executive to another. They recognize that the outcome of everyone doing his or her part may be that someone else is elected, and they intend that if the majority vote goes the other way, the person against whom they voted becomes president.

Not all eligible voters may vote, and not all citizens may be eligible, though we still say that the country (we, the citizens) elected a new president. But this can be explained as a matter of the eligible voters who exercise their franchise in turn representing the country in the election, both the voters who do not vote, and those citizens who are not eligible to vote, in the way a spokesman, as we noted above, may represent and act for an organization. In this case, explicitly opting out of the representation suffices for not being an agent of the event. Thus, if a political party boycotts an election, it seems right to say that they did not participate, and were not part of the group that elected the president. That is the point of a boycott. It puts a group in a position to claim that they were not represented in the election. It is explicitly opting out. And here we will feel uncomfortable with the thought that the nation elected a new president, as opposed to a new president being elected for the nation (for there is no doubt that there would be a new constitutionally elected president even if a political party or substantial portions of the electorate boycott the election). We have in this case a nation divided against itself.

What of the case, though, of citizens, or proto-citizens, children say, who cannot be said in any sense to have accepted even implicitly membership in the nation which goes to vote? It seems doubtful that they should be considered members of the group that elects the president. So, in addition to the above considerations, when we speak of the nation voting, we speak loosely, idealizing, perhaps, or think of the nation as constituted not by official citizenry but by minimally self-conscious citizenry.

I turn now to question (Q4). Sometimes people undertake to do something thinking that they are doing it with certain others, whom they know individually, or whom they pick out by some description, though in the end what is done is not done by way of the participation, or intentional participation

of all the members.²⁸ Let us take a simple case. Suppose that two brothers intend that they and their other brother paint the house together, but that the third brother does not intend that they all do it. The two set out to paint the house, thinking the third is contributing, when in fact he is not. They paint the house. Did the two of them paint it together intentionally? If the answer is yes, then this is a counterexample to my analysis of joint intentional behavior. I think the answer is ‘no’, but that there is a good reason not to say flatly that they painted the house unintentionally. In a case like the above, it may be difficult to imagine that they get through the whole thing without realizing that they are the only ones doing it. To screen for this, let us imagine that two people are told that members of a certain group that includes a lot of people are all going to do something together, and the group is picked out so as to include both of them. They are asked to do with the group a certain task (perhaps it is raising a certain amount of money for a cause) and given a role to play, and they agree to participate. They don’t know any other members of the group, and their roles do not require them to come into contact with any other members of the group. Suppose that in fact there are just two of them who are doing the thing in question, and they do it together. Did the two of them do it together intentionally? In this case, it seems to be a mistake to say that they did it together intentionally, for they were each intending to do it with a group that included many more members. They did their parts. It would come as a surprise to them, however, that what they did was enough for the task to be carried off. However, it also seems unhappy just to say that they did it together unintentionally. If we are to accept that they did not do it together intentionally, then we need an explanation for why it is not felicitous to say that they did it unintentionally. There is, however, a straightforward explanation of this. For to say that they did it unintentionally together, in a neutral context, without background, would suggest or implicate that neither one of them intended to do what they did together as a part of a group, i.e., that they did not have any relevant *we-intentions*. But that is not so in this case.

Finally, I turn to (Q5). In a firing squad of ten men, the shots of any two of them might be sufficient to bring about the death of the prisoner. Should we say, then, focusing on the two rightmost members of the squad, that they killed the prisoner? On my account, the answer is ‘no’, because they are not the only direct agents of it. We can grant the shots of any two of the members of the firing squad would have been sufficient by themselves, and would have been the cause of his death if only their shots had been fired. However, it is not clear that it follows that in the actual circumstances they were a cause of or caused his death. For, intuitively, it is a mistake to say that any two of the members of the firing squad were the cause of the prisoner’s death or that they caused it. The cause of his death was the complex event consisting of the individual firings of each member of the squad. Though the death was

overdetermined, this does not mean that we can speak of each pairing of two shots as an independent cause of the death. Consider an analogy with what we say about the cause of the compression, say, of a can, when more force is applied than is needed. If twice as much pressure is applied to a can in crushing it as needed, we would not say that there was more than one cause of its being crushed, even if we are able to analyze the force into separate components. Now think of the pressure as being applied by several people pushing the can against the ground. Perhaps what any two of them does is enough force, but the total force is the sum of the forces applied by the several individuals, and it is that total force is responsible for the can being crushed.

VII

A final point I wish to make is that this analysis of joint intention and joint intentional action does not require any special attention to mutual beliefs, or any robust belief requirement at all.²⁹ That is, it does not fall out of the concept of *our* intending, or *our* doing something intentionally, that the participants in the action believe that others are doing their parts, or that the action will be successful. A belief requirement might fall out of the concept of *intending*. But no special belief condition would be required as a part of the analysis of *joint* intentional action. I think that, in fact, reflection on cases shows that there is no mutual or common belief requirement on joint intentional action, whatever the source. Members of a group need not believe they will succeed in order to rationally intend to do something and to intentionally do it, nor even to believe that other members of the group will do or even attempt their parts. It seems clear, first of all, that we may try to do something together without much hope of success, and succeed despite our doubts. For example, members of an expedition to the North Pole might rationally believe that they had only a small chance of succeeding, but intend and try to reach the North Pole, and succeed against the odds.³⁰ This shows that no belief in success on the part of participants is required to jointly intend and to do something together intentionally. But even more strongly, it is not required that members of a group even believe that others will play their parts in a joint action which they nevertheless intend to do and perform intentionally. Consider the following example. Suppose that country *X* launches a pre-emptive nuclear strike against country *Y*. After the initial strike, some missile silos in country *Y* are still operative. However, country *Y* has established an elaborate procedure for firing its missiles as a safeguard, which requires two on-site operators, who are physically isolated from one another, and one remote operator, all to punch in a secret code and turn a firing key at their locations in order to launch a missile. Consider the team charged with this for surviving silo 451. After the strike, which interrupts communications between them, none of them knows whether the others have

survived, and have some reason, perhaps even preponderant reason, to think that they have not. Nonetheless, they intend to launch the missile. Each of them intends that they do it, and so each of them intends to do his part in launching the missile. Each punches in his code, and then turns his key, hoping that there are still others who are doing their parts, however unlikely it may seem; and so they launch the missile in silo 451 together, and they do so intentionally.

Notes

¹ See the pioneering work of (Tuomela, 1991a, 1991b, 2005; Tuomela & Miller, 1988; J. R. Searle, 1990; Gilbert, 1990, 1992, 1996b, 2000; Bratman, 1992, 1993, 1999a, 1999b; Velleman, 1997).

² Among those who think that groups can be the subjects of intentional states (in one or another sense), and intentions in particular, are Margaret Gilbert (Gilbert, 1990, 1992, 1996a), David Velleman (Velleman, 1997), Deborah Tollefsen (Tollefsen, 2002), Philip Pettit (Pettit, 2003) and David Copp (Copp, 2006).

³ In this I am accepting a 'coarse grained' individuation of actions, which I will not defend here. I think there are adequate defenses of this in the literature. See in particular (Hornsby, 1980, chs. 1 and 2). The tremendous fruitfulness of this proposal for linguistics and semantics constitutes support for it as well. See (Pietroski, 2003; Schein, 1993, 2002).

⁴ See (Castañeda, 1967), as well as (Jackendoff, 1987; Parsons, 1985, 1990). Thematic roles are associated with particular argument places in event verbs. In 'John hit Bill', John is marked as the agent of a hitting by the position of 'John' in the sentence and Bill as its patient or object by the position of 'Bill.' In many languages, thematic role is indicated by case inflection rather than, as in English, position or preposition.

⁵ There is probably some unavoidable indeterminacy in exactly how this condition is represented. Given this, I have chosen one that I hope will make for perspicuity in the sequel.

⁶ This is connected with the distinction between the transitive and intransitive senses of a verb. Let us subscript a verb with 'T' to indicate its transitive sense and with 'I' to indicate its intransitive sense. It is widely accepted that the truth of a sentence of the form '*A* ϕ_T -ed *B*' requires that *A* caused *B* to ϕ_I . If I move my finger, then my finger moves. If a falling branch breaks the window, then the window breaks. My finger may move without my moving it, and the window may break without anything's breaking it. In both these cases we also will hold that I caused my finger to move and that the branch caused the window to break. The intransitive sense corresponds in the above analysis to the contribution of the predicate of the event that specifies its type and the specification of the patient (or object) of the event. The transitive sense adds the role of the agent. It is, however, not clear that 'cause' is general enough to capture the relation between *A* and *B*'s ϕ_I -ing, when *A* ϕ_T -ed *B*. While I think that one way to be an agent of an event is to cause it, I do not assume that this is the only way. In particular, I would wish to allow us to be agents of our basic actions without saying that we cause them, for otherwise it appears that we must accept an unanalyzed notion of agent causation, or an infinite regress of actions. Cf. (Hornsby, 1980; Pietroski, 2000). In addition, we will, I think, want to allow one can be an agent of an event by doing something partly constitutive of it, or something of which it is a part. See section V for further discussion.

⁷ Of course not everything we speak of ourselves as doing is what philosophers have had in mind by actions. We sleep, breathe, and sweat, cough and sneeze, smile, blush and stammer. We jangle the keys in our pockets, we move our feet restlessly in colloquia, and so on. Perhaps the extension of 'action', fixing one of its senses, does include these as well as those actions of ours that are intentional under a description. But that we work with a narrower notion seems

in any case plausible, since whether an action in some broader sense is intentional under some description or not marks a natural division among activities expressing states of mind. That there may be a sense of 'action' which corresponds is suggested by the fact that when we ask, in 'A-ing, did he act?' substituting the relevant verbs above, we answer, it seems to me, 'no'—unless we think in fact the thing in question was intentional under some description.

⁸ See (Ginet, 1989, 1990; Wilson, 1985, 1989) on the con side and (Davidson, 2001a; Davis, 1984) among others on the pro side.

⁹ When there are more quantifiers in a sentence, the event quantifier can be read as having more than just two positions, as in 'All of the boys danced with all of the girls'; on one reading, for each of the boys there was a dancing with all of the girls. Arguably some sentences require reading the event quantifier as a plural quantifier, as in 'Four architects collaborated on two entries for the design competition', where there were two teams of two architects working separately on the two entries. We should not need to touch on these issues in this study.

¹⁰ As I noted above, my account does not eliminate the use of the plural referring term, and I will take no stand in this paper on whether the use of plural referring terms itself commits us to plural entities. George Boolos, among others, has argued that it does not, but it is a matter of controversy (Boolos, 1984). Whatever the right thing to say is about this, if my account of the logical form of plural action sentences is correct, we will not be committed to plural *agents*.

¹¹ See (Schein, 1993, 2002) for some important and sophisticated work on quantification and action sentences.

¹² 'Together' is, however, ambiguous, and in one sense means 'in proximity', as in 'We were together the entire time, except for the stop at Puerto Vallarta.' See note 16 below.

¹³ Gilbert argues the plural subject is irreducible (Gilbert, 1992, pp. 435–436).

¹⁴ See (Parsons, 2000) for a survey and an interesting argument for underlying states. The view that propositional attitude verbs in particular introduce underlying states is often assumed in accounts of their semantics.

¹⁵ Often judgments of ungrammaticality are based on the sense that something is wrong with a sentence that turns out to be not ungrammaticality but disguised semantic incoherence. Thus, 'John will go to the store yesterday' or 'Tomorrow John went to the store' sound "ungrammatical" though the fault really is that they are semantically incoherent, since the tense puts a requirement on the event time relative to the time of utterance that the restriction imposed by the adjuncts is in conflict with.

¹⁶ We might say here that I intend that we go to New York together, but this is the sense of 'together' which means 'in a group' rather than 'as a group' or 'jointly.' We can also say in this sense that we drove to New York together. This clearly does not imply that we performed a joint action together. Rather, this is the sense of driving together—parallel to the sense of 'flew together' in which we say that we flew to San Francisco together, being passengers on the same flight—in which Bob and Sue, hiring a limousine and driver, drove to Niagara Falls for their honeymoon. The confusing factor is that in this case one of us is also the agent of a driving.

¹⁷ The account of joint intention or group intention given here is not that of a shared intention discussed by (Bratman, 1993, 1999b). The requirement on intentional collective action that falls out of this is a generalization of Bratman's VIEW 1. My goal is to give the truth conditions for the collective reading of plural sentences about action, including about collective intentional actions. This does not require the conditions of common knowledge or dependence on the others' intentions that Bratman discusses, and I doubt that, except perhaps in a very abstract sense, it requires the members of the group to intend that they perform the joint action by way of meshing subplans (see section IV below). This does not represent a disagreement between us, for Bratman does not state his aim in the same way I have stated mine. His is introduced by way of an example of the sort of 'shared intention' he aims to provide an analysis of. Mine is introduced by way of a certain semantic reading of a plural action sentence. The notion of a shared intention is made for the notion of shared cooperative activity which Bratman discusses in (Bratman, 1992). See the last paragraph of section V for discussion of the difference between shared cooperative activity and collective intentional behavior more generally. Even

if there are more robust and involved notions of joint intention and joint intentional action which are important for understanding some sorts of collective behavior of interest to us, it is important to have a clear account of the minimal requirements on joint intention and action, for that tells us what the most general form of the phenomenon is of which there are more specific varieties, and saves us from overgeneralizing from the more specific sorts.

¹⁸ Searle offers this as a counterexample to the account of (Tuomela & Miller 1988), but it looks as if he misunderstood their intent. Tuomela and Miller held that members of a group that do something intentionally must intend to do their parts, and they conceived of this as their thinking of their doing their part in a joint action. While Searle entertains this response, he takes it to be unavailable to Tuomela and Miller because it relies on one of the notions in the family of notions involved in understanding joint action and intention. Tuomela and Miller, however, have disavowed the aim of providing a fully reductive account.

¹⁹ This is adapted from an example in (Vermazen, 1993), though I reach a different conclusion about it. This paper will not be able to address fully Vermazen's argument for some intentions not being directed toward an agent's actions.

²⁰ It is this point, together with the reflections in the previous paragraph, which are my reasons for rejecting Bratman's suggestion that *intentions* that do not have to be about their possessor's actions, on the grounds that they may play a role in coordinating planning (Bratman, 1999b). It seems to me that they are not intentions but expectations, if that is the only role that they play.

²¹ This thought experiment is inspired by, though more elaborate than, an example from (Bratman, 1992, p. 332).

²² See (Baier, 1997; Stoutland, 1997; Velleman, 1997) for objections along these lines. See (Bratman, 1999b) for an extended response to this sort of objection.

²³ Since there may be different sorts of causal explanations, I stipulate that I here mean a causal explanation by subsumption under causal laws.

²⁴ What counts as a positively causally relevant factor is to be governed by our ordinary notion of a cause, which is more restrictive than something's just being an INUS conditions, in J. L. Mackie's sense (Mackie, 1974), that is, an insufficient but necessary part of an unnecessary but sufficient condition. For example, the manufacture of an automobile would not be counted as a cause of its rusting in a junk yard years later. Similarly, God's creation of the universe would not make him an agent of every subsequent event.

²⁵ We would not ordinarily count a decision to do something as an action, for the contemplated action requires more of us. In the case of a negative action, the decision not to do something is all that is required to let something happen or not to intervene, and so it seems the appropriate event with which to identify the action.

²⁶ Sara Chant has suggested to me that we might think of the aggregate of the primitive actions of the individuals engaged in collective behavior as the collective primitive action.

²⁷ The key point appealed to here about the Battle of Champions is that the decision procedure adopted by the two sides specifies that the fight is to be between 300 champions of each side, and so not between the armies as a whole: that is the point of the decision procedure, to decide without the armies as a whole fighting. Might a battle between subgroups count, however, as a battle between the groups from which they are drawn, by a conventional arrangement? I think the answer is 'no.' The reason is that, though conventions may govern how a battle is fought, what counts as a battle is not a matter that is constituted even in part by conventions. The cases in which what a person or group not identical with the whole of a particular group can do something that counts as the group's doing it are cases in which the action type is at least partially governed by constitutive rules or conventions. It is this that makes possible someone acting with my power of attorney doing things that count as my doing them. The things so done involve bringing about changes in social facts which are constituted by actions performed in accordance with conventions in which the giving of a power of attorney and acting on it have a place. Similarly, in the case next discussed of the spokesperson for the administration performing an action that counts as the administration's announcing something,

there are constitutive conventions which govern what is to count as the group announcing something in virtue of what the spokesperson says under certain conditions.

²⁸ I owe this objection to Kevin Savage.

²⁹ For a contrasting view, see (Tuomela, 2005).

³⁰ I believe something stronger, namely, that they need not even believe that it is possible for them to succeed. I have argued for this in the case of individual actions, and I believe that the argument is extendable in a straightforward way to the case of collective action (Ludwig, 1992).

References

- Adams, F. (1986). Intention and Intentional Action: the Simple View. *Mind and Language*, 1, 281–301.
- Anscombe, G. E. M. (1957). *Intention*. Oxford: Blackwell.
- Anscombe, G. E. M. (1979). Under a Description. *Noûs*, 13(2), 219–233.
- Baier, A. (1997). *The Commons of the Mind*. Chicago: Open Court.
- Bishop, J. (1990). *Natural Agency: An Essay on the Causal Theory of Action*. Cambridge: Cambridge University Press.
- Boolos, G. (1984). To Be Is To Be a Value of a Variable (or to Be Some Values of Some Variables). *Journal of Philosophy*, 81, 430–450.
- Bratman, M. (1984). Two Faces of Intention. *The Philosophical Review*, 93(3), 375–405.
- Bratman, M. (1992). Shared Cooperative Activity. *The Philosophical Review*, 101(2), 327–341.
- Bratman, M. (1993). Shared Intention. *Ethics*, 104(1), 97–113.
- Bratman, M. (1999a). *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge: Cambridge University Press.
- Bratman, M. (1999b). I Intend that We J. In *Faces of Intention: Selected Essays on Intention and Agency* (pp. 142–162). Cambridge: Cambridge University Press.
- Castaneda, H.-N. (1967). Comments on Donald Davidson's 'The Logical Form of Action Sentences'. In N. Rescher (Ed.), *The Logic of Decision and Action* (pp. 104–112). Pittsburgh: University of Pittsburgh Press.
- Chan, D. K. (1995). Non-Intentional Actions. *American Philosophical Quarterly*, 32(2), 139–151.
- Chan, D. K. (1999). A Not-So-Simple View of Intentional Action. *Pacific Philosophical Quarterly*, 80(1), 1–16.
- Copp, D. (1980). Hobbes on Artificial Persons and Collective Actions. *Philosophical Review*, 89(4), 579–606.
- Copp, D. (2006). On the Agency of Certain Collective Entities: An Argument from 'Normative Autonomy'. *Midwest Studies in Philosophy: Shared Intentions and Collective Responsibility*, 30, 194–221.
- Davidson, D. (2001a). Actions, Reasons, and Causes. In *Essays on Actions and Events* (2nd ed., pp. 3–20). New York: Clarendon Press.
- Davidson, D. (2001b). Agency. In *Essays on Actions and Events* (2nd ed., pp. 43–62). New York: Clarendon Press.
- Davidson, D. (2001c). The Logical Form of Action Sentences. In *Essays on Actions and Events* (2nd ed., pp. 105–121). New York: Clarendon Press.
- Davis, W. (1984). A Causal Theory of Intending. *American Philosophical Quarterly*, 21, 43–54.
- Gilbert, M. (1990). Walking Together: A Paradigmatic Social Phenomenon. *Midwest Studies in Philosophy*, 15, 1–14.
- Gilbert, M. (1992). *On Social Facts*. Princeton: Princeton University Press.
- Gilbert, M. (1996a). Concerning Sociality: The Plural Subject as Paradigm. In J. Greenwood (Ed.), *The Mark of the Social*. Rowman and Littlefield.
- Gilbert, M. (1996b). *Living Together: Rationality, Sociality, and Obligation*. Lanham: Rowman & Littlefield Publishers.

- Gilbert, M. (2000). *Sociality and Responsibility: New Essays in Plural Subject Theory*. Lanham: Rowman and Littlefield.
- Ginet, C. (1989). Reasons Explanation of Action: An Incompatibilist Account. *Philosophical Perspectives*, 3, 17–46.
- Ginet, C. (1990). *On Action*. Cambridge: Cambridge University Press.
- Goldman, A. I. (1976). *A Theory of Human Action*. Princeton, N.J.: Princeton University Press.
- Harman, G. (1976). Practical Reasoning. *Review of Metaphysics*, 79, 431–463.
- Hornsby, J. (1980). *Actions*. London: Routledge & Kegan Paul.
- Jackendoff, R. (1987). The Status of Thematic Relations in Linguistic Theory. *Linguistic Inquiry*, 18(3), 369–411.
- Lepore, E., & Ludwig, K. (2003). Outline for a Truth Conditional Semantics for Tense. In Q. Smith & A. Jokic (Eds.), *Tense, Time and Reference* (pp. 49–105). Cambridge: MIT Press.
- Ludwig, K. (1992). Impossible Doings. *Philosophical Studies*, 65(3), 257–281.
- Mackie, J. L. (1974). *The Cement of the Universe: A Study of Causation*. Oxford: Clarendon Press.
- McCann, H. (1991). Settled Objectives and Rational Constraints. *American Philosophical Quarterly*, 28, 25–36.
- Melden, A. I. (1961). *Free Action*. London: Routledge & Kegan Paul.
- Mele, A. R., & Moser, P. K. (1994). Intentional Action. *Noûs*, 28(1), 39–68.
- O’Shaughnessy, B. (1973). Trying (as the Mental ‘Pineal Gland’). *The Journal of Philosophy*, 70(13), 365–386.
- Parsons, T. (1985). Underlying Events in the Logical Analysis of English. In E. Lepore (Ed.), *Actions and Events: Perspectives on the Philosophy of Donald Davidson* (pp. 235–267). Oxford: Blackwell.
- Parsons, T. (1990). *Events in the Semantics of English: A Study in Subatomic Semantics*. Cambridge: MIT Press.
- Parsons, T. (2000). Underlying States and Time Travel. In J. Higginbotham (Ed.), *Speaking of Events*. New York: Oxford University Press.
- Pettit, P. (2003). Groups with Minds of their Own. In F. Schmitt (Ed.), *Socializing Metaphysics* (pp. 167–193). Rowman and Littlefield.
- Pietroski, P. M. (2000). *Causing Actions*. Oxford: Oxford University Press.
- Pietroski, P. M. (2003). Semantics and Metaphysics of Events. In K. Ludwig (Ed.), *Donald Davidson*. Cambridge: Cambridge University Press.
- Schein, B. (1993). *Plurals and Events*. Cambridge: MIT Press.
- Schein, B. (2002). Events and the Semantic Content of Thematic Relations. In G. Preyer & G. Peter (Eds.), *Logical Form and Language* (pp. 261–344). Oxford: Clarendon Press.
- Searle, J. (1983). *Intentionality*. Cambridge: Cambridge University Press.
- Searle, J. (1990). Collective Intentions and Actions. In P. R. Cohen, J. Morgan & M. E. Pollack (Eds.), *Intentions in Communication*. Cambridge, Mass.: MIT Press.
- Smith, A. (1976). *An Inquiry into the Nature and Causes of the Wealth of Nations*. Oxford: Clarendon Press.
- Stoutland, F. (1997). Why are Philosophers of Action so Anti-Social? In Alanen, Heinamaa & Wallgreen (Eds.), *Commonality and Particularity in Ethics* (pp. 45–74). New York: St. Martin’s Press.
- Tollefsen, D. (2002). Collective Intentionality and the Social Sciences. *Philosophy of the Social Sciences*, 32(1), 25–50.
- Tuomela, R. (1991a). Intentional Single and Joint Action. *Philosophical Studies* 62(3), 235–262.
- Tuomela, R. (1991b). We Will Do It: An Analysis of Group Intentions. *Philosophy and Phenomenological Research*, 249–277.

- Tuomela, R. (2005). We-Intentions Revisited. *Philosophical Studies*, 125(3), 327–369.
- Tuomela, R., & Miller, K. (1988). We-Intentions. *Philosophical Studies*, 53(3), 367–389.
- Velleman, D. (1997). How to Share an Intention. *Philosophy and Phenomenological Research*, 57(1), 29–49.
- Vermazen, B. (1993). Objects of Intention. *Philosophical Studies*, 71(3), 223–265.
- Wilson, G. M. (1985). Davidson on Intentional Action. In E. LePore (Ed.), *Actions and Events: Perspectives on the Philosophy of Donald Davidson* (pp. 29–43). Oxford: Blackwell.
- Wilson, G. M. (1989). *The Intentionality of Human Action*. Stanford, Calif.: Stanford University Press.