

FOUNDATIONS OF SOCIAL REALITY IN COLLECTIVE INTENTIONAL BEHAVIOR

Kirk Ludwig

INTRODUCTION

The foundation of any account of social reality is an account of the nature of collective behavior, and in particular collective intentional behavior. Social practices and social interaction of any sort involve some form of collective intentional behavior essentially. If we do not understand the latter, we do not understand the former.

Collective intentional behavior is not just the aggregate or summation of individual intentional behavior. John Searle provides a vivid example of this in his paper “Collective Intentions and Actions” (1990; all citations to page numbers alone are to this article in the reprint in Searle 2002).¹

Imagine that a group of people are sitting on the grass in various places in a park. Imagine that it suddenly starts to rain and they all get up and run to a common, centrally located shelter. Each person has the intention expressed by the sentence “I am running to the shelter.” But for each person, we may suppose that his or her intention is independent of the intentions and behavior of others. In this case there is no collective behavior; there is just a sequence of individual acts that happen to converge on a common goal. Now imagine a case where a group of people in a park converge on a common point as a piece of collective behavior. Imagine that they are part of an outdoor ballet where the choreographer calls for the entire corps de ballet to converge on a common point. We can even imagine that the external bodily movements are indistinguishable in the two cases; the people running for shelter make the same types of bodily movements as the ballet dancers. (p. 92)

The same movements, by the same people, may, on different occasions, be something they do intentionally as a group or only as individuals. The fundamental difference must lie with the differing intentions (and perhaps other attitudes) of the individuals on the two occasions, for everything else is constant.

To set out the basic issues that arise in the project of analyzing collective intentional behavior, let us begin by contrasting [1]–[3] with [4]–[6].

¹ This paper provides the fullest description of Searle’s account of collective intentions. Later discussions, in (Searle 1995, pp. 23–26) and (Searle 1998, pp. 118–120), are shorter summaries of the argument of (Searle 1990).

- | | |
|--|---|
| [1] I intend to sing the national anthem | [4] We intend to sing the national anthem |
| [2] I sang the national anthem | [5] We sang the national anthem |
| [3] I sang the national anthem intentionally | [6] We sang the national anthem intentionally |

There is a reading of each of [4]–[6] on which they express merely that each member of the group picked out by ‘we’ intended to sing the national anthem, sang it, and sang it intentionally, that is, on which each member of the group could say truly [1]–[3]. But there is also another reading of each of [4]–[6] which implies that what we intend we intend jointly, and what we did and did intentionally we did together and did together intentionally. I will call the former reading of [4]–[6] the distributive reading and the latter the collective reading. To take [5] as an example, the distributive reading can be represented as in [5’].

[5’] Each of us sang the national anthem.

If [5] on its collective reading is true, then [5’] is true, but not vice versa. To understand collective intentional behavior, we need to understand the difference between what [1]–[3] express on the one hand, and what [4]–[6] express, on their collective reading, on the other. This can also be put as the problem of understanding the difference between the distributive and collective readings of [4]–[6].

There are two main problems in the philosophy of collective action.

First, what is the ontology of collective intentional action? In particular, must we admit the existence of irreducibly group or plural agents over and above individual agents? A powerful argument for the existence of group agents derives from the observation that the only difference between [1]–[3] and [4]–[6], on their collective reading, appears to be the presence of a plural subject term in [4]–[6] in place of the singular subject term in [1]–[3]. As the referent of the subject term in [1]–[3] is the intender and agent of the action in the individual case, so it seems, by analogy of form, we must accept that the referent of the subject term in [4]–[6] is the intender and agent in the collective case. Some philosophers have embraced this conclusion and sought to explain how we can make good sense of it (see, e.g., Copp 2006; Gilbert 2000; Tollefsen 2002; Stoutland 1997). With Searle, I think this conclusion should be avoided if at all possible. We start with individuals and their intentions. In certain circumstances, typically those in which they wish to cooperate with one another, individuals conform their intentions to a common purpose. In these circumstances (which remain to be precisely delineated) we say they did something together intentionally. The difficulty is to see how to understand sentences such as [4]–[6] on the collective reading compatibly with this individualistic perspective.

Second, what is the psychology of collective intentional action? There are two principal subquestions that arise here.

First, if there are irreducibly group or plural agents, then it seems we must attribute to them psychological states such as intention, and presumably belief and desire. How can we make sense of this? Along with Searle, I accept a form of the connection principle, according to which nothing is a mental state unless it is a conscious mental state or a disposition, *inter alia*, to produce a conscious mental state. While it seems conceptually possible that there should be a conscious mind supervening on certain groups of individuals in certain circumstances, I think that we have no reason to believe that anything of the sort occurs or that we suppose anything of the sort when we use collective action sentences. This provides one of the most powerful reasons against admitting irreducibly group or plural agents.

Second, even if we eschew collective agents, there is an important and difficult question about the psychology of collective intentional behavior. I will call the sort of intention that an individual has in performing a singular action an I-intention. [1] expresses an I-intention, for example. Whatever is involved in collective action, it involves among other things individuals doing things with intentions. I will call the intention an individual has in participating in a group's doing something intentionally a we-intention. Thus, a we-intention is not an intention of a group, but of an individual who is a member of a group. If [4] is true on its collective reading, then each of the members of the group has a we-intention directed toward singing the national anthem. We need this notion even if we think there are group agents with intentions, for, of course, no one thinks that collective intentional behavior occurs without individual members of the group having intentions connected with their actions which contribute to what the group does. Our second question then is: what is the difference between I-intentions and we-intentions? In particular, are we-intentions a distinctive sort of intention? Can we understand we-intentions in terms of the concepts which we already deploy in understanding individual intentional action?

I will argue that once we understand properly the logical form of plural action sentences on their collective reading, we will be able to see that they do not carry any commitment to irreducibly plural or group agents and that all the resources needed for understanding we-intentions are already available in our understanding of individual intentional action. This is not the same as saying that we can analyze we-intentions in terms of I-intentions—or even I-intentions with mutual beliefs about a common goal. It is to say, however, that no conceptual resources not already deployed in understanding I-intentions are required to understand we-intentions.

I begin by looking at Searle's original and distinctive account of collective intentional behavior. It was through Searle's seminars at Berkeley in the late 1980s that I was originally introduced to the problem of collective intentional action. There is much that I admire and agree with in Searle's account, and I have learned a lot from it. There are also, however, some aspects of it which are not completely clear to me. In addition, if I understand his account correctly, it locates the difference between I-intentions and we-intentions in the

psychological mode of these attitudes rather than in their contents. As I understand it, it is for this reason that Searle holds that collective intentionality is a “primitive phenomenon” (p. 90). I am more optimistic that we can capture what is distinctive about we-intentions by a difference in contents of a single psychological mode.

The paper divides into two parts. In the first part, I look at Searle’s account and the motivations for it. I will aim to say how I understand it, and raise some questions about it. This will set the stage for the second part, in which I will present my positive proposal, for it will give us a set of constraints and problem cases which any adequate account should be able to meet. My method will be to take a close look at the proper way of projecting the standard account of the logical form of singular action sentences to plural action sentences in light of the ambiguity between collective and distributive readings of plural action sentences, and then extend this to the analysis of plural intention sentences, such as [4], on their collective reading, to arrive at an account of the structure and content of we-intentions. This will provide us with an account which shows what is distinctive about we-intentions, meets our constraints, takes care of problem cases, and deploys only notions already deployed in our understanding of individual intentional action.

1. SEARLE’S ACCOUNT OF WE-INTENTIONS

Searle places two constraints on any adequate account of we-intentions.

Constraint 1

It must be consistent with the fact that society consists of nothing but individuals. Since society consists entirely of individuals, there cannot be a group mind or group consciousness. All consciousness is in individual minds, in individual brains.

Constraint 2

It must be consistent with the fact that the structure of any individual’s intentionality has to be independent of the fact of whether or not he is getting things right, whether or not he is radically mistaken about what is actually occurring. (p. 96)

I think both of these constraints are correct and important. I think the best defense of them against objectors is to show that we can understand collective intentional behavior compatibly with them, for no one would be inclined to deny either one except insofar as they were thought to be incompatible with facts about collective action.

Searle’s account of collective intentional behavior aims to meet constraints 1 and 2 compatibly with the thesis that

We-intentions cannot be analyzed into sets of I-intentions, even I-intentions supplemented with beliefs, including mutual beliefs, about the intentions of other members of a group. (p. 93)

I will call this *the irreducibility thesis*.

A question of interpretation arises about the irreducibility thesis, connected with the interpretation of 'I-intentions'. I believe that my characterization of 'we-intentions' as intentions that individuals have in participating in group intentional behavior conforms to Searle's use. There may be a difference in our use of 'I-intentions', however. I have characterized I-intentions as intentions individuals have when involved in individual, as opposed to group, intentional behavior. In some places Searle uses 'I-intention' compatibly with the way in which I have characterized it. However, in some places there is the suggestion that Searle intends 'I-intention' to cover any intention expressed with a sentence of the form 'I intend. . . .' These are not equivalent because the former way of characterizing 'I-intentions' leaves it open that what distinguishes I-intentions from we-intentions has to do with what replaces '. . .' in sentences of the form 'I intend. . . .' My positive proposal in the second part of the paper takes this route. This is important because it determines whether I can accept the irreducibility thesis as Searle intends it. I am inclined, on the whole, to think that Searle intends the official characterization of 'I-intentions' to go by way of reference to the intentions individuals have when engaging in singular action, rather than by way of what sentences of a certain form express. This does not build into the categories in which the discussion is conducted a view about what distinguishes I-intentions from we-intentions. In this case, I can accept the irreducibility thesis as Searle intends it.

Searle's account of we-intentions is developed in two stages, first briefly and informally, and then more fully using the resources first made available in his *Intentionality* (1983).

The informal account is

. . . there are intentions whose form is: We intend that we perform act *A*; and such an intention can exist in the mind of each individual agent who is acting as part of the collective. In cases like that of the football team each individual will have further intentional content, which in ordinary English he might express in the form "I am doing act *B* as part of our doing act *A*." For example, "I am blocking the defensive end as part of our executing a pass play." We need only note that all the intentionality needed for collective behavior can be possessed by individual agents even though the intentionality in question makes reference to the collective. (pp. 96–97)

On this account, as I understand it, there are two components to the intentionality of individuals participating in group intentional action. First, each individual who is a member of a group which does something intentional has an intention of the form 'We intend that we perform act *A*'. Second, each also has an intention he or she could express using the sentence 'I intend to perform act *B* as part of our performing act *A*' (plus appropriate beliefs). This is the "further intentional content."

The first part of this, however, is not fully clear to me. In saying that 'We intend that we perform act *A*' is the form of an intention each individual has, it is unclear what state is being attributed to each individual, none of whom

is identical to the group. Each individual could think *we intend that we perform act A*, and intend that he does his part in their performing act A. But this attributes only one *intention* to each of the agents together with a *belief* that they intend to do something. So it is not equivalent to what is said above. Since any intention that an individual has should be attributable to him in a sentence in which the subject term refers to him, I suggest, at least for the moment, that we take it that the intention that each individual has in virtue of participating in a group that does something intentionally could be expressed by him by a sentence of the form ‘I intend that we perform act *A*’, or, somewhat more generally, ‘I intend that we *F*’ where *F* is replaced by a verb phrase whose main verb is an action verb. I believe this to be compatible with Searle’s detailed development of the account in the second stage. If we interpret the above passage in this way, then it is also compatible with reading the irreducibility thesis with my characterization of I-intentions and we-intentions.

The second stage of Searle’s account aims to say in more detail what the form is of the intentions that individuals who participate in collective intentional action have. If what I have just said is right, then this is a matter of making clearer the form of the intention expressed by ‘I intend that we *F*’. Searle’s method, which seems to me to be the right one, is to start with the form of singular intentions and ask how it is to be extended to the collective case, compatibly with our two constraints. Searle focuses on intention-in-action, but I think the same issues arise for prior intention, so I will just use the verb ‘intend’ in what follows. Searle uses the notation ‘*S(p)*’, where *S* stands for the type of psychological state, and *p* for its content. This leaves it unclear whose psychological state we are talking about. In the case of collective intentional behavior, it will be helpful to make this explicit. I will therefore use the following notation, which is equivalent to Searle’s, except for making explicit whose state we are talking about: *x* intends (*p*).

Let us take now Searle’s example of someone who intends to fire a gun by pulling the trigger. In *Intentionality*, Searle introduces a principle for identifying the content of a propositional attitude with mind-to-world direction of fit (e.g., belief) or world-to-mind direction of fit (e.g., desire). The content of the attitude is given by the proposition which has to be true for the propositional attitude to be satisfied, where this is a general term for the positive valuation of a bivalent evaluation of the attitude. For example, a belief is satisfied if it is true, and an intention is satisfied if it is carried out. Then to identify the full content of an intention, we must ask under what conditions it is carried out or not. This is an important principle, which I will make use of in my positive proposal. I will call this *the satisfaction principle*. Searle argues that using this principle reveals an important component of the content of an intention which is not represented in the complement clause of reports of intentions. Consider ‘John intends to fire the gun by pulling the trigger’. Since the attitude expressed has a bivalent evaluation condition, its content must be given by a full sentence. In this case, the infinitival phrase indicates that the subject of the complement clause is the same as the subject of

the main clause, namely 'John'. Representing this, we get, 'John intends (he fire the gun by pulling the trigger)', where the 'he' is understood as coreferring with 'John'. But it is not enough for John to carry out his intention that as a matter of fact he fires the gun by pulling the trigger. He may do that accidentally, for example, while reaching into a drawer to retrieve a pair of scissors. For him to carry out his intention, that very intention must play a causal role in what he does. If it is a prior intention, it must lead to what Searle calls an intention-in-action that is coeval with the action, directly about it, and which guides it while it is being carried out. That intention-in-action must itself have the same feature, i.e., it must, to be carried out, be causing the action which it is about.

Thus, we need to fill in the content of the intention above to accommodate this. As Searle articulates it, it goes as follows: John intends (this intention causes it to be the case that the trigger pulls, which causes it to be the case that the gun fires). I understand 'trigger pulls' and 'gun fires' here as indicating event types which could occur independently of any agents.

The question, then, in light of the method, is how this is to be extended to 'We intend to *F*'. As I understand Searle, this in turn is the question how this is to be extended to '*x* intends that we *F*' where *x* takes on as a value an individual, not a group.

Take now an example (Searle's example). Jones and Smith are preparing a hollandaise sauce. They cooperate. Jones stirs while Smith pours, and they intend thereby to mix the sauce. What are the intentions each has in doing this in virtue of which they are doing it together intentionally? As Searle notes, it is not enough that Jones intends (this intention causes it to be the case that the ingredients are stirred) and Smith intends (this intention causes it to be the case that the ingredients are poured), for this is compatible with, for example, Jones noticing that Smith is pouring the ingredients for hollandaise sauce and opportunistically stirring. This could occur without Smith and Jones intentionally making hollandaise sauce. Smith might care less or even fail to attend to what Jones is doing and so not realize they are making hollandaise sauce. It does not help if we add that Jones/Smith intends (this intention causes it to be the case that the ingredients are stirred/poured, which causes it to be the case that the sauce is mixed). For this could fit a scenario in which each thinks he is opportunistically taking advantage of what the other is doing, and, in that case, they would not be making hollandaise sauce together intentionally.

Searle suggests and rejects another proposal, on which there is "a special class of collective intentions and that is all that is needed" (p. 101). He continues, "On this account, from Jones's point of view the intentionality is this:

collective i.a. (this collective i.a. causes: ingredients are stirred,
causes: sauce is mixed)."

The content here is read: this collective intention-in-action causes it to be the case that the ingredients are stirred, which causes it to be the case that the sauce is mixed. Since on Searle's view only individuals have intentionality, I take it that we are to think that each of Jones and Smith has such an intention, so that we

could put it this way (using ‘we-intend’ in the place ‘collective i.a.’): Jones/Smith we-intend (this we-intention causes it to be the case that ingredients are stirred/poured, which causes it to be the case that the sauce is mixed). The complaint Searle has with this is that “it leaves out the fact that Jones/[Smith] is making an individual contribution to a collective goal . . . this account leaves it as mysterious how the collective intentionality can move the body”; each has to “personally . . . intend to do something if the sauce is ever going to get mixed” (p. 101). Since these intentions are intentions of the individuals, it is not clear to me why *how it moves them to action* should be more mysterious than how singular intentions move them to action. However, perhaps the difficulty may be put in this way: intuitively, when Jones and Smith each participates in their mixing the sauce, each of them does have particular individual intentions directed at their parts of what they are doing together, and this is a general feature of group intentional behavior. So what is left out of this account is the role of these individual intentions to do their parts in the collective action.

If individual intentions (or singular intentions) to do their parts in a group action are always present when individuals act as parts of groups, then any we-intention they have (taking this now as a *sui generis* type) would be satisfied only if the group action is brought about by way of such singular intentions. This suggests then reading into the content of the we-intention that it bring about such singular intentions, and this is the next proposal that Searle considers: Jones/Smith we-intends (this we-intention causes it to be the case that Jones/Smith has an I-intention which causes it to be the case that the sauce is stirred/poured, which causes it to be the case that the sauce is mixed).

The difficulty with this, Searle suggests, is that it misrepresents the relation between the collective intention and the singular intention. If I understand, the basic charge is that intuitively when I intend, in the relevant sense, that we do something, it is not part of what I intend that I form further singular intentions, even if I do form further singular intentions as a result. This seems right. Further intentions arise as a result of our intentions through rational deliberation given that one has settled on the intended end, not as a result of being part of what is intended. So we have still not quite got hold of the right relation between the collective intention and what “I have personally to intend to do” (p. 101).

The next move that Searle makes is original and clever, but I also find it puzzling. Let me try to explain carefully how I think it goes.

We begin by rethinking the account of a singular intention that involves a by-means-of relation, such as the intention expressed by: *x* intends (this intention causes it to be the case that the trigger pulls, which causes it to be the case that the gun fires). Searle suggests that this does not fully represent the form of the intention. He suggests that this “isn’t just any old type of i.a. [intention-in-action], it is an achieve-B-by-means-of-A type of i.a.” and that we

. . . might think of the notation that represents this type of i.a. as containing two free variables, “A” and “B”; and these variables are then bound by clauses inside the brackets that function as nouns. What we are trying to say is that I have an achieve-B-by-means-of-A sort

of intention whose content is that that-the-trigger-pulls-as-A causes it to be the case that-the-gun-fires-as-B. (p. 102)

Searle suggests we can represent this as

i.a. B by means of A (this i.a. causes: A trigger pulls, causes: B gun fires)

I wish to rewrite this to try to make more perspicuous what I take the suggestion to be.

$(\exists A)(\exists B)(\text{Jones/Smith intends B by means of A (this intention causes a trigger pulling A, which causes a gun firing B)})$

I am not sure this fully corresponds to Searle's intentions. My reasons for this representation are as follows. I have introduced explicit existential quantifiers with wide scope because we need to bind variables both inside and outside the content clause, and so the scope of the binding operators has to cover both. I have not included as part of the quantifiers any restriction on a type (e.g., ' $(\exists A: A \text{ is a trigger pulling})$ '), since that A is a trigger pulling and B a gun firing is properly part of the content of the intention, which may misfire, after all. One might execute the intention but be opening a drawer by pulling on its handle instead of firing a gun by pulling its trigger. One could easily modify this if it were thought to matter. I will postpone further comment on this until after we have seen how it is to be used to help with the problem of we-intentions.

The suggestion based on this view of singular by-means-of intentions is that we think of the we-intention as a by-means-of intention as well, but one where the means is a singular action—not a singular intention—and the goal is the collective goal, and that we write in the nature of the actions as collective and singular in the 'by means of' portion of the verb phrase, as follows:

i.a. collective B by means of singular A (this i.a. causes: A stirred, causes: B mixed)

Rewriting this in accordance with my suggestion above, we have:

$(\exists A)(\exists B)(\text{Jones/Smith intends collective B by means of singular A (this intention causes a stirring/pouring A, which causes a mixing B)})$

This accomplishes two things: (i) It keeps the idea that the collective intention involves the idea that a collective action is accomplished by singular actions. (ii) It avoids putting this into the content of the intention and makes it rather a modification of the type of intention expressed. (i) takes care of the idea that the collective intention requires for its satisfaction the formation of an individual intention. It does this not by making explicit mention of a singular intention as a means, but of a singular action as a means, which in turn will require the formation of a singular intention. (ii) secures that the type of intention is distinct from singular intentions without relying on a difference in content. This looks desirable because if we put this into the content, as in

Jones/Smith intends (this intention causes it to be the case that there is singular stirring/pouring, which causes collective mixing),

it would again seem subject to counterexamples in which each individual takes opportunistic advantage of what the other is doing, without their intending in the appropriate sense to do something together.

I admire the creativity of Searle's suggestion and how he works his way to this proposal. Yet, I am not at all certain that I understand it. And I have a number of reservations about it. I will now explain what these are. This will motivate an additional constraint on a proper solution. And I will try to meet all the constraints in the second part of this paper.

A key feature of the proposal is to locate the difference between I-intentions and we-intentions not in their content but in the kind of state expressed. In considering this, I want to start first with a proposal Searle rejects, before considering his own proposal. I will raise some objections to this proposal different from those we considered above, which I think don't apply fully to Searle's own, but which will help to set the stage for its discussion.

The proposal that Jones, for example, we-intends (this intention causes it to be the case that the sauce is stirred, which causes it to be mixed) locates the difference with I-intentions not in the content of the intention but in the psychological mode. A first problem here is that this seems more to label a problem than to solve it, for we are left with the question what the difference is between such a we-intention—which is a psychological state had by an individual—and an I-intention, if it does not lie in the content. What account is to be given of this? A second problem is that intuitively the sort of state that individuals have when engaging in intentional cooperative behavior with others is not of a distinct type. Imagine that you are intentionally pushing a car, and then someone else comes and joins you, and you are now pushing the car together intentionally. What happens? There are some changes of course. You come to believe that the other person intends to push the car with you, you believe you are pushing it with him, you accept his help, and you continue pushing it as you were before, intending now to do it with him. But is it really plausible that there is some change in the psychological mode of the intention? If we can find any other solution, it would be preferable. A third problem is that if there were genuinely a distinct kind of intention involved, it would seem plausible that this would be reflected in a distinct verb form, for in other cases where we have differences in psychological modes, we also have a difference in vocabulary to express them. But this requires us to suppose that there is a kind of ambiguity in the verb 'intend' which is not made explicit syntactically, and not just in English. This is not a knockdown argument, but it does give us reason to pause, and to look for another solution. Fourth, from the standpoint of practical deliberation, reasoning about whether to join with others to do something looks to be the same sort of process as reasoning about what to do generally. It looks as if it ought to issue in an intention to do something of the same sort as well—with a content

appropriate for the practical problem and its solution—and then if others decide similarly, we have a group of people who intend to do something together. It looks, in other words, from the standpoint of practical deliberation, like a single type of intention is involved as the issue of practical reasoning about either individual or group action.

However, Searle's final proposal has some important advantages over this initial one. For one of the charges I have leveled above is that it makes it look as if the we-intention is *sui generis* and not continuous in any intelligible way with singular intentions, and therefore more a label for a problem than a solution to one. But the final proposal in fact represents the form of collective and singular intentions as the same except for a small point. It seeks to make intelligible what the difference is by appeal to a difference in one point in a common structure. Each has the form

x intends . . . by means of ___ (p)

The difference resides in what fills in the ' . . . '. In the case of a means-end singular intention, it is 'B' for the first and 'A' for the second, or, we might even write, 'singular B' and 'singular A'. Then for the collective case, the only difference lies in replacing 'singular B' with 'collective B'. This makes the account of we-intentions continuous with the account of singular intentions, and represents sentences attributing we-intentions as sharing the same logical form as sentences attributing singular intentions.

Why am I dissatisfied then? First, I am doubtful of the motivation for the representation of the singular intention as involving a modification of the verb. An intention is a commitment to carrying out a plan of action. If the plan is to bring about B by doing A, to fire a gun by pulling its trigger, then we have said enough about what intention someone has when we represent him as intending (this intention causes it to be the case that a trigger moves which causes it to be the case that a gun fires). The type of intention is not different when the agent simply intends to pull the trigger but not to fire the gun by doing that. If we are not justified in representing the singular means-end intention in the special form Searle suggests, then the suggestion for the collective case is not a simple step from the individual case, and the objections canvassed above to the first proposal apply.

Second, the existential quantifiers that, if my representation of the proposal is correct, are required would make it impossible for someone to intend to do B by doing A without there being some actions corresponding to 'B' and 'A': but if we allow that a brain in a vat could have the same intentionality as any of us, this cannot be a requirement on our having the intentions we have. If I intend to fire a gun by pulling a trigger, and do so, then the values of 'B' and 'A' for me are the firing of the gun and the trigger pulling; but there are no such actions to appeal to for a brain in a vat, though it could have the same intentions I do. This worry extends to the collective intention case. It may be that I have read into Searle's notation something he does not intend. But in that case I am uncertain how to understand the proposal.

Third, setting aside these worries, it is not clear to me that the proposal solves the problem. Since the intentionality in collective intentional behavior consists entirely in states of individuals, the proposal is that when members of a group of people all have intentions with appropriately corresponding contents, they collectively intend to do something, and then if they all carry out their intentions, they do something together intentionally. The intentions, on this proposal, which they all have, are of the form

$(\exists A)(\exists B)(x \text{ intends collective B by means of singular A (this intention causes a F-ing A, which causes a G-ing B)})$

where the same thing goes in for 'B' for each member of the group and different things typically for 'A'. The difficulty I have in mind is that this seems to allow that it could accidentally happen that a group of people have intentions of this form, and do something together, though none of them know about any of the others, and their reasons for thinking that they were doing something with others were ungrounded, confused, or mistaken. Let us suppose that, at three institutions for the criminally insane, in three widely separated locations, three different inmates, Jules, Jim, and Jill, each coincidentally dream of being visited by the angel of God and being told that he or she is to cooperate in a mission to blow up the Golden Gate Bridge. Jules believes he has been assigned the job of planting explosives at the base of the north tower; Jim believes that he has been assigned the job of planting explosives at the base of the south tower; Jill believes that she has been assigned the job of planting explosives halfway along the main span. They all believe that they are working with others, who are to remain unknown to them, to prevent any of them giving away the others if captured. They each make their escapes, make their way to San Francisco, set explosives, coincidentally on the same day, but without coming in contact with each other, and then set them off, again, coincidentally, at about 5 p.m. in the evening, blowing up the Golden Gate Bridge together. Yet, intuitively, they did not do this together intentionally. Something still is missing.

Fourth, is it even clear that the account handles the case of Jones and Smith each taking advantage, as each thinks, of what the other is doing, to mix the hollandaise sauce? The reason the question arises is that in this case each is undoubtedly aiming at their doing something together, i.e., at a collective action, only not one that is collectively intended. They each expect that they will mix the hollandaise sauce, but neither thinks they will do it together intentionally. So they each have a means-end intention to bring it about that they do something by his doing something. It looks then as if, in line with Searle's proposal for how to represent means-end intentions, they each have an intention of the form

i.a. collective B by means of singular A (this i.a., causes: stirring/pouring A, causes mixing B)

But they are not mixing the hollandaise sauce together intentionally.

I do not think that these concerns that I have raised are decisive. Some may rest on misunderstandings. Others may be met by modifying the account in

some way. But I am inclined to ask whether a different approach may prove more fruitful. One thing that is central to my concern about Searle's approach is that it works by way of introducing a distinctive collective mode of intending. Other things being equal, it seems to me that we should look for a way to understand collective intentional behavior without introducing a special mode of intending. The only alternative is to try to understand the difference in terms of the contents of the intentions—and perhaps beliefs—of individuals who are engaged in collective intentional behavior. I therefore propose a third constraint:

Constraint 3

An account of we-intentions must locate the difference between I-intentions and we-intentions in the contents of the intentions, not in their mode.

As in the case of the first two constraints, beyond its intuitive appeal, I think the best defense of this one is to show that it can be met, and I undertake that next.

2. THE LOGICAL FORM OF PLURAL ACTION SENTENCES²

Let us return to the contrast between [1]–[3] and [4]–[6] (repeated here) with which we began.

- | | |
|---|--|
| [1] I intend to sing the national
anthem | [4] We intend to sing the national
anthem |
| [2] I sang the national anthem | [5] We sang the national anthem |
| [3] I sang the national anthem
intentionally | [6] We sang the national anthem inten-
tionally |

I follow Searle in thinking that the proper way to approach the question of the structure of plural action sentences is to determine the proper thing to say about singular action sentences and see what the implications of this are for the plural case. I also follow Searle in rejecting the attribution of psychological states to groups, and so of intentional actions to groups. Further, I accept with Searle the satisfaction principle, and so that the content of an intention is to be determined by what conditions must be met for it to be said properly to be carried out.

The account I will give rests on two things. The first is the event analysis of singular action sentences. The second is the observation made at the outset that plural action sentences are ambiguous between a distributive and collective reading.

I will begin with the analysis of [2] and how to project that to [5]. Then I will turn to [3] and [6], and finally to the vexed case of [1] and [4].

The event analysis of singular action sentences treats [2] as introducing an implicit existential quantifier over events. It was introduced to explain how adverbs contribute to determining the truth conditions of action sentences in Donald Davidson's 1967 paper "The Logical Form of Action Sentences"

² In the following I draw on (Ludwig, forthcoming).

(reprinted in (Davidson 2001)). The action verb is treated as having an extra argument place for an event which is bound by the existential quantifier, and the adverbs are treated as contributing predicates of the event variable bound by the quantifier. This explains their compositional structure in terms of standard logical devices. This is now the standard analysis of the logical form of action sentences, and, more generally, of sentences whose main verb is an event verb. In linguistics, it is standard to separate out different case or thematic roles for arguments as conjuncts in the scope of the existential quantifier, and I will follow that practice in the following. In addition, I will explicitly require in the case of singular action sentences that there be a unique agent of the event expressed by the action verb. To keep the formulations less complex, I will not in the following incorporate an analysis of the tense structure of the sentences under analysis. This can be added without affecting any of the conclusions reached below (see Lepore and Ludwig 2003). The resulting analysis of [2] is given in [2a].

[2a] $(\exists e)(\text{singing}(e) \text{ and } \text{agent}_{\chi}(e, I) \text{ and } (\text{no } y: y \neq \text{me})(\text{agent}_{\chi}(e, y)) \text{ and } \text{of}(e, \text{the national anthem}))$

I have subscripted ‘agent’ because action verbs may require a particular form of agency. I take the agency relation to be a determinable, which has various determinate forms. The subscripted ‘ χ ’ represents then a parameter which is set by the action verb which determines the particular form of agency required. The need for this is brought out by the contrast between causing someone to die and killing him. If one kills someone, one causes him to die, but it does not follow from the fact that one caused someone to die that one killed him. If I hire an assassin to kill a rival, I cause my rival to die, but it is the assassin who kills him. In this case, the verb ‘kill’ imposes some additional constraints on the relation of the agent to the event (in this case a death) beyond his just causing the death: the causing must not go primarily by way of another’s agency, i.e., it must be not be in this sense indirect. So we may say that ‘kill’ requires direct agency. There are at least these ways in which one can be an agent of an event: the event can be a primitive action of the agent’s, something he does but not by doing anything else (raising one’s arm); the event can be something which an action of the agent causes—of which a subvariety is direct causation (as in killing someone); the event can be one which something that the agent does constitutes in whole or in part (adjourning a meeting by saying ‘This meeting is hereby adjourned’); the event can be one which is part of something the agent does (contracting one’s forearm muscles by clenching one’s fist). This suggests a further analysis of ‘ $\text{agent}_{\chi}(e, x)$ ’ as ‘ $(\exists f)(\text{primitive-agent}(f, x) \text{ and } R_{\chi}(f, e))$ ’, which we will see is particularly helpful for thinking about plural action sentences. However, for the most part I will elide this further articulation of the agency relation (see Pietroski 2000 for discussion).

If we replace ‘I’ with ‘we’ in [2a], we get [5a], which looks like the most straightforward application of this analysis to the case of plural action sentences (using ‘Y’ as a variable that takes as values whatever we take plural terms

to refer to—as we have to make a change in the quantifier to reflect a change in the sort of referring term that appears on the right of the identity sign).

[5a] $(\exists e)(\text{singing}(e) \text{ and } \text{agent}_\chi(e, \text{we}) \text{ and } (\text{no } Y: Y \neq \text{us})(\text{agent}_\chi(e, Y)) \text{ and of}(e, \text{the national anthem}))$

This looks problematic because it seems to commit us to a group being an agent. But it is here that the observation that [5] is ambiguous plays a crucial role in arriving at the right analysis. On the distributive reading of [5], we understand it to be saying not that we sang together, but that, instead, *each of us* sang the national anthem. That is, we interpret ‘we’ as a restricted quantifier, a quantifier that ranges over members of the group. Thus, we can represent [5] on the distributive reading as

[Each x : x is one of us](x sang the national anthem),

where ‘ x ’ is a variable that takes individuals as values. ‘ x sang the national anthem’ gets the analysis represented in [2a]. So this comes out as:

[5b] [Each x : x is one of us]($\exists e)(\text{singing}(e) \text{ and } \text{agent}_\chi(e, x) \text{ and } (\text{no } y: x \neq y) (\text{agent}_\chi(e, y)) \text{ and of}(e, \text{the national anthem}))$

[5b] has two quantifiers in it, and we know that in natural languages where scope relations are not explicitly represented in a sentence typically we can read two quantifiers in a sentence in either order, as in ‘Everyone loves someone’, which can be read as ‘Someone is such that everyone loves him’ or ‘Everyone is such that there is someone he loves’. What then happens when we construe the ‘we’ in [5a] as ‘[each x : x is one of us]’ and give the event quantifier wide scope? The idea here will be that in effect we are saying that there is an event such that each of us is an agent of it, and no one who is not one of us is, and that it is a singing of the national anthem. This is straightforward except for how to handle ‘[no Y : $Y \neq \text{us}](\text{agent}(e, Y))$ ’. If we treat ‘us’ here as ‘[each x : x is one of us]’, which binds the argument place occupied by ‘us’, then we need to also have a “step-down” quantifier from the plural variable ‘ Y ’ so that we have a singular quantifier binding the position occupied by ‘ Y ’. Thus, we need to represent this as follows: ‘[no y : y is one of any Y & [for any x : x is one of us]($y \neq x$)](agent(e, y))’. This says no member of any group which is not a member of us is an agent of e . Let us abbreviate this as ‘[only $x \in \text{us}](\text{agent}(e, x))$ ’. Similarly let us abbreviate ‘[each x : x is one of us]’ as ‘[each $x \in \text{us}$]’. This gives us [5c].

[5c] $(\exists e)(\text{singing}(e) \text{ and } [\text{each } x \in \text{us}](\text{agent}_\chi(e, x)) \text{ and } [\text{only } x \in \text{us}](\text{agent}_\chi(e, x)) \text{ and of}(e, \text{the national anthem}))$

This says then: there is an event which is a singing of the national anthem of which each of us but no others are agents (in a particular way). And that is simply to say that it is something which we did together, i.e., this is the collective reading of [5]. This is motivated by two independently established semantic observations, one about the logical form of singular action sentences, and the

other about how we need to understand the contribution of ‘we’ to plural action sentences on the distributive reading. Putting these two things together, this account of the collective reading falls out. And it is evident that here we are not committed to any group agents at all. This shows that the argument from the similarity of surface form of [1]–[3] and [4]–[6] discussed in the introduction to plural agents is unsuccessful, for the surface form does not reveal fully the logical form; on the contrary, a proper analysis of the logical form of plural action sentences shows that we are not committed to any but individual agents.

So far so good. Let us turn now to [3] and [6]. In line with the motivation for the event analysis, it will be natural to take ‘intentionally’ in [3] to be introducing a predicate of the event variable introduced by the event quantifier, and the most straightforward approach is to treat ‘intentionally’ as introducing the verb ‘intend’ and the sentence it modifies as providing material which figures in the content of the intention, which will be an intention of the agent of the action at the time of the action. This yields:

[3a] $(\exists e)(\text{singing}(e) \text{ and } \text{agent}_\chi(e, I) \text{ and } (\text{no } y: y \neq \text{me})(\text{agent}_\chi(e, y)) \text{ and } \text{of}(e, \text{the national anthem}) \text{ and } I \text{ intend-}[\text{singing}(e) \text{ and } \text{agent}_\chi(e, I) \text{ and } (\text{no } y: y \neq \text{me})(\text{agent}_\chi(e, y)) \text{ and } \text{of}(e, \text{the national anthem})])$

The event quantifier binding the variable inside the scope of ‘intends’ represents the intention as a *de re* intention about the action (i.e., directly about and not via a description), and this is intuitively right, for one could have an intention to do something even at the time one does it without one doing it with that intention, for it might be directed at a different coeval event. I used the ‘[. . .]’ notation here rather than parentheses because I want to indicate that this is not yet a full analysis of the content of the intention. I will come back to this when we turn to [1] and [4]. A fuller analysis of [1] and [4] will then fill in what is left unarticulated in this.

The extension to [6] on the collective reading is straightforward:

[6a] $(\exists e)[\text{each } x \in \text{us}](\text{singing}(e) \text{ and } \text{agent}_\chi(e, x) \text{ and } [\text{only } x \in \text{us}](\text{agent}_\chi(e, x)) \text{ and } \text{of}(e, \text{the national anthem}) \text{ and } x \text{ intends-}[\text{singing}(e) \text{ and } \text{each } x \in \text{us}](\text{agent}_\chi(e, x)) \text{ and } [\text{only } x \in \text{us}](\text{agent}_\chi(e, x)) \text{ and } \text{of}(e, \text{the national anthem}))$

Here I have pulled the quantifier binding the second argument place in ‘agent_χ(e, x)’ out front to let it bind also the subject position of ‘intends’. I have not bound the argument positions for agents inside the scope of ‘intends’ with this initial quantifier, however, for that would, incorrectly, require agents who do something together intentionally to have an intention that just he be an agent of it.

I turn now to the real crux of the matter, [1] and [4]. In the case of [1], the complement in logical form should be represented with a sentence with the same subject as the superordinate clause, and, likewise, [4]. We have rejected genuinely plural agents, and so likewise the idea that we attribute intentions to such agents.

This means that in [4] ‘We’ should be read distributively, as ‘[each $x \in us$]’. This gives us [1a] and [4a], where in [4a] ‘us’ and ‘we’ are understood to corefer.

[1a] I intend that [I sing the national anthem]

[4a] [Each $x \in us$](x intends that [we sing the national anthem])

The matrix in [4a] has a singular variable in the subject position. When we instantiate [4a] to a member of the group, we get a sentence that attributes to him—on one reading of it—a we-intention. The analysis of we-intentions then comes down to the analysis of the matrix in [4a] on its collective reading. The same verb is evidently used in both [1] and [4]. Therefore, the difference must lie in a difference in the complement. A first observation is that evidently we must not give the complement sentence in [4a] its distributive reading, but rather its collective reading. This gives us [4b].

[4b] [Each $x \in us$](x intends that $[(\exists e)(\text{singing}(e) \text{ and } [each\ x \in us](\text{agent}_\chi(e, x)) \text{ and } [only\ x \in us](\text{agent}_\chi(e, x)) \text{ and } of(e, \text{the national anthem}))])$)

This is not our final analysis. But we can already observe how it handles some cases. This is enough to handle Bratman’s (1992: 333) case of my intending that we go to New York, as he puts it, in the Mafia sense—by way of my kidnapping you and putting you in the trunk. For in that case, I do not intend that each of us be an agent of our going to New York, but only me. This also is enough to handle Searle’s lovely case (p. 94) involving the graduates of a business school who have been taught Adam Smith’s theory of the invisible hand, and who each intend to help humanity by pursuing his own self-interest “without reference to anyone else”—knowing that each of the others so intends and likewise knows that the others do. For in that case, none of them intend that each of them be an agent of the helping, for what they intend they intend without reference to anyone else.

But the analysis is still not adequate. One problem we have to face, a problem which Searle highlights, is how this we-intention is connected with what the agent who has it does. Let us turn our attention back to [1a]. Our analysis of this is incomplete. When we turn to an account of the conditions under which this intention is carried out, we must read more into its content, as we have noted. I am going to represent this slightly differently from the way Searle does, but I think equivalently. For my intention to sing the national anthem to be carried out, it must be the case that as a result of so-intending: I do something primitively which brings it about directly that a singing of the national anthem occurs in accordance with a plan I have or will develop (perhaps in the course of doing it). I add the requirement that it come about as planned to handle difficulties with deviant causal chains. I will suppose that in the case of a prior intention this includes that the action be brought about by the intention evolving into an intention-in-action which guides the action (as needed) while it occurs and which is *de re* about the action. Representing this, we have [1b].

[1b] I intend (as a result of my so intending: $(\exists e)(\text{singing}(e) \text{ and } (\exists f)(\text{primitive-agent}(f, I) \text{ and cause}_{(\text{directly})}(f, e) \text{ in accordance with a plan I have or will develop) and } (\text{no } y: y \neq \text{me})(\text{agent}_{(\text{directly})}(e, y)) \text{ and of}(e, \text{ the national anthem}))$)³

As an aside, it is worth noting that if we allow that ‘intend’ introduces a quantifier over states, as event verbs introduce a quantifier over events, we can make use of that to represent more explicitly what is meant by ‘as a result of my so intending’. The idea would be roughly:

$(\exists s)(\text{intention}(s) \ \& \ \text{subject}(s, I) \ \& \ \text{content}(s, s \text{ brings it about directly that } (\exists e)(\text{singing}(e) \ \& \ (\exists f)(\text{primitive-agent}(f, I) \ \& \ \text{cause}_{(\text{directly})}(f, e) \ \& \ \text{in accordance with a plan I have or will develop) and } (\text{no } y: y \neq \text{me})(\text{agent}_{(\text{directly})}(e, y)) \ \& \ \text{of}(e, \text{ the national anthem}))))$

In the interests of keeping the length of the formulations from getting out of hand, I will not use this elaboration below.

Now let us turn to [4b]. A difficulty in the case of [4b] is that the complement does not have a term that refers to the subject of the intention, and this signals that its form cannot be read straight off from [1b]. To see what to do, we should then turn our attention to singular intentions with complement clauses which do not refer to the subject of the attitude. Let us then apply the satisfaction principle to [7].

[7] I intend that Sue not drink too much at the party.

For this intention to be carried out, it is not enough that Sue not drink too much at the party. I may invite her, and intend that she not drink too much, but forget about it and do nothing toward that end. It may happen that she does not drink too much, perhaps because she came with that resolve, or left early for another appointment, or was dissuaded by someone else. My intention was not carried out. For my intention that Sue not drink too much to be carried out, I have to do something that has some positive effect on her not drinking too much. I have to do something that helps to secure (in the way planned) *that* she not drink too much. I have to, for example, hide the liquor bottles when I see her glass is empty, or ask her what her New Year’s resolutions were (knowing one was not to drink), or leave in a place where she will see it a magazine with a cover story on alcoholism, expecting that these things, in the circumstances, will have the right result. Thus, in logical form, [7] should involve my doing something primitively to bring it about as a result of my so intending, in accordance with a plan I have or will develop, that Sue not drink too much at the party. We can represent this as in [7a].

[7a] I intend (as a result of my so intending: $(\exists f)(\text{primitive-agent}(f, I) \ \& \ f \ \text{brings it about in accordance with a plan I have or will develop that Sue not drink too much at the party}))$

³ A quick note on ‘cause_(directly)(f, e)’: this should not be read as ‘f is the one and only direct cause of e’ but rather as ‘f is a direct cause of e’. We need to allow that more than one agent can stand in this relation to an event, as when we lift a bench together.

Return to [4b], repeated here.

[4b] [Each $x \in us$](x intends that $[(\exists e)(\text{singing}(e) \text{ and } [\text{each } x \in us](\text{agent}_\chi(e, x)) \text{ and } [\text{only } x \in us](\text{agent}_\chi(e, x)) \text{ and of}(e, \text{the national anthem}))])$)

How should we apply the pattern used in [7a] for [7] to [4b]? A first thought is that we should instantiate [4b] to a member of the group—let us say me—and then put the that-clause in the place of ‘that Sue not drink too much at the party’ in [7b]. This gives [4c].

[4c] I intend (as a result of my so intending: $(\exists f)(\text{primitive-agent}(f, I) \text{ and } f \text{ brings it about in accordance with a plan I have or will develop that } [(\exists e)(\text{singing}(e) \text{ and } [\text{each } x \in us](\text{agent}_\chi(e, x)) \text{ and } [\text{only } x \in us](\text{agent}_\chi(e, x)) \text{ and of}(e, \text{the national anthem}))])$)

Unfortunately, this does not work. Seeing why, and where we went wrong, will show, I think, both what the key is to understanding the structure of we-intentions, and why it has been so elusive.

The difficulty is not that I do not intend to do my part. For I do intend that we do something, and that requires me to do my part in it, since I am part of the group. The difficulty rather is that my planning that we do something and your planning that we do something may not amount to our planning to do something. One case of this is that of Jones and Smith where each intends to take advantage, as he thinks, of the other’s inattention to do something that will result in their mixing a hollandaise sauce. Another example is provided by two painters who think they differ on which is the best brand of paint to use, Sherman-Williams or Benjamin Moore, in painting a house. They each think the other prefers Benjamin Moore and will insist on using only Benjamin Moore for the project. But each intends that they should paint using only Sherman-Williams. To this end, each plans a deception. Each brings to the work site a Benjamin Moore can filled with Sherman-Williams. Each also brings in the trunk of his car another can of Sherman-Williams. Each then surreptitiously empties the other’s Benjamin Moore can and replaces the paint in it with Sherman-Williams paint from the can he had hid in his trunk. They then proceed to paint using Sherman-Williams, each thinking he is fooling the other into doing so with him. They each have an intention which fits the pattern in [4c] with appropriate adjustments for content. But they did not have we-intentions with respect to painting with Sherman-Williams, and *they* did not paint using Sherman-Williams *intentionally*.⁴ Intuitively, what is missing is that neither intended that they should be doing something in accordance with a shared plan, but rather each intended something that presupposed that they did not intend to act in accordance with a shared plan.

⁴ I am not saying that they did not paint the house together intentionally, but that they did not paint it together using Sherman-Williams intentionally.

Where did we go wrong? Return to [4], which should pattern after [1], rather than [7], because it, like [1], has an infinitival complement. Thus, we should use the pattern in [1b], which gives us [4d].

[1] I intend to sing the national anthem

[1b] I intend (as a result of my so intending: $(\exists e)(\text{singing}(e))$ and $(\exists f)(\text{primitive-agent}(f, I))$ and $\text{cause}_{(\text{directly})}(f, e)$ in accordance with a plan I have or will develop) and $(\text{no } y: y \neq me)(\text{agent}_{(\text{directly})}(e, y))$ and of(e , the national anthem)))

[4] We intend to sing the national anthem

[4d] We intend (as a result of our so intending: $(\exists e)(\text{singing}(e))$ and $(\exists f)(\text{primitive-agent}(f, we))$ and $\text{cause}_{(\text{directly})}(f, e)$ in accordance with a plan we have or will develop) and $(\text{no } Y: Y \neq us)(\text{agent}_{(\text{directly})}(e, Y))$ and of(e , the national anthem)))

Here, of course, we want to read the content as involving a collective action, and so representing the first ‘we’ and the ‘us’ as quantifying over members of the group and taking narrow scope with respect to the quantifier over events. Likewise ‘our so intending’ should be read as ‘each of us so intending’. In the case of ‘a plan we have or will develop’, we must read ‘we’ as it appears here as independent of the first, as distributive, but as taking narrow scope with respect to ‘a plan’, since the key idea is that there is a shared plan. The subject term of course is interpreted as a quantifier over members of the group. When we put these ideas together, we get:

[4e] [Each $w \in us$](w intends (as a result of each of us so intending: $(\exists e)(\text{singing}(e))$ and [each $x \in us$]($\exists f$)(primitive-agent(f, x) and $\exists p$: plan(p) and [each $y \in us$](y has or will develop p))(cause_(directly)(f, e) in accordance with p)) and [only $z \in us$](agent_(directly)(e, z)) and of(e , the national anthem))))

This says that each of us intends *that* as a result of each of us so intending there is an event of singing of the national anthem of which each of us is a direct agent and there is a plan (the same one) which each of us has or will develop to do this and that no one who is not one of us is a direct agent of this singing.

We still though have to connect what is expressed by [4e] with actions of the individual. We have in [4e] an intention in which the complement would be expressed in English by a *that*-clause without the subject of it being identical to the subject of the superordinate clause. So we must apply at this point the pattern of analysis appropriate for [7], ‘I intend that Sue not drink too much at the party’; that pattern is given in [7a]. Instantiate then [4e] to a member of the group—let it be me—and we have finally [4f].

[7a] I intend (as a result of my so intending: $(\exists f)(\text{primitive-agent}(f, I))$ and f brings it about in accordance with a plan I have or will develop that Sue not drink too much at the party))

[4f] I intend (as a result of my so intending: $(\exists f)(\text{primitive-agent}(f, I))$ and f brings it about in accordance with a plan I have or will develop that as a result of each

of us so intending: $(\exists e)(\text{singing}(e) \text{ and } [\text{each } x \in \text{us}](\exists f)(\text{primitive-agent}(f, x) \text{ and } [\exists p: \text{plan}(p) \text{ and } [\text{each } y \in \text{us}](y \text{ has or will develop } p)] (\text{cause}_{(\text{directly})}(f, e) \text{ in accordance with } p)) \text{ and } [\text{only } z \in \text{us}](\text{agent}_{(\text{directly})}(e, z)) \text{ and of}(e, \text{ the national anthem}))))))$

This is the final analysis of a we-intention.⁵ A couple of quick remarks may help to ward off misunderstandings. First, to have such an intention does not require that one think that one has the power to determine what other people intend (see Baier 1997; Stoutland 1997; Velleman 1997). It is no more required in this case than it is in the case of my intending that Sue not drink too much at the party. What is required is that one think that what one can do will, given the rest of how one believes things are or will develop, which may include of course what one takes the dispositions of others to be, have a reasonable chance, weighted by the value of the outcome, of achieving its goal (see Bratman 1999 for discussion). Second, one might worry that the requirement of a shared plan is too strong, for surely when people do things together they often do so without having in full detail exactly the same plan in mind that they will follow. This is true enough, but the requirement is not that everyone share all the same plans in all their detail, but that there is a plan that they share, and at a suitable level of abstraction, e.g., reform campaign financing, this seems easy to meet.

3. CONCLUSION

By way of conclusion, let us ask how this account fares with our examples and with our constraints.

Does the analysis handle the cases that gave us trouble before? Yes. It handles the case of the business school graduates, as noted above, because none of them

⁵ It is worth noting that this account of we-intentions includes no mutual belief requirement. That requirement attempts to secure a shared goal, but as Searle's example discussed above involving the business school graduates who have learned the theory of the hidden hand, a mutual belief requirement falls short of what is required. The shared goal is secured through the intentions of the participants in group intentional action. Further, it seems clear that mutual belief is not just too weak in one respect, but too strong in another. Suppose A and B are partners in crime who believe in honor among thieves, and they agree that if they are captured and interrogated by the police, neither will give up the other by confessing but rather stick to their mutually agreed upon cover story. Suppose they are captured and presented with a classic prisoner's dilemma. Each is told that if neither confesses, each will receive some jail time but not commensurate with the crime of which they are accused. If both confess, then each will receive a sentence at the lower end of the range appropriate for the crime. However, if one confesses and the other doesn't, then he will receive immunity while the other will receive a much longer jail sentence, while if he does not and the other does, then he will receive the longer jail sentence and the other will receive immunity. Suppose A and B under the stress of interrogation each come to believe that the other will most likely confess, given the structure of the payoffs. But suppose that each is nonetheless determined to observe the pact he made, hoping the other will as well, but expecting otherwise, and willing to accept the consequences. They both act in accordance with the plan which they had agreed on, and so carry it out together. Intuitively they have acted together intentionally to frustrate the police investigation by sticking to their cover story. But a mutual belief requirement to the effect that the others are doing their parts would rule this out, incorrectly, and so is too strong a requirement on having we-intentions.

intend anything with regard to their all doing something together—this is ruled out by the description. In the case of Jones and Smith taking advantage, as each thinks, of the other's inattention to mix hollandaise sauce, neither intends that they do what they do by way of a shared plan, for each thinks the other unaware of the other's plan that they mix hollandaise sauce. Similarly, in the case of the two painters, neither intends that they should be painting with Sherman-Williams by way of their sharing a plan to do so, since each thinks the other plans (and thinks they plan) something incompatible with that. What about the case of Jules, Jim, and Jill, who each break out of prison thinking they are blowing up the Golden Gate Bridge with other saboteurs and do so together? What they lack are intentions that are about the group that actually blows up the Golden Gate Bridge. None of them has any way of thinking of them as a group. Their dreams of being approached by the angel of God provide no route in thought that connects them with *any* other people engaged in a task with them. And so their intentions are not about any group and so not about *them*. Contrast this case with one in which someone approaches the three of them independently and gives each a coordinated role in a plan to blow up the Golden Gate Bridge, indicating that there are others who will participate but who must remain unknown. In this case, they do have a way of thinking about the group of which they are a part, namely, as the individuals the organizer has recruited for the job. In this latter case, if they blow up the bridge according to plan, they do it intentionally.

Does this analysis meet our three constraints? Yes. It is clearly consistent with the fact that society consists of nothing but individuals, and requires no psychological states except those of individuals. In the case of the requirement that it be consistent with an individual's intentions being independent of whether he is getting things right, the account is certainly consistent with no one else having the plans or intentions that he intends that they have. It will be consistent with no one else existing provided that we take the appearances of the plural referring terms in the representation of the content of the intentions to involve in the individual's thought some sort of mode of presentation of the members of the relevant group that makes no *de re* reference to anything external to the individual. And it is clearly consistent with the third constraint, for the distinction between I-intentions and we-intentions has been drawn by way of a difference in the contents of the relevant sorts of intentions.

Finally, if it is successful, it also shows that we-intentions can be understood without invoking any concepts which are not already available in our understanding of I-intentions, for all of the concepts expressed by terms used in the analysis of we-intentions are used in the analysis of I-intentions. Yet given how we-intentions and I-intentions are introduced, it does this without reducing we-intentions to I-intentions, or to I-intentions together with mutual beliefs.

REFERENCES

- Baier, A. (1997). *The Commons of the Mind*. Chicago: Open Court.
- Bratman, M. (1992). Shared cooperative activity. *The Philosophical Review*, 101(2): 327–341.
- Bratman, M. (1999). I intend that we *J*. In his *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge: Cambridge University Press.
- Copp, D. (2006). On the agency of certain collective entities: an argument from ‘normative autonomy’. *Midwest Studies in Philosophy*, 30: 194–221.
- Davidson, D. (2001). The logical form of action sentences. In his *Essays on Actions and Events*, Second Edition. Oxford: Clarendon Press. Originally published in N. Rescher (ed.), *The Logic of Decision and Action*. Pittsburgh: University of Pittsburgh Press, 1967.
- Gilbert, M. (2000). *Sociality and Responsibility: New Essays in Plural Subject Theory*. Lanham, MD: Rowman & Littlefield.
- Ludwig, K. (forthcoming). Collective intentional behavior from the standpoint of semantics. *Noûs*.
- Lepore, E., and Ludwig, K. (2003). Outline of a truth conditional semantics for tense. In Q. Smith and A. Jokic (eds.), *Tense, Time and Reference*. Cambridge, MA: MIT Press.
- Pietroski, P. M. (2000). *Causing Actions*. Oxford: Oxford University Press.
- Searle, J. R. (1983). *Intentionality*. Cambridge: Cambridge University Press.
- Searle, J. R. (1990). Collective intentions and actions. In P. R. Cohen, J. Morgan, and M. E. Pollack (eds.), *Intentions in Communication*. Cambridge, MA.: MIT Press. Reprinted in Searle 2002.
- Searle, J. R. (1995). *The Construction of Social Reality*. New York: Free Press.
- Searle, J. R. (1998). *Mind, Language, and Society: Philosophy in the Real World*. New York: Basic Books.
- Searle, J. R. (2002). *Consciousness and Language*. Cambridge: Cambridge University Press.
- Stoutland, F. (1997). Why are philosophers of action so anti-social? In L. Alanen, S. Heinamaa, and T. Wallgreen (eds.), *Commonality and Particularity in Ethics*. New York: St. Martin’s Press.
- Tollefsen, D. (2002). Collective intentionality and the social sciences. *Philosophy of the Social Sciences*, 32(1): 25–50.
- Velleman, D. (1997). How to share an intention. *Philosophy and Phenomenological Research*, 57(1): 29–49.