

Philosophy of the Social Sciences

<http://pos.sagepub.com>

Joint Actions and Group Agents

Philip Pettit and David Schweikard
Philosophy of the Social Sciences 2006; 36; 18
DOI: 10.1177/0048393105284169

The online version of this article can be found at:
<http://pos.sagepub.com/cgi/content/abstract/36/1/18>

Published by:



<http://www.sagepublications.com>

Additional services and information for *Philosophy of the Social Sciences* can be found at:

Email Alerts: <http://pos.sagepub.com/cgi/alerts>

Subscriptions: <http://pos.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations <http://pos.sagepub.com/cgi/content/refs/36/1/18>

Joint Actions and Group Agents

Philip Pettit

Princeton University, Princeton, New Jersey

David Schweikard

University of Cologne, Germany

Joint action and group agency have emerged as focuses of attention in recent social theory and philosophy but they have rarely been connected with one another. The argument of this article is that whereas joint action involves people acting together to achieve any sort of result, group agency requires them to act together for the achievement of one result in particular: the construction of a centre of attitude and agency that satisfies the usual constraints of consistency and rationality in adequate measure. The main discovery in the recent theory of group agency is that this result is not easily achieved; no regular voting procedure will ensure, for example, that a group of individually consistent agents will display consistency in group judgments.

Keywords: *groups; group agents; collective agents; joint action; joint intention*

Joint action and group agency have emerged as focuses of attention in recent social theory and philosophy. But they have not been connected with one another in most discussions, and the relationship between the two may not be properly appreciated. This article is an attempt to provide an overview of the more salient issues covered in the literature on joint action and to relate them to the question with which the theory of group agency engages.

1. The Problems of Joint Action

Every discussion in this area has to begin from the fairly banal observation that without any suggestion of joint action or group agency, people routinely

Received 5 October 2005

Authors' Note: Our thanks for comments received from the guest editors, and from Michael Bratman and Christian List.

give rise to collective outcomes. People each cheer at a football match and combine to raise the roof. People each pursue bargains on an open market and combine to drive prices to the competitive level. People are each careless about littering and combine to make the neighborhood a mess. The examples seem.

In all of these cases, people act as independent agents, without anyone intending—certainly without anyone necessarily intending—that the effect they combine to produce should materialize. They may act without any thought to what others are doing, as if others were part of the parametric environment. Or they may act in strategic awareness of how their actions may interact with one another (Elster 1983). Thus the litterers may reason in the fashion of free-riders that if others are careful, their own efforts will not be needed; that if others are careless, their own efforts will not do any good; and that their own choice will not affect how others behave. But whether the members of the plurality are strategic or nonstrategic in their individual choices, the important point is that they can and often do produce a collective effect without any one's intending that they do so.

Without having such an intention, of course, some individuals may foresee the effect that the plurality will combine to produce, and others may not. And among those who foresee it, some may even want it to occur, as is plausible in the case of the cheering crowd or the competitive market. But no matter how widespread this pattern of belief and desire, it is in no way responsible for the appearance of the aggregate effect. It is socially epiphenomenal.

Joint action requires different people to produce an effect together, as in the cases just considered. But it also requires more. Not only is there an effect that people own in common as something they combine to realize. There is also an action that people own in common as something they combine to perform. People will make their particular contributions to that joint action and will be separately responsible—causally responsible—for those individual efforts. But, so the idea goes, they will also be responsible together for an action in which they each participate. That is why it is appropriate to speak of their acting jointly, and not just of their producing a joint effect.

Examples of joint action are not hard to find. You and I lift the table together, exchange ideas, or go for a pleasant stroll. A number of us combine to write a letter of protest, or to sing together in a choir, or to play for victory as a team, or to coordinate efforts in order to help a swimmer in difficulty. In none of these cases is it enough to say that we each act independently, where the sum of those individual actions involves the realization of a certain effect. In each case there is something that we together do; apart from our individual actions, there is an action that we together perform. In particular, there is an action that we together perform without any one of us being manipulated or

coerced by others; it is cases of such unforced cooperation that will concern us here.

Action is always performed with a certain intention, whether it be an intention formed in advance or an intention that materializes with the behavior; otherwise it need not amount to anything more than reflexive behavior. And that is why there are problems of joint action. There are at least three problems, or families of problems, that naturally arise.

The *I-to-we issue*: can I as a separate agent be rationally moved to think in we-terms and act as the member of a plurality? Or is the shift to we-thinking essentially subrational?

The *we-as-acting issue*: is the primary intention in joint action an intention that we do something together, acting as one? Or is it an intention on the part of each to do his or her bit?

The *we-as-intending issue*: is the primary intention a single state of ours, intending as one? Or is each of us moved only by a separate, individual intention?

The theory of joint action tries to resolve issues of this kind, removing the sense of paradox that may surround the idea that people can come to intend and enact things together. A large literature has grown up around these questions, and we cannot hope to review it fully.¹ In the following sections we describe a brisk path across the territory and indicate, usually in footnotes, where others stand on the questions addressed. We then go on to connect our favored view of joint action with the theory of group agency.

The article is in seven sections. In section 2, we sketch out an analysis of joint action that we find attractive, providing an initial, minimal defense of the claims it makes. In the three sections after that, we look at how each of the three problems of joint action can be resolved under that analysis, displaying the further attractions of the analysis in the process. In the sixth section, we try to show how the issues in the theory of group agency connect with issues already discussed but take us at the same time into essentially new territory. And then, in a brief, concluding section, we situate these issues within the broad domain of social ontology.

The path we take may be contested at any of a number of points, but it is especially novel, we believe, in connecting the theory of joint action with the

1. The main contributions to the debate on joint action are (without claiming completeness) the following: Brooks (1981), Tuomela (1984, 1989a, 1989b, 1991a, 1991b, 1995, 2000, 2002), Tuomela and Miller (1988), Tuomela and Tuomela (2003), Bratman (1999, chs. 5-8), Gilbert (1989, 1996, 2000), Searle (1990), Cohen and Levesque (1991), and Miller (1992, 1995, 2001). For critiques and extensions, see especially Baier (1997), Stoutland (1997), Velleman (1997), Meijers (1994, 2003), and the collections by Holmström-Hintikka and Tuomela (1997); Lagerspetz, Ikaheimo, and Kotkavirta (2001); Meggle (2002); Sintonen, Ylikoski, and Miller (2003); Koepsell and Moss (2003); and Schmitt (2003).

theory of group agency. We hope that our article may stimulate some of the many people working on joint action to take an interest in this new, rapidly evolving area of study.

2. An Analysis of Joint Action

If people perform a joint action intentionally, then they must each be focused on a common target: the behavior that they display together, in which each makes his or her contribution to the overall performance. Call this single, multiply supported behavior the joint performance. Each must combine with others to enact that performance: to take some time together on the stroll, to get the letter of protest finished, to let the choral sound ring loud and clear.²

Not only must everyone be focused on the joint performance, but also each must act with intention in enacting that performance. And this is where the air of paradox blows strongly. Each of us may intend to do something that is within his or her control, as when I scratch my shoulder or wave at a friend. But how can I or you or anyone else intend to do something that depends on all of us if it is to be accomplished? Does the occurrence of a joint performance require that the performance in toto be intended by one and only one acting subject?

It is true that I cannot intend *to* X, where X-ing is a joint performance. But I may still be able to intend *that* we X together. Intending-that is different in this respect from intending-to.³ I can intend that my child go to university, or I can intend that a student do some background reading, where the subject of the action is someone else. And equally I can intend that we, a plurality of individuals, do something of which I am not the author, or not the sole author.⁴

A first requirement of joint action, so we propose, is that we each have the appropriate intention-that. Each of us in the plurality intends that we together enact the relevant performance. Or equivalently, to stretch ordinary usage, each of us intends us to enact the performance. If this condition were not fulfilled, then how could one speak of our acting intentionally together? Those

2. Performances of the sort envisaged here are behavioral patterns. See section 3 for a discussion of whether they are represented in joint agents as the products of intention and whether that gives rise to a circularity.

3. This difference between intending-to and intending-that is reflected in Bratman's account (Bratman 1999, ch. 8).

4. The assumption that the ? "we" appears only in the content of the intention, not in the subject place—"It is I who intend that we . . ."—is defended in section 5.

of us who lacked that intention would be cut out of the action. And if all of us lacked it, then there would be no action from which to cut anyone out.

It makes sense to say that I intend that something happen only if two conditions are fulfilled. First, I must in some sense want it to happen. And second, I must be in a position to do something about it. No one would say that I intend that something happen, when I am known overall to be averse to the prospect. And no one would say that I intend something to happen, when I can have little or no effect on whether it happens as a matter of fact. There is no problem with seeing how a number of individuals should each want that they together enact a certain performance. But what is it that they can each do in order to make this happen? What is it that makes room for talk of intention rather than just talk of want?

The answer is that in the normal case, each of us will have a certain part to contribute to anything that counts as a joint performance. And so each of us can intend that the joint performance happen, so far as each of us is disposed to do all that we can to make it happen; each of us is disposed to do our bit, recognizing it as our bit, as our contribution to the joint performance.⁵

Doing one's bit is not the only contribution that a person might make to a joint performance. We can imagine scenarios in which one or another member of a plurality is in a position to exploit, manipulate, or coerce others and so is specially empowered to intend that they together enact a certain performance. But this would not be sufficient for joint action in the unforced sense that interests us here. There can be nothing underhanded or overbearing involved in unforced joint action; people must voluntarily contribute whatever is required for the desired performance.

So far we have identified two conditions of joint action, then. First, each member of the plurality intends that they together enact the relevant performance. And, second, they intend this so far as each intends to do his or her bit on the usual, unforced basis. But with these conditions in place, it is easy to see that more needs to be added.

Why might each of us intend to do his or her bit in bringing off some performance together? It just might be that we each thought that others were zombies who would automatically, as if under hypnotism, do what was required of them. It just might be, in other words, that we each thought of ourselves as the only properly intentional agent involved. But that would not suffice for anything deserving in an ordinary sense of being called joint action. Each of us would think of acting jointly with others, only in the sense in

5. Raimo Tuomela distinguishes between the cases in which an agent intends to act but does not take this action to be a contribution to a joint performance and cases in which, as he says, he "intends to perform his part of X as his part of X" (Tuomela forthcoming). In the following, we will not mention this specification but take it to be entailed in what is said.

which we might think of acting jointly with the wind in steering a yacht, or jointly with the fire in toasting a marshmallow.

We need to add to the conditions given a stipulation to the effect that the members of the plurality who act together each intend to do their bit because of believing that others intend to do their bit in the preconceived performance. The members of the plurality must see themselves, in other words, as members who in this respect are on a par with others; they are involved on an equal footing in the enterprise of joint action.

This extra stipulation involves two separate clauses: that each believes that others intend to do their bit, and that this is the reason why each intends to do his or her bit. Thus we are pointed towards four elements for an analysis. A number of people in a plurality perform a joint action in enacting a certain performance together only if

1. they each intend that they enact the performance;
2. they each intend to do their bit in this performance;
3. they each believe that others intend to do their bit; and
4. they each intend to do their bit because of believing this.⁶

Do we need to add a further clause? Yes, for familiar reasons. Suppose you and I each intend that we together paint the house. Suppose we each intend to do our bit, with you taking the front and me taking the back. Suppose we each believe that the other intends to do his or her bit, and suppose, finally, that this belief helps to explain why we each intend to do our own bit. For all that these clauses stipulate, it just might be that you do not believe that I believe that you intend to do your bit, let alone that I intend to do my bit because of believing this; for example, you might regard me, wrongly, as someone who takes you to be acting like a zombie, as if under hypnotic suggestion. And in such a case, there would be a reasonable ground for denying that our painting together should count as an unforced, joint action.

In order to block this possibility, we need to introduce a fifth clause that serves to silence it and any other possibility of the same kind. The clause stipulates that everything amongst the parties is above board. Take any question that might arise as to whether I am really aware of what is happening, or really aware of our each being aware of what is happening, or really aware of our each being aware of our each being aware of what is happening, and so on. The clause stipulates that we are each disposed to give an affirmative response to such a question; in that sense, we are aware of what is happening,

6. On related issues to do with conditional and unconditional intentions, see Tuomela and Miller (1988) and Velleman (1997).

aware of our each being aware of this, and so on in the usual hierarchy of common belief (Lewis 1969). Thus we should add as a further clause that

5. they each believe in common that the other clauses hold.⁷

We think that the five clauses given here are not just individually necessary but jointly sufficient in order for the enactment of a joint performance—for the behavior involved in a joint performance—to count as a properly joint action. We shall proceed on the assumption that that is so, though nothing much turns on that claim. We now go on to look in turn at the three problems of joint action that were distinguished at the end of the last section. We can articulate and address those problems with greater clarity, in the light of the analysis sketched.

3. The I-to-We Issue

The first of our problems or families of problems bears on the issue of whether I can rationally shift from thinking in my own name to thinking in the name of a plurality. In normal action, I am the unit of agency in play when I ask about what should be done; I am asking about what I should do. In joint action, the unit of agency becomes the plurality, as I ask after what we should do and draw conclusions that may not require the same action that I-reasoning would have supported. On the face of it, there is something mysterious in the alleged shift from the I-mode to the we-mode, and this is the shift interrogated in the I-to-we question.⁸

It is standard in economics to assume that the theory of rational agency is exhausted by investigations in decision theory and game theory. Decision theory focuses on what it is rational of an agent to do in decisions generally, given his or her desires and credences, and game theory looks at what is rational for different parties in a variety of interdependent decisions. Both sorts of theory take the individual to be the only unit of agency involved (Hurley 1989). Each individual is an agent who is representable by means of a utility function and a credence function, and the rational action for such agents is always to maximize expected utility, with their credences determining the expectation; rational agents, in more common language, always satisfy their desires according to their beliefs.

7. The whole analysis of the requirements of joint action is probably closest to Bratman's; cf. Bratman (1999, chs. 5-8).

8. For an analysis of the notions of I-mode and we-mode, see Tuomela (2003).

This theory allows for the possibility that agents may have quite altruistic desires and that this may often lead them to behave for the benefit of others and at some cost to themselves. But it does not allow for the possibility that agents might shift the center of control away from themselves to this or that plurality of agents. Hence, the problem of whether I can rationally shift from I-thinking to we-thinking has had most salience among writers from an economics background (Sugden 1993, 2000, 2003; Bacharach 1998).

We do not think that the I-to-we problem goes very deep. There are really three aspects to the problem, none of them very troubling. The first is the question of how there can be an option open to individuals of giving up on thinking as a singleton on some issue in favor of joining with others in thinking as a plurality. The second is the question as to how individuals, when they do this, can identify the joint performances possible, each with its own allocation of parts to individual agents. And the third is the question of whether, assuming the first two issues are resolved, people can rationally opt for performing in a plurality with others.

The first two questions raise issues in empirical psychology. Do people routinely identify options of acting jointly with others? And do they routinely recognize the parts that they are required to play under this or that different option? Behavioral economics suggests that they do, revealing a readiness on the part of experimental participants to go for 'cooperative' solutions to various predicaments. And in doing this, it merely provides backing for common sense. Clearly we do often see joint action as an option. And clearly we are very good at recognizing what this or that joint performance requires of us; think of the versatility of footballers as they concoct a joint move, or of jazz players as they improvise on a standard piece.⁹

We should not be surprised at human adaptability and versatility on these fronts. We are a social species, and it is hardly a surprise that our natural and perhaps cultural history should have equipped us to be highly sensitive to possibilities and proposals of joint action. It is as if we go around, advertising to others conditional intentions of the "I will, if you will" kind; that we are disposed to assume the presence of various such intentions in others; and that we routinely respond in the "I will" fashion, thereby triggering cooperation (cf. Velleman 1997). We do this even when we do something as banal as

9. The jazz example might not seem to be in line with the other cases mentioned, since it might be held that it is an individual improviser who does the job. But improvisation is more complex than this: on one hand, it requires a multifaceted knowledge of various rhythmic, melodic, and stylistic structures on the part of the improviser. On the other hand, it requires the accompanying musicians to listen and decide what sort of background fits the spontaneously created patterns, or the improviser might even "ask for" a certain type of backing by "alluding" to genre-specific phrases.

reaching out to shake hands with another, or asking another a question with a view to exchanging information.

But, to turn to the third aspect of the I-to-we question, can it really be rational for me to do this? It cannot, we would say, under a common but crude picture of what individual rationality involves. If being rational means maintaining the first-person singular focus, and asking the question with every choice made—for example, with every initiative taken jointly with others—as to whether this answers satisfactorily to one's own desires, then joint action probably is inconsistent with rationality. At least, that will be so under a plausible empirical assumption. This is that joint action can be sustained across the complexities of circumstances that confront pluralities only if the participants are willing to put aside their personal bookkeeping and think as members of the collectivity—in particular, willing to think about which is the best way for the plurality to go, now in this context, now in that, without constantly checking back with their personal interests.

We think that there is nothing rationally problematic about joint action, however, once it is recognized that in many contexts the rational way for individuals to conduct themselves may be to put aside constant, case-by-case calculation and to go on automatic pilot. Assume that there are benefits to be won by individuals as a result of acting jointly with others: benefits in what they can achieve jointly together, and benefits of an attitudinal and related kind in the links that joint action can forge. Assume further that those benefits cannot be obtained unless individuals are willing, within temporal and other limits, to relinquish the practice of monitoring their every move for whether or not it makes rational sense. In that case, the rational thing for individuals to do may be to form and maintain various joint actions without constant calculation about the returns they personally enjoy.

Rationality will still require that they are ready to look at the personal returns whenever the “red lights” go on: whenever there are signs, as there will often be, that the cooperation may be damaging to their own interests. But it will not support constant monitoring. In individuals who cooperate with one another over any stretch of time, without calculating the returns from joint action, rational interest will not have an *active* presence or influence. But it will have a *virtual* presence and it will exercise a *virtual* form of control (Pettit 1995, 2000). The agents will be ready to endorse the behavior required of them only so long as the red lights do not go on, and, assuming that the red lights are a reliable mechanism, that means that the behavior maintained by cooperators will generally be in their rational interest. Rational interest will not have been abandoned under this dispensation; it will have been recruited to a different role.

4. The We-as-Acting Issue

The second set of problems we have to confront surrounds the issue as to whether the primary intention in joint action is that we together do something. The analysis sketched suggests that the primary intention does indeed bear on the plurality in this way. According to that account, I intend like everyone else that we do so and so, and only come to intend to do my allotted bit in the joint performance as a result of that intention. Primacy is given to the intention over the we-action rather than the intention over the I-action. The challenge now is to show that this line is defensible.¹⁰

The challenge can be sharpened by considering the following question. Do we each intend to do our bit in a certain performance, because we each intend that we together do so and so? Or is it the other way around? Is it the case that we each intend that we together do so and so, because we each intend to do our bit in that performance? Our analysis suggests that the intention over the we-action is primary, but it might seem that the other would be a less problematic claim.

The other claim would be that in virtue of our each intending to do our bit in a certain performance—in virtue of our believing that others will act in a certain way and of our intending to act in a complementary manner—we can each be said to intend that we do so and so together; or that this is the case, at any rate, so far as the pattern is a matter of common belief. Thus the only intending that materializes in any literal sense would be regular, I-directed intending. Talk of our each intending that we do something would be a harmless way of speaking, since it would summarize the truth about the aligned I-directed intendings. But it would not carry any independent weight and would not point us to anything irreducibly plural in character.

This claim is false, we think, in familiar cases of joint action. By our earlier comments, to intend that we do so and so will require two things: one, wanting that we do so and so, and, two, being able to do something to help make it the case that we do so and so. In the cases we envisage, from lifting a table together, to going for a stroll, to singing in choral harmony, the reason why we will each intend to do our bit in that performance is that we will each want that we together enact the performance. Thus the intention that we do so and so, considered just as a desire-like state, will be primary. It will help to explain why we each intend to do our bit, not the other way around. We will

10. This issue has been dealt with in the literature, yet in divergent ways. It figures prominently in Raimo Tuomela's and Kaarlo Miller's analysis of we-intentions (Tuomela and Miller 1988); in John Searle's thesis that collective intention is a primitive phenomenon and the corresponding individual intentions are "derivative" (Searle 1990, 403); and in Michael Bratman's works on these matters (Bratman 1999).

each intend and want to sing our part because of each wanting that we together sing in harmony. We will not want to sing in harmony, not at least in the ordinary, nonnarcissistic case, because of each intending and wanting to sing our part.

But why describe the desire that we do such and such as a case of intention? And can we continue to think of it, qua intention, as the moving state, not the moved?

My desire over the we-action will constitute or give rise to an intention so far as I am able to do something about bringing about the we-action desired and am led by the desire to do precisely that.¹¹ It will count as an intention, as we may say, when it leads me to do my bit and, of course, to do it with the I-directed intention required for controlling that contribution. It will count as an intention in the way any desire for an end counts as an intention once it leads me to adopt a means to that end, and to form a secondary intention controlling the precise means I adopt.

Does this imply, then, that even qua intention over the we-action, the state of wanting that we do so and so is primary, explaining why I form the intention to do my bit? We think it does. That desire counts as an intention only so far as it produces suitable action, with a subsidiary intention appropriate to that action. But it remains the state that explains the appearance of that action and that intention. Consider a parallel. All will agree that the measles virus is what produces the measles syndrome: the characteristic spots and fever. The causal and explanatory role of that virus is not in any way diminished by the fact that it is only in virtue of producing the syndrome that it is called the measles virus. In the same way, the causal and explanatory role of my intention that we do so and so is not in any way diminished by the fact that it is only because of its effect in getting me to do my bit, and intend to do my bit, that we can speak of an intention that we do so and so, not just a desire that we do so and so.

If all this is right, then our analysis takes the correct line in putting up front in the characterization of joint action an intention on the part of each of us in a plurality that we do something together; that we enact a joint performance, as it was put earlier. But some may still balk, on the grounds that the line taken involves a troubling circularity. Is not the notion of a joint performance already a notion of joint action, they will ask? And does not that mean that we are effectively invoking the notion of joint action in the analysis of joint action? The suggestion is that we might have done better, after all, to treat the first clause in the analysis given as being implied by the second, rather than

11. Although this way of speaking may sound partisan, we hope that it can be reconciled with different theories of intention, including for example that which is defended in Bratman (1987).

having an autonomous status and providing the explanation of why the second holds.

We reject the suggestion, because we do not think that there is a troublesome circularity in play. The joint performance intended can be conceptualized just as a pattern of behavior in which our different efforts combine to effect a certain result. Thus there need be no circle involved in intending that that pattern of behavior be realized in which we lift a table or sing together, or indeed play chess or dance the tango. The content of our intention can be conceptualized at the behavioral level, without recourse to the notion of joint action that we are trying to analyze (see Bratman 1999, 114; and for a discussion, Petersson 2005).

Each one of us involved in a joint action will have the concept of an intention, of course, given that the analysis requires us to have beliefs about intentions. And so each of us will be in a position to know, and know a priori, that the performance intended in joint action, if it materializes appropriately, will have materialized as a result of that intention. But this is no problem. True, we may not be able to intend a joint performance except as something that, so we know, will be produced as the result of the intention we each have that it be realized. We may not be able to intend that we play chess or dance the tango except as something that, so we know, will come about as a result of such an intention. But this is not to say that we cannot conceive of what we intend to produce except so far as we already have access to the notion of an intentional joint action. There is no obstacle to one of us first conceiving of a joint performance independently of its being intentional, then realizing that we each have an intention that it be realized, then recognizing the truth of the other clauses in the analysis, and only at that point coming to understand fully what a joint intentional action is.

5. The We-as-Intending Issue

And so, finally, to the third family of problems that arise with joint action. The issue is whether joint action involves the plurality of agents having a single state of intending that can be expressed in the words such as “we, the group, collectively intend that we . . .” rather than “we, these individuals, severally intend that we . . .”. The analysis sketched suggests that a single plural state of intending need not be involved in joint action. And so the issue here is whether the analysis takes a defensible line.

When we act together intentionally, it might be thought that we should intend together to perform that act. But on the analysis provided, we do not figure as a single intending subject, at least not in a collective sense; it is we

severally who intend that we act together, not we in the sense in which we might constitute a subject proper. When we act together intentionally, it might be thought that a single state of intending must have taken possession of us. But on the analysis provided, there are really many particular intentions in play at the origin of the action: my intention that we do such and such, your intention that we do such and such, and so on. The analysis guarantees that we each have the same type of intention, being intent on realizing the same content. But it does not give any countenance to the intuitive idea that there ought to be a single instance or token of that type, a single state of intending, that is at the source of the act.

As the analysis fails to point us to a single collective subject that is causally responsible for the action, then, so it fails to direct us to a single state of intending. Is it inadequate for failing in these ways? We do not think so. We see no metaphysical reason why a joint intentional action has to be the product of a single agent or a single state of intending. We do not deny that it is possible for a number of agents to construct a single agent, where that agent is characterized by particular states of intending. But we believe that that possibility materializes only when group agents make an appearance; it is not associated with joint action as such. We believe that those who associate it with joint action are probably focused, without being aware of the fact, on group agency.

We will only be able to display the full significance of this position after the discussion of group agency in the next section. But the general approach should be clear. In joint action, as our analysis brings out, a number of different agents come together to perform a joint action and pursue a joint effect. In doing that, they go beyond the case where a number of different agents perform different actions and bring about a joint effect; this is the sort of scenario we mentioned at the outset. But they do not, or need not, go so far as to constitute anything deserving to be called a novel plural subject, with its own distinctive states of intention. Joint action is just that: the joint and intentional production of some effect. It need not yet involve the joint construction of a novel center of intentional attitude and action.

The argument for this point of view can be put as follows. Assume that, while it may be impossible to state a lower limit, any agent will have to display a modicum of desires and judgments. If joint action involves the appearance of a novel agent, distinct from the individuals involved in the action, then that agent will have to satisfy the constraints associated with the possession and updating of such attitudes. One important constraint is that it will have to be responsive to demands for consistency in the judgments it forms: judgments as to the ends to be sought in action, the relative order and urgency of those ends, the opportunities available for their realization, the most effec-

tive means for realizing them, and so on. But people may come together in joint action, under any plausible analysis of the phenomenon—certainly under any analysis in the literature—without giving rise to a subject that is fitted to meet such a constraint.

Consider the discursive dilemma (Pettit 2001, 2003). This predicament, which is an instance of a more general problem (List and Pettit 2002), shows that it is no easy matter for a plurality of people to generate a common, consistent body of judgments by which to navigate in action. Ask three people, A, B, and C, to generate a common body of judgments even on a simple set of issues such as whether *p*; whether if *p*, then *q*; and whether *q*. Suppose they decide to follow majority voting to resolve each issue. It is then quite possible that they will conclude as a group that *p*; that if *p*, then *q*; and that not-*q*. A and C may support *p*, with B against. B and C may support if *p*, then *q*, with A against. Thus the group will support both *p* and if *p*, then *q*. Will the group support *q*? Not necessarily. A may vote against *q*, not believing that if *p*, then *q*. B may vote against *q*, not believing that *p*. Only C is required to vote for *q*. And so the group may find itself committed to holding that *p*; that if *p*, *q*; and that not-*q*.

No plausible analysis of joint action, and none in the literature, requires that just for purposes of acting jointly people have to take precautions against the appearance of such inconsistencies. The analyses do not require, in other words, that agents take those steps that are essential for the appearance of a novel, consistency-sensitive center of intentional attitude and agency. Thus the notion of joint action as that is analyzed in the literature does not entail the appearance of a novel subject, or the appearance of the token states of intention that such a subject might instantiate. And so it can be no complaint against the analysis we have offered that it does not support an entailment of that kind.

Before leaving this discussion, it may be worth mentioning two proposals that are occasioned, so we think, by the desire to link joint action as such with a single we-agent or at least a single we-intending. They have received a good deal of prominence in the literature.¹²

John Searle (1990) rejects the idea that there is any true subject constructed in the course of joint action. So far, we are with him. But, rehearsing the intuition that a joint action cannot be the result of just individual intendings, even intendings directed to a we-action, Searle argues for a distinctive we-intending. His idea is that amongst those involved in a joint action, each head contains a state of intending that is distinctively plural: a

12. Those who are familiar with the debate on (the intentionality of) joint action will have noticed that our account has a lot in common with Michael Bratman's (1999) and Raimo Tuomela's (esp. 1991b and 1995) accounts of joint action.

we-intending. What makes this distinctively plural? Not the fact that a number of people instantiate it, since Searle allows that I may instantiate such a state in the mistaken belief that others do so too. So what, then? We see no answer in Searle's work and find his position on this issue inherently obscure.¹³

Margaret Gilbert (1989, ch. 5; 1996, ch. 6) takes a more interesting line that leads her to argue both that a plural subject appears with every joint action and that there is a single token intention at the origin of every action.¹⁴ She starts from the observation that whenever two or more people are involved in a joint action, then, intuitively, it is a matter of common acceptance or belief that any one of them may reasonably complain about the failure of another to do his or her bit and may expect an explanation or apology in return. She argues that it is in virtue of this state of affairs—this joint commitment by the parties to one another, and the resulting mutual obligations and entitlements—that we can speak of their properly acting together out of a shared intention. And she holds that this joint commitment, all on its own, means that there is a joint subject present, in particular a subject that can instantiate token intentions to act in this or that manner.

We agree with Gilbert that joint action goes, intuitively, with the sort of joint commitment that she describes. But we do not see that the joint commitment alleged does anything to support the presence of a joint subject, capable of instantiating states of intention; the argument from the discursive dilemma applies in this case too. Nor do we favor the line that she takes in making the joint commitment basic. We think that it is in virtue of their jointly acting together that people are jointly committed to one another, not that it is in virtue of their being jointly committed that they jointly act together. The satisfaction of clauses like those given in our analysis of joint action means that people rely on one another and form expectations in regard to one another, and that this mutual reliance or expectation is manifest to all. And that sort of manifest, mutual reliance or expectation will be a source of obligation under a range of different normative theories (Scanlon 1990; Pettit and Smith 2004). Thus the position we adopt on this issue is not particularly controversial, though it is not one that Gilbert herself would endorse (cf. Gilbert 2004).

6. The Theory of Group Agency

There are some group agents in the natural world that emerge under selective pressures and in such a way that we need not think of the members as

13. For a similar critique of Searle, see Zaibert (2003).

14. We take this to be Gilbert's view in "What Is It for Us to Intend?" (Gilbert 2000, ch. 2).

having any intentional attitudes in the domain of the group's attitudes, or perhaps any intentional attitudes at all. Consider, for example, the way bees in a swarm may behave like an agent (Seeley 2001). Under a recent story, half the bees in a hive move away with the old queen at the end of each season. Initially, they swarm outside the old nest, then move to a nearby tree that provides a little shelter, and then send off scouts in search of a new location for a hive. The scouts return at regular intervals, and each performs a dance that indicates the distance, direction, and quality of a candidate location for the new hive that it has found. At any time a number of such dances are being performed for the benefit of the swarm. After a couple of days, however, the dances converge and then the swarm moves off to the chosen location. With a group of bees like this, it is natural and useful to ascribe swarm-level goals, representations, and intentions: that is, to adopt the intentional stance towards the group (Dennett 1987). But there is little or no reason to think that the individual bees have corresponding attitudes.

We do not think that there are group agents of that kind formed by human beings. We see no evidence for groups in the formation of which individual intentional attitudes are irrelevant in the same way; indeed we have trouble imagining what might plausibly constitute such evidence (Pettit 1993, Chap. 3). But we think that while individual agents continue to be guided by their own intentional psychology, still they may come together to construct novel agents: agents that have a distinct intentional profile from the profiles of their members (Pettit 2001, ch. 4; Pettit 2003).¹⁵ We defend that possibility claim in this final section.

A group of individuals will constitute an agent, plausibly, if it meets conditions like the following. First, the members act jointly to set up certain common goals and to set up a procedure for identifying further goals on later occasions. Second, the members act jointly to set up a body of judgments for rationally guiding action in support of those goals, and a procedure for rationally developing those judgments further as occasion demands. And third, they act jointly to identify those who shall act on any occasion in pursuit of the goals, whether they be the group as a whole, the members of the group individually, certain designated members, or certain agents that the group hires.

Did individuals come together in the manner characterized, then they would be in a position as a group to mimic or simulate the performance of an individual agent. The group would have goals corresponding to individual desires, judgments corresponding to individual beliefs, and just as rational

15. For a different argument towards a similar conclusion, see Tollefsen (2002a, 2002b). A completely different, reductive account of corporate action is developed by Miller (2001, ch. 5; 2003).

individuals act so as to satisfy their desires according to their beliefs, so this group would be able to act rationally so as to achieve its ends according to its judgments.

There is evidence on all sides that individuals often do come together to form group agents of the sort envisaged. Coauthorships, partnerships, civic associations, commercial companies, churches, courts, and cabinets: all are examples of group agents that act on something more or less approximating the pattern described. There may be great variations in how far group procedures require the participation of individual members but the pattern remains essentially the same. Under widely accepted conventions and expectations, entities such as those mentioned have a place in social life that parallels that of individual persons. They own property and enter contracts, they are responsible for their commitments, and they are treated as agencies with which it is possible to reason, as one person may reason with another.

But we hold that not only can such group agents exist; they can also exist as agents in their own right, distinct in a significant way from the agents who are their members. A lesson of the discursive dilemma is that, if a group uses majority voting to make judgments on several interconnected issues, then it may run into inconsistencies in its resulting body of judgments. If the group seeks to maintain consistency, it will have to abandon its majority judgment on some issues and form a judgment on those issues that most of its members do not individually endorse (List and Pettit forthcoming). It will have to develop a mind of its own (Pettit 2003).

Consider a group agent consisting of A, B and C, as in the earlier example. Imagine that under the pressure of decision and action, they have to form judgments, now on whether p, now on whether q, now on whether r, and yet again on whether $p \& q \& r$. All but A might vote for p; all but B for q; all but C for r; and, consequently, none for $p \& q \& r$: each would reject it because of rejecting one conjunct. These votes would have the group holding that p, that q, and that r, but that not- $p \& q \& r$. Now in order to function as a group, the members would have to decide on rejecting at least one of those four propositions. But were they to reject p or q or r, they would establish a group view that broke on that issue with the views of a majority of its members. And were they to reject not- $p \& q \& r$ —were they to accept that $p \& q \& r$ —then they would establish a group view that none of them individually held on that issue.

The cost of forming a perfectly possible form of group agent, so this illustrates, is that under certain circumstances the group will have to adopt a body of judgments that breaks on some issue with the judgments of a majority, perhaps even a unanimity, of its members. This means that there can be group agents formed, then, that are distinct in an absolutely clear sense from their

own members. Such an agent will have to espouse an intentional profile that marks it off on some issues from the profile, not just of a minority of members, and not just of a majority of members, but even of every single individual in the group. In quite a literal sense, it will have to evolve a mind of its own.

This line of thought establishes that while joint action does not in itself necessitate the appearance of a novel subject, contrary to the views of some authors, joint action of a specific sort can do so. Let individuals come together in the pattern of joint action described here, and distinct group agents will certainly emerge. Let them act jointly so as to set up suitable ends, a suitable body of judgments for guiding the pursuit of those ends, and a structure of agency for promoting those ends according to those judgments, and they will give rise to a group agent of that kind. That agent will pursue various actions, and do so with intention. And the intentions it enacts will be its own: they will be liable to emerge in the same discontinuous way that judgments may materialize (cf. Rovane 1997).

We hope that enough has been said to establish the possibility that certain patterns of joint action can lead to the formation of group agents that are distinct as agents from their own members. And we hope that the distinctive character of this possibility will serve to reinforce our argument in the last section that joint action on its own is not sufficient for the appearance of such a group agent. The theory of group agency is a theory about a certain domain of joint action. It presupposes whatever truths hold of joint action in general but it is not derivable from those truths alone; it represents an independent body of knowledge.

7. In Conclusion, Some Ontology

There are two major debates that have dominated social ontology: one is centered on individualism, the other on atomism (Pettit 1993). Both relate to possibilities involving the intentional psychology that we ascribe to ourselves when we take our actions to be the more or less rational products of more or less rationally held beliefs and desires; they both address the connection between mind and society. The individualism debate concerns the question of whether our individual intentional psychologies are compromised in any way by social regularities: whether we are predetermined or predestined, notwithstanding our apparent intentional powers, to behave so that the regularities are sustained. The atomism debate concerns the question of whether there are any aspects of our individual intentional psychology such that we depend noncausally on having certain relations with one another for

instantiating those features; as usually formulated, the issue is whether we depend on such relations for having the capacity to reason and think in a purposeful way.

Where do the issues we have been discussing fit on the broad canvas of social ontology? They point us, we think, to a third debate, one involving what may be called singularism rather than individualism or atomism (Pettit 2004). Like the other two controversies, this also relates to a possibility involving intentional psychology; it also connects with mind and society. It is the debate as to whether there are centers of intentional attitude and action over and beyond singular agents; whether singular subjects ever combine to form novel, plural centers of intentional life: minds of their own (Pettit 2003). The discussions of joint action and group agency are all directed, one way or another, towards this question. Those discussions culminate, if we are right, in the claim that, yes, there can be group agents of this relatively novel variety. Singularism, quite simply, is false.

The reason for saying that the singularism debate is distinct from the other two is that it is possible to be a singularist or a nonsingularist consistently with adopting any of the four possible positions on the other two debates: consistently with being an individualist atomist, with being a nonindividualist nonatomist, or with holding either of the two mixed positions.¹⁶ There is no difficulty in seeing how singularism itself, as the denial of plural centers of intentional life, is consistent with any of these positions. But what about nonsingularism of the sort that we ourselves endorse? Does it presuppose atomism or nonatomism, individualism or nonindividualism?

It does not presuppose atomism or nonatomism, since the argument for nonsingularism does not turn on whether atomism is true or false. There are no logical connections in evidence between the doctrines. The most that can be said is that a certain sort of nonatomism does give a very salient role to joint action, if not to the formation of group agents. This is the position according to which people are noncausally dependent on relations with one another for the capacity to reason and think so far as they use one another as means by which to triangulate on rules that their dispositions are meant to make salient: in particular, rules for the use of semantically basic predicates (Pettit 1993; Pettit 2002, ch. 4). The triangulation that this story requires involves joint action at the most basic level of human interaction; it requires that now in this instance, now in that, we aspire implicitly to co-determine right ways to go on in the use of basic concepts.

Does nonsingularism presuppose individualism or nonindividualism? It might seem to presuppose individualism, because the story we told in argu-

16. For an argument that individualism and atomism are themselves independent, leaving the four possibilities open, see Pettit (1993, 1996).

ing for nonsingularism certainly involves people coming together on the basis of their individual, intentional attitudes. But consistently with that story, of course, nonindividualism might still be true. It might be true in areas of behavior that are removed from group formation. And, more interestingly, it might even be true of group formation itself. Did groups form and operate among human beings in the manner of the apian swarm mentioned earlier, then the regularities associated with that group behavior would exist independently of the intentional attitudes of individual members—we would be selectionally predestined to enact the regularities—and anti-individualism would be vindicated.

We conclude that the singularism issue has a rightful, autonomous place in social ontology, side by side with the more traditional issues of individualism and atomism. Discussions of joint action, in particular of the joint action that may lead to group formation, are not just of interest in doing detailed, social theory. Like discussions of individualism and atomism, they bear on the central concern of social ontology: the relationship, or rather the relationships, between mind and society.

References

- Bacharach, M. 1998. Interactive team-reasoning: A contribution to the theory of co-operation. *Research in Economics* 58:117-47.
- Baier, A. C. 1997. Doing things with others: The mental commons. In *Commonality and particularity in ethics*, edited by Lilli Alanen, Sara Heinamaa, and Thomas Wallgren, 15-44. Basingstoke, UK: Macmillan.
- Bratman, M. 1987. *Intention, plans, and practical reason*. Cambridge, MA: Harvard University Press.
- . 1999. *Faces of intention: Selected essays on intention and agency*. Cambridge: Cambridge University Press.
- Cohen, P. R., and H. J. Levesque. 1991. Teamwork. *Noûs* 25:487-512.
- Dennett, D. 1987. *The intentional stance*. Cambridge, MA: MIT Press.
- Elster, J. 1983. *Explaining technical change*. Cambridge: Cambridge University Press.
- Gilbert, M. 1989. *On social facts*. Princeton, NJ: Princeton University Press.
- . 1996. *Living together: Rationality, sociality, and obligation*. Lanham, MD: Rowman & Littlefield.
- . 2000. *Sociality and responsibility: New essays in plural subject theory*. Lanham, MD: Rowman & Littlefield.
- . 2004. Scanlon on promissory obligation: the problem of promisees' rights. *Journal of Philosophy* 101:83-109.
- Holmström-Hintikka, G., and R. Tuomela, eds. 1997. *Social action*, vol. 2 of *Contemporary action theory*. Dordrecht: Kluwer Academic.
- Hurley, S. 1989. *Natural reasons*. New York: Oxford University Press.
- Koepsell, D. R., and L. S. Moss, eds. 2003. *John Searle's ideas about social reality: Extensions, criticisms and reconstructions*. Oxford: Blackwell.
- Lagerspetz, E., H. Ikaheimo, and J. Kotkavirta, eds. 2001. *On the nature of social and institutional reality*. Jyväskylä, Finland: SoPhi.
- Lewis, D. 1969. *Convention*. Cambridge, MA: Harvard University Press.

- List, C., and P. Pettit. 2002. The aggregation of sets of judgments: An impossibility result. *Economics and Philosophy* 18:89-110.
- . Forthcoming. Group agency and supervenience. *Southern Journal of Philosophy*.
- Meggle, G., ed. 2002. *Social facts and collective intentionality*. Frankfurt am Main: Deutsche Bibliothek der Wissenschaften.
- Meijers, A. 1994. *Speech acts, communication and collective intentionality: Beyond Searle's individualism*. Utrecht: de Jonge.
- . 2003. Can collective intentionality be individualized? In *John Searle's ideas about social reality: Extensions, criticisms and reconstructions*, edited by D. R. Koepsell and L. S. Moss, 168-83. Oxford: Blackwell.
- Miller, S. 1992. Joint action. *Philosophical Papers* 21:1-23.
- . 1995. Intentions, ends and joint action. *Philosophical Papers* 24:51-66.
- . 2001. *Social action*. Cambridge: Cambridge University Press.
- . 2003. Against collective agency. In *Social facts and collective intentionality*, edited by G. Meggle, 273-98. Frankfurt am Main: Deutsche Bibliothek der Wissenschaften.
- Petersson, Bjorn. 2005. *Collectivity and circularity*. University of Lund.
- Pettit, P. 1993. *The common mind: An essay on psychology, society and politics*. New York: Oxford University Press.
- . 1995. The virtual reality of Homo economicus. *Monist* 78: 308-29.
- . 1996. *The common mind: An essay on psychology, society and politics*. Paperback ed. New York: Oxford University Press.
- . 2000. The virtual reality of Homo economicus. Expanded version. In *The world of economics*, edited by U. Maki. Cambridge: Cambridge University Press, 2000. (Reprinted in Pettit 2002.)
- . 2001. *A theory of freedom: From the psychology to the politics of agency*. Cambridge and New York: Polity and Oxford University Press.
- . 2002. *Rules, reasons, and norms: Selected essays*. Oxford: Oxford University Press.
- . 2003. Groups with minds of their own. In *Socializing Metaphysics*, edited by F. Schmitt. New York: Rowman & Littlefield.
- . 2004. *Penser en société: Selected essays*. Paris: PUF.
- Pettit, P., and M. Smith. 2004. The truth in deontology. In *Reason and value: Themes from the moral philosophy of Joseph Raz*, edited by R. J. Wallace, Philip Pettit, S. Scheffler, and M. Smith. Oxford: Oxford University Press.
- Rovane, C. 1997. *The bounds of agency: An essay in revisionary metaphysics*. Princeton, NJ: Princeton University Press.
- Scanlon, T. M. 1990. Promises and practices. *Philosophy and Public Affairs* 19 (3): 199-226.
- Schmitt, F., ed. 2003. *Socializing metaphysics: The nature of social reality*. Lanham, MD: Rowman & Littlefield.
- Searle, J. 1990. Collective intentions and actions. In *Intentions in communication*, edited by P. Cohen, J. Morgan, and M. E. Pollack. Cambridge, MA: Bradford Books/MIT Press.
- Seeley, T. D. 2001. Decision making in superorganisms: How collective wisdom arises from the poorly informed masses. In *Bounded rationality: The adaptive toolbox*, edited by G. Gigerenzer and R. Selten, 249-61. Cambridge, MA: MIT Press.
- Sintonen, M., P. Ylikoski, and K. Miller. 2003. *Realism in action: Essays in the philosophy of the social sciences* (festschrift for Raimo Tuomela). Dordrecht: Kluwer Academic.
- Stoutland, F. 1997. Why are philosophers of action so anti-social? In *Commonality and particularity in ethics*, edited by Lilli Alanen, Sara Heinamaa, and Thomas Wallgren, 45-74. Basingstoke, UK: Macmillan.

- Sugden, R. 1993. Thinking as a team: Towards an explanation of nonselfish behavior. *Social Philosophy and Policy* 10:69-89.
- . 2000. Team preferences. *Economics and Philosophy* 16:175-204.
- . 2003. The logic of team reasoning. *Philosophical Explorations* 6:165-81.
- Tollefsen, D. 2002a. Collective intentionality and the social sciences. *Philosophy of the Social Sciences* 32 (1): 25-50.
- . 2002b. Organizations as true believers. *Journal of Social Philosophy* 33 (3): 395-401.
- Tuomela, R. 1984. *A theory of social action*. Dordrecht: Reidel.
- . 1989a. Actions by collectives. *Philosophical Perspectives* 3:471-96.
- . 1989b. Collective action, supervenience, and constitution. *Synthese* 80:243-66.
- . 1991a. We will do it: An analysis of group intentions. *Philosophy and Phenomenological Research* 51:249-77.
- . 1991b. Intentional single and joint action. *Philosophical Studies* 62:235-62.
- . 1995. *The importance of us: A philosophical study of basic social notions*. Stanford, CA: Stanford University Press.
- . 2000. *Cooperation: A philosophical study*. Dordrecht: Kluwer Academic.
- . 2002. *The philosophy of social practices: A collective acceptance view*. Cambridge: Cambridge University Press.
- . 2003. The we-mode and the I-mode. In *Socializing metaphysics*, edited by F. Schmitt. New York: Rowman & Littlefield.
- . Forthcoming. We-intentions revisited. *Philosophical Studies*.
- Tuomela, R., and K. Miller. 1988. We-intentions. *Philosophical Studies* 53:367-89.
- Tuomela, R., and M. Tuomela. 2003. Acting as a group member and collective commitment. *Protosociology* 18-19:7-65.
- Velleman, J. D. 1997. How to share an intention. *Philosophy and Phenomenological Research* 57:29-51.
- Zaibert, L. A. 2003. Collective intentions and collective intentionality. In *John Searle's ideas about social reality: extensions, criticisms and reconstructions*, edited by D. R. Koepsell and L. S. Moss, 209-32. Oxford: Blackwell.

Philip Pettit is L.S. Rockefeller University Professor of Politics and Human Values in Princeton University. Among his recent publications are a collection of his own papers, *Rules, Reasons and Norms* (Oxford University Press, 2002); a collection of joint pieces with Frank Jackson and Michael Smith, *Mind, Morality and Explanation* (OUP, 2004); plus a new book, *The Economy of Esteem* (with Geoffrey Brennan: OUP, 2004).

David Schweikard is completing his Ph.D. from his position as lecturer in philosophy at the University of Cologne.