

Planning for Collective Agency

Stephen A. Butterfill
University of Warwick, UK
<s.butterfill@warwick.ac.uk>

Abstract

Which planning mechanisms enable agents to coordinate their actions, and what if anything do these tell us about the nature of collective agency? On the leading, best developed account, Michael Bratman's, collective agency is explained in terms of interconnected planning. For our plans to be interconnected is for them to concern not just facts about our environment and goals but also facts about each others' plans. This chapter contrasts interconnected with parallel planning. In parallel planning, we each individually plan all of our actions and so are in a position to conceive of our own and each other's actions as parts of a single plan or exercises of a single ability. (The very idea of parallel planning may initially seem incoherent; the chapter examines this issue.) Could parallel rather than interconnected planning underpin collective agency? Some considerations in favour of a positive answer are provided by appeal to recent evidence on the role of motor representation in coordinating exercises of collective agency.

Acknowledgements: Much of what follows has been shaped by objections and suggestions from readers of drafts and from audiences at talks, including Olle Blomberg, Michael Bratman, Gergely Csibra, Naomi Elia, Chris Frith, Mattia Gallotti, Eileen John, Guenther Knoblich, Guy Longworth, John Michael, Marlene Meyer, Catrin Misselhorn, Elisabeth Pacherie, Wolfgang Prinz, Johannes Roessler, Thomas Sattig, Hans Bernhard Schmid, Natalie Sebanz, Corrado Sinigaglia, Thomas Smith, Joel Smith, Matthew Soteriou, Anna Strasser, Cordula Versper, Hong Yu Wong, and some peculiarly helpful anonymous referees.

1. Introduction

Which planning mechanisms enable agents to coordinate their actions, and what if anything do these tell us about the nature of collective agency? In this chapter I shall address these questions. But first I want to step back and consider how we might get a pre-theoretical fix on a notion of collective agency.

In everyday life we exercise agency both individually and collectively. Now Hannah is climbing the tree alone but later she will return with Lucas and they will exercise collective agency in climbing the tree together. When two or more people dance or walk together, when they kiss, when they move a rock together or when they steal a spaceship together—then, typically, they are exercising collective agency. But what is collective agency? Fortunately philosophers have offered some suggestions. Unfortunately philosophers have offered many different, wildly incompatible suggestions about collective agency.¹ Do we just have to pick a view and hope for the best, or can we get a fix on the notion in advance of choosing a favourite philosopher?

2. Three Contrasts

One way to anchor our thinking is by contrasting paradigm cases involving collective agency with cases that are as similar as possible but do not involve collective agency.² Suppose lots of commuters pile into an elevator and, because of their combined weight, cause it to get stuck. Then they have collectively broken the elevator, and this is something they have done together. (This story is not one in which each commuter enters the elevator alone on a different occasion and causes it to get stuck.) It does not follow that their actions involve collective agency. After all, there are different ways we could extend the story, not all of which involve collective agency. First, we could extend the story by specifying that each individual is acting alone, unconcerned about the others. In fact, each commuter might have attempted to prevent others from getting into the elevator. In this case, the

¹ Recent contributions include Bratman (2014); Gilbert (1990); Gallotti (2011); Gold and Sugden (2007); Kutz (2000); Ludwig (2007); Miller (2001); Schmid (2009); Searle (1990); Seemann (2009); Smith (2011); Tuomela and Miller (1988); Tuomela (2005).

Even the terminology is fraught with pitfalls. I take ‘collective agency’ and ‘shared agency’ to be synonymous; likewise ‘collective action’ and ‘joint action’. Use of the term ‘collective’ occasionally cues audiences to expect discussion of large-scale activities, although in this chapter ‘collective’ should be understood, as it is in discussions of plural prediction, as contrasting with ‘distributive’ (see, e.g., Linnebo 2005).

² The use of contrast cases to draw conclusions about collective agency is not new (compare Searle 1990). The strategy is familiar from Pears (1971), who used contrast cases to argue that whether something is an ordinary, individual action depends on its antecedents.

commuters only exercise merely individual agency in parallel. For a second way of extending the story, imagine a case that is as similar as possible except that the people involved are not really commuters but are accomplices of a master criminal posing as commuters. In accordance with an earlier agreement they made, they all pile into the elevator in order to put it out of order during a robbery. This would typically involve an exercise of collective agency on the part of the accomplices.

Contrasting the genuine commuters with the criminal accomplices indicates something about what collective agency isn't. It indicates that our exercising collective agency cannot be a matter only of our doing something together, nor only of there being something that we collectively do.³

At this point one might easily be struck by the thought that there is one goal to which each criminal accomplice's actions are directed, namely the breaking of the elevator. Since none of the genuine commuters aim to break the elevator, perhaps this gives us a handle on collective agency. Can we say that collective agency is involved where there is a single outcome to which several agents' actions are all directed?

Unfortunately we probably cannot. To see why not, consider a second pair of contrasting cases. In the first case, two friends decide to paint a particular bridge red together and then do so in a way typical of friends doing things together. This is a case involving collective agency. But now consider another case. Two strangers each independently intend to paint a large bridge red. More exactly, each intends that her painting grounds⁴ or partially grounds⁵ the bridge's being painted red. They start at either end and slowly cover the bridge in red paint working their ways towards the middle where they meet. Because the bridge is large and they start from different ends, the two strangers have no idea of each other's involvement until they meet in the middle. Nor did they expect that anyone else would be involved in painting the bridge red. This indicates that they were not exercising collective agency.

Contrasting the two painting episodes, the friends' and the strangers', indicates that our exercising collective agency cannot be just a matter of there being a single outcome to which each of our actions is directed. After

³ If this is right, Gilbert is wrong that '[t]he key question in the philosophy of collective action is simply ... under what conditions are two or more people doing something together?' (Gilbert 2010, p. 67).

⁴ Events D_1, \dots, D_n *ground* E just if: D_1, \dots, D_n and E occur; D_1, \dots, D_n are each part of E ; and every event that is a part of E but does not overlap D_1, \dots, D_n is caused by some or all of D_1, \dots, D_n . This notion of grounding is adapted from Pietroski (1998).

⁵ Event D *partially grounds* event E if there are events including D which ground E . (So any event which grounds E thereby also partially grounds E ; I nevertheless describe actions as 'grounding or partially grounding' events for emphasis.) Specifying the intentions in terms of grounding ensures that it is possible for both people to succeed in painting the bridge, as well as for either of them to succeed alone.

all, this feature is common to both painting episodes, the friends' and the strangers', but only the former involves collective agency.

What are we missing? If we want to try to keep things simple (as I think we should), a natural consideration is that the friends painting the bridge red care in some way about who is involved in the painting whereas the strangers do not. Maybe what is needed for collective agency is not just a single outcome to which all agents' actions are directed but also some specification of whose actions are supposed to bring this outcome about. Consider this view:

The Simple View. Collective agency is involved where there is a single outcome, G , and several agents' actions are all directed to the following end: they, these agents, bring about this outcome, G , together.

Note the appeal to togetherness in specifying the end to which each agents' actions are directed. Is this circular? No. As we saw in discussing the first contrast (the one with the elevator), there can be something we are doing collectively and together without our exercising any collective agency at all.

One way—perhaps not the only way—to meet the condition imposed by the Simple View involves intention: the condition is met where several agents are acting in part because⁶ each intends that they, these agents, bring about G together. In the case of the friends painting the bridge red, the idea is that each would intend that they, the two of them, paint the bridge red together.

The Simple View does enable us to distinguish the two painting episodes: the friends painting the bridge exercise collective agency and their actions are each directed to an end involving them both, whereas this is untrue of the strangers. But is the view right?

Probably not. We can see why not by adapting an example from Gilbert and Bratman. Contrast two friends walking together in the ordinary way, which paradigmatically involves collective agency, with a situation where two gangsters are walking together but each is forcing the other. It works like this. The first gangster pulls a gun on the second gangster and says, sincerely, 'Let's walk!' Simultaneously, the second does the same to the first. Each intends that they, these two gangsters, walk together. This ensures that the condition imposed by the Simple View is met. Yet no collective agency is involved. At least, it should be clear that no collective agency is involved unless you think the central events of *Reservoir Dogs* involve collective agency (Tarantino 1992).

⁶ Here I ignore complexities involved in accurately specifying how events must be related to intentions in order for the events to involve exercises of collective agency; these parallel the complexities involved in the case of ordinary, individual agency. (On the individual case, see Chisholm 1966.)

Contrasting the ordinary way of walking together with what we might call walking together in the Tarantino sense indicates that the Simple View is too simple. Apparently we cannot distinguish actions involving collective agency just by appeal to there being an outcome where several agents' actions are directed to the end that they, these agents, bring about that outcome together.⁷

Maybe this is too quick. Maybe we should be neutral about whether there is no collective agency involved when the gangsters walk together in the Tarantino sense. It should be clear, of course, that there is a contrast with respect to collective agency between the ordinary way of walking together and walking together in the Tarantino sense. But what can we infer from this contrast? Suppose we are neutral on whether there are degrees of collective agency, or on whether there are multiple kinds of it. Then we cannot infer from the contrast that no collective agency is involved in walking together in the Tarantino sense. The contrast may be due to a difference in degree or kind rather than to an absence of collective agency. So rather than saying that the Simple View does not allow us to distinguish actions involving collective agency from actions that do not, it would be safer to say that this view does not enable us to make all the distinctions with respect to collective agency that we need to make.

In any case, the Simple View cannot be the whole story about collective agency. Focus on cases like two friends walking together or painting a bridge together in the ordinary way. To introduce an arbitrary label, let *alpha* collective agency be the kind or degree of collective agency involved in cases such as these. (And if there are no kinds or degrees of collective agency, let alpha collective agency be collective agency.) What distinguishes exercises of alpha collective agency from exercises of other kinds or degrees of collective agency (if there are any) and from parallel but merely individual exercises of agency?

3. Bratman's Proposal

The best developed and most influential proposal is due to Michael Bratman. Others defend conceptually radical proposals, proposals which involve introducing novel attitudes distinct from intentions (e.g. Searle 1990), novel subjects distinct from ordinary, individual subjects (e.g. Schmid 2009), or novel kinds of reasoning (e.g. Gold and Sugden 2007). Bratman's proposal aims to be, as he puts it, conceptually conservative. He proposes that we can give sufficient conditions for what I am calling alpha collective agency merely by adding to the Simple View.

⁷ This argument is adapted from Bratman's (1992, pp. 132–4; 2014, pp. 48–52).

The key part of Bratman's proposal is straightforward.⁸ It is not just that we each intend that we, you and I, perform the action (as the Simple View requires); further, we must have, and act on, intentions about these intentions. As Bratman writes:

‘each agent does not just intend that the group perform the [...] joint action.⁹ Rather, each agent intends as well that the group perform this joint action in accordance with subplans (of the intentions in favor of the joint action) that mesh’ (1992, p. 332).

This appeal to interlocking intentions enables Bratman to avoid counterexamples like the Tarantino walkers. If I am acting on an intention that we walk by way of your intention that we walk, I can't rationally also point a gun at you and coerce you to walk.

Why not? If you are doing something by way of an intention to do that very thing, then, putting any special cases aside, I can't succeed in coercing you to do it. (True, I might be coercing you to intend to do it; but that is different from coercing you to do it. After all, someone in a particularly contrary mood might coherently coerce you to intend to do something while simultaneously preventing you from acting.) Accordingly, if I intend both that we walk by way of your intention that we walk and by way of my coercing you to walk, then I have intentions which cannot both be fulfilled. Since the incompatibility of these intentions in any particular case is, on reflection, obvious enough to me, I can't rationally have this combination of intentions unless perhaps I am unreflective or ignorant.

Bratman's proposal hinges on adding second- to first-order intentions (give or take some common knowledge, interdependence and other details that need not concern us here). On the face of it, it would be surprising if this worked. After all, one of the lessons from a parallel debate concerning ordinary, individual agency appears to be that if something can't be captured with first-order mental states, invoking second-order mental states will not suffice to capture it either (compare Watson 1975, pp. 108–9 on Frankfurt 1971). Relatedly, anyone familiar with a certain drama surrounding Grice on meaning would have to be quite optimistic to bet on Bratman's proposal. Nevertheless, to my knowledge no one has yet succeeded in showing that Bratman's proposal is wrong.

⁸ His proposal has been refined and elaborated over more than two decades (for three snapshots, compare Bratman 1992, p. 338; Bratman 1997, p. 153; and Bratman 2014, p. 84). Here I skip the details; the discussion in this chapter applies to all versions.

⁹ It may be tempting to think that invoking joint action here is somehow circular. But Bratman is using 'joint action' as I am using 'collective action'; and, as illustrated in Section 2, there are collective (or 'joint') actions which do not involve collective agency.

4. A Fourth Contrast

I seek a fourth contrast to add to the earlier three (see Section 2). This will be a pair of cases in both of which Bratman's conditions are met, although one involves alpha collective agency whereas the other does not. I do not claim that the contrast considered in this section shows that Bratman's proposal is wrong. My aim is only to motivate considering alternatives to Bratman's proposal by indicating the sort of case that might yield a counterexample.

As background, recall that there can be something that several people do collectively without exercising collective agency, as in the story about commuters breaking the elevator from Section 2. One consequence is that an individual can unilaterally have intentions concerning what several people will collectively do. (By saying that someone has such an intention *unilaterally*, I mean she intends irrespective of whether others intend the same.) To illustrate, suppose that Ayesha, one of the commuters, knows that all the others will get into the elevator irrespective of whether she elbows her way in or not. Suppose, further, that she knows that the elevator will break only if all of the commuters including herself are in it. Then she can intend, unilaterally, that they, the commuters, collectively break the elevator. (Detailed defence of this sort of possibility against objections about what individuals can intend is given in Bratman (1997).) I shall exploit this possibility in describing a pair of cases that contrast with respect to collective agency.

Ayesha and Beatrice each have one wrist handcuffed to the steering wheel of a moving car. They are matched in strength closely enough that neither can decide the car's course alone: its movements will be a consequence of both of their actions. Ayesha, determined that Beatrice should die and wishing to die herself, is wondering how she could bring this about. Thinking that she could pull her gun on Beatrice to force her to cooperate, she intends, unilaterally, that they, Ayesha and Beatrice, drive the car off the road and over a cliff. But a sudden jolt causes the gun to fly from her hand and land far out of reach. Just as it seems she will have to abandon her intention, it strikes her that Beatrice has an intention which renders the gun unnecessary. For Beatrice, whose thoughts and actions mirror Ayesha's, plainly intends what Ayesha intends, namely that they drive the car over the cliff. So Ayesha retains this intention and changes her mind only about the means. Whereas before she intended that they do this by way of her gun, she now intends that they do it by way of her and Beatrice's intentions that they drive the car over the cliff. Beatrice's intention renders Ayesha's gun unnecessary. Now Beatrice, continuing to mirror Ayesha, also forms an intention about their intentions. Further, all this is common knowledge between them. So this is a case in which the conditions imposed in Bratman's proposal about collective agency are met: the agents not only each intend that they perform the action (driving the car over the cliff) but also intend that they do so in

accordance with these intentions and meshing subplans of them.

Contrast this case with Thelma and Louise's better known and more romantic intentional car crash (the two friends evade capture by driving off a cliff together; Khouri 1992). Whereas Thelma and Louise's escape is a paradigm case of collective agency, the episode involving Ayesha and Beatrice does not seem to involve collective agency at all, and certainly not alpha collective agency. For Ayesha conceives of Beatrice's intention merely as an opportunity to exploit, something that achieves what she would prefer to have ensured by force of arms. And Beatrice does likewise. Each is willing to be exploited in order to be able to exploit the other. While not all forms of exploitation are incompatible with collective agency, it seems unlikely that collective agency can consist only in the kind of mutual exploitation exemplified by Ayesha and Beatrice. This indicates that Bratman's proposal does not yield sufficient conditions for collective agency. Or at least the conditions are not sufficient for alpha collective agency, that is, for the degree or kind of collective agency the proposal is supposed to capture. Apparently, appealing to intentions about intentions does not enable us to distinguish actions involving alpha collective agency from actions not involving it.

This fourth contrast is a challenge to Bratman's proposal just as the third contrast (the one with the gangsters walking in the Tarantino sense) is a challenge to the Simple View. Someone who wanted to reject the Simple View and endorse Bratman's proposal would need to hold that something like the third contrast is genuine while insisting the fourth contrast is merely apparent.

Is this a decisive objection to Bratman's proposal? Clearly it is not. After all, one might deny that the fourth contrast is genuine and insist that, whatever intuitions anyone might have to the contrary, Ayesha and Beatrice are exercising alpha collective agency. This form of response could be made to any of the contrasts offered; the objection to the Simple View, for instance, is similarly non-decisive. Other responses to the contrasts are surely conceivable too.

What, then, can we conclude from the four contrasts? It would be a mistake simply to disregard them. True, they do not provide decisive grounds for rejecting the views we have so far considered. But none of the narrowly philosophical considerations offered in this area have provided decisive reasons for accepting or rejecting a view about collective agency. The contrasts are valuable because they indicate that something is missing from even the best, most carefully developed attempt to distinguish alpha collective agency from other kinds or degrees of collective agency and from parallel but merely individual agency. What is missing?

5. Interconnected Planning vs Parallel Planning

Bratman's proposal about collective agency is part of an attempt to provide a planning theory of agency, both individual and collective. Let us say that our plans are *interconnected* just if facts about your plans feature in mine and conversely. Then Bratman's proposal about collective agency (see Section 3 on page 5) implies that where our actions are guided by appropriately interconnected planning, we are exercising alpha collective agency. In the previous section I attempted to motivate doubt about whether interconnected planning is sufficient for alpha collective agency by appeal to Ayesha and Beatrice, whose actions are driven by interconnected planning but who, apparently, do not exercise alpha collective agency. Could we deny that interconnected planning is sufficient for collective agency while nevertheless holding on to a more general insight about planning—the insight that we can understand something about how collective agency differs from parallel but merely individual agency by considering the kinds of planning characteristic of the former? My aim in the rest of this chapter is to show that we could.

Many collective actions are subject to a certain kind of relational constraint. Suppose that you and I are tasked with moving a table out of a room, through a narrow doorway. When, where and how you should grasp, lift and move depends on when, where and how I will do these things; and conversely. How could we meet such relational constraints on our actions?

One way is by means of interconnected planning. When moving the table, interconnected planning would require two things of you (and likewise of me). You would need to form a view about my plans for my part of our action. And you would need to make plans for your part which mesh with what you know of my plans. In this case we each have in mind two separate plans—one for your actions and one for mine—that need to be interconnected. The need for interconnection means that the plan you end up with must involve not only facts about the size and weight of the table and the aperture of the doorway but also facts about my plans. Interconnected planning is meta-planning.

But there is another, perhaps more natural way of proceeding that doesn't require interconnected planning. Instead of considering separately how you and I should each act and then needing to interconnect these two things, you might simply consider how two people in our situation should move the table through the door; and I might do likewise. In this case, the question that guides preparation for action is not, How should I act? but, How should we act? Answering this question results in us each having in mind one plan for our actions (rather than two separate plans, one for your actions and one for mine). So no interconnection is needed. This second approach to moving the table exemplifies *parallel planning*: we each make a plan for all of our actions.

So the difference between interconnected and parallel planning concerns how many plans we each have in mind. In interconnected planning, we each have two plans in mind, one for my actions and one for yours. In parallel planning, we each have just one plan in mind, a single plan that describes how we, you and I,¹⁰ should act.

The distinction between having one plan in mind and having two plans in mind may initially seem too subtle to matter. A plan is a structure of representations of outcomes, where the structure is subject to the normative requirements that the outcomes represented can be partially ordered by the means-end relation yielding a tree with a single root, and that the way the representations are structured corresponds to this tree. Any two separate plans that you could both execute can be combined into a single plan. The single plan is simply to implement each of the two formerly separate plans. This might seem to indicate, misleadingly, that there is no significant difference between making a single plan and making two separate plans.

But there are significant differences. One difference arises from the ways plans can be structured. This can be seen by reflection on individual agency. Suppose you have to organise two events, a workshop and a wedding. If the two events were largely independent of each other, it would be simplest to make two separate plans for them. But suppose there are many constraints linking the two events. Some participants at the workshop are also wedding guests, and it happens that many choices of transport, entertainment and catering for one event constrain possible choices for the other event. Envisaging the complex interdependence of these two plans, you might reasonably be inclined to make a single plan for the wedding and workshop. This could be simpler than making two separate plans because it allows for greater flexibility in structuring subplans than would making separate plans for the workshop and wedding and then combining them. For instance, if you make separate plans for the wedding and workshop, you need subplans for transport in each plan and so are forced to divide the problem of transporting participants into the problem of transporting them for the workshop and the problem of transporting them for the wedding. Making a single plan for both events allows you to treat transport as a single problem (but it does not require you to do so, of course). This illustrates one way in which our each having two plans in mind, as happens in interconnected planning, can be significantly different from our each having a single plan for all of our actions in mind, as happens in parallel planning.

¹⁰ In specifying that a plan describes how we, you and I, should act, I do not mean to imply that the plan must be, or include, a plan for how you should act where this is something over and above being a plan for how we should act. (Distinguishing senses in which a plan on which I act might be a plan for how you should act requires some care.) I specify 'you and I' just to emphasise that the plan is supposed to answer, partially or wholly, the question, What should we do? and not only the question, What should I do?

Perhaps distinguishing collective agency from parallel but merely individual agency requires us to focus, not on interconnected planning, but on parallel planning. Some readers may already wish to object that invoking the notion of parallel planning is somehow incoherent. I will eventually consider grounds for this objection (see the end of Section 6) and how it might be overcome (see Sections 7 and 8). But for now please suspend disbelief and let me first explain why, assuming the notion is coherent, reflection on parallel planning may provide insights about collective agency.

How does parallel planning enable agents to coordinate their actions? Suppose you and I are parents about to change our baby's nappy. This involves preparing the baby and preparing the nappy. You're holding the baby and I'm nearest the pile of clean nappies, so there's a single most salient way of dividing the task between us. Preparing the baby is, of course, a complex action with many components. Now there are relational constraints on how the baby and nappy should be prepared; how you clean constrains, and is constrained by, how I prepare the clean nappy (because we don't want to get pooh on it). How do we meet these relational constraints? Suppose that we engage in parallel planning, and that, thanks to environmental constraints, our locations and planning abilities, we predictably and non-accidentally end up with plans that match.¹¹ Your having a single plan for our actions, yours and mine, means that your plan for your actions is constrained by your plan for my actions. And the fact that our plans match means that your plan for my actions is, or matches, my plan for my actions. So thanks to our parallel planning—to the fact that we each plan the whole action—your plan for your actions is indirectly constrained by my plan for my actions; and conversely.

As this example illustrates, parallel planning sometimes enables us to meet relational constraints on our actions not by thinking about each other's plans or intentions but, more directly, by planning each other's actions. By virtue of parallel planning, you can use the same planning processes which enable you to meet constraints on relations between several of your own actions in order to meet constraints on relations between your own and others' actions.

The fact that parallel planning enables agents to coordinate their actions in this way does not imply that it can help us to understand collective agency. After all, if there were no more to exercising collective agency than merely

¹¹ Two or more agents' plans *match* just if they are the same, or similar enough that the differences don't matter in the following sense. First, for each agent's plan, let the *self part* be the bits concerning the agent's own actions and let the *other part* be the other bits. Now consider what would happen if, for a particular agent, the other part of her plan were as nearly identical to the self part (or parts) of the other's plan (or others' plans) as psychologically possible. Should this agent's self part be significantly different? If not, let us say that any differences between her plan and the other's (or others') are not relevant for her. Finally, if for some agents' plans the differences between them are not relevant for any of the agents, then let us say that the differences don't matter.

coordinating actions, the difficulties in characterising collective agency partially surveyed in Section 2 would hardly arise. Why suppose that the notion of parallel planning—assuming for the moment that it is even coherent—might help us to understand collective agency?

6. Parallel Planning and Collective Agency

Suppose you and I are about to set up a tent in a windy field. If we first collectively plan how we will do this and then act on our plan, it seems clear enough that we are exercising collective agency. Can we invoke collective planning in order to explain collective agency? To do so might involve circularity because collective planning might itself involve collective agency. My suggestion is that invoking parallel planning enables us to capture much the same insight about collective agency that invoking collective planning would, but without the risk of circularity and with greater generality. Let me explain.

It is a familiar idea that, in thinking about your own future actions, there are two sorts of perspective you can adopt. You can adopt the perspective of an outsider and think about your own actions in a theoretical way. Alternatively, you can adopt the perspective of the agent and think about your own actions in a practical way. One good indicator that the theoretical and practical perspectives are distinct is that, as several philosophers have noted, in considering actions practically you should disregard biases, limits and non-rational quirks which mean your actions sometimes fail to conform to your intentions, even if these are highly reliable.

Can you adopt a practical perspective on actions which, we both know, I will perform? Initially it might seem that you could not do this without irrationality or confusion. But suppose I ask you how I could get around this obstacle (say). To answer this question you might imaginatively put yourself in my place and deliberate about how to get around the obstacle. In doing this you are taking a practical perspective on actions which I, not you, will eventually perform. You are not, of course, engaged in practical reasoning concerning those actions: your reasoning is, after all, not actually directed to action. Nevertheless, there is a fine line between what you are doing—call it practical deliberation—and practical reasoning. For suppose that, unexpectedly, you find yourself actually in my place and needing to get around the obstacle. You do not now need to start planning your actions: the planning has already been done. Indeed, for you to consider afresh how to act now that it turns out to be you rather than me who is facing the obstacle would indicate that you had not done your best to answer my question. So when imaginatively putting yourself in my place, you can adopt a practical perspective on actions which I will eventually perform and doing so actually

prepares you to act.

Invoking the notion of parallel planning involves going just a tiny step further. In parallel planning, it is not just that we each adopt a practical perspective on actions some of which the other will eventually perform. It is also that we each adopt a perspective from which both of our actions, yours and mine, are parts of a single plan. We are thereby thinking practically about all of these actions simultaneously. To return to setting up our tent in the windy field, rather than collectively planning our actions we might each individually deliberate about how to proceed and, realising that just one plan is most salient to both of us, spring into action without needing to confer. It is not that we are thinking of each other's actions exactly as if they were our own, of course; but we are also not thinking of them in a merely theoretical way, as the actions of someone who is just passing by. From the perspective we each adopt in parallel planning, our actions, yours and mine, have a certain kind of practical unity.

Why is this significant? Return for a moment to the two gangsters walking in the Tarantino sense and what I am calling the Simple View (see Section 2). According to the Simple View, for some agents' walking together to involve collective agency it is sufficient that their walking be appropriately guided by intentions, on the part of each agent, that they, these agents, walk together. Now, as you may recall, earlier I noted that the Simple View seems incorrect because gangsters might act on such intentions while forcing each other to walk at gunpoint. It is this sort of problem that Bratman uses to motivate complicating the Simple View with appeal to interconnected planning (compare Bratman 1992, pp. 132–4; Bratman 2014, pp. 48–52). But the problem might also be overcome by invoking parallel planning. Consider the view that for us to exercise (alpha) collective agency in walking together it is sufficient that:

1. we each intend that we, you and I, walk (the Simple View);
2. we pursue these intentions by means of parallel planning; and
3. our plans predictably and non-accidentally match.

This view correctly implies that the gangsters' walking is not an exercise of (alpha) collective agency. This is because pursuing an intention by means of parallel planning means taking a practical attitude towards each other's actions. Where the above conditions, (1)–(3), are met, for one of the agents to point a gun at the other to force her to walk would be almost like her pointing a gun at herself in order to force herself to do something she plans to do. This would involve a form of irrationality. So if the gangsters were pursuing their intentions that they walk by means of parallel planning, they could not rationally be forcing each other to walk at gunpoint.

Invoking parallel planning also enables us to distinguish the case of Ayesha and Beatrice from that of Thelma and Louise, which Bratman's proposal was unable to do (see Section 4 on page 7). Ayesha and Beatrice, who are handcuffed to the steering wheel of a moving car, each rely on the other's intention only for want of a gun. From their perspectives, the other's intention is just one among many factors that might have justified relying on the other to perform actions that bear a certain relation to her own. This shows that they each adopt a theoretical perspective when thinking about the others' actions and intentions. So they cannot be pursuing their intentions by means of parallel planning. By contrast, Thelma and Louise decide what they will do, which indicates that they each adopt the perspective required for parallel planning. So appeal to the above conditions, (1)–(3), but not to Bratman's proposal (see Section 3 on page 5), plausibly gives us what seems, pre-theoretically, to be the right result: Ayesha and Beatrice are not, whereas Thelma and Louise are, exercising (alpha) collective agency.

This is one reason for supposing that reflection on parallel planning may yield insights into what distinguishes collective agency (or the alpha variety of it), insights that cannot be got from reflection on interconnected planning. Whereas engaging in interconnected planning is consistent with thinking of all others' plans and intentions merely as opportunities to exploit or constraints to work around, engaging in parallel planning involves each agent taking a practical perspective on everyone's actions simultaneously and so conceiving of them as having a certain kind of practical unity.

There's just one tiny problem. Engaging in parallel planning seems, on the face of it, to be always either irrational or else dependent on confusion about whose actions are whose. Why? For us to engage in parallel planning is for each of us to plan both of our actions, yours and mine. This implies, of course, that we each plan actions that are not our own: if things go well, your plan for our actions will include actions the agent of which is not you but me; and likewise for my plan. But assuming (for now) that the elements of plans are intentions, this means that among the things you would be intending are some actions that I will eventually perform; and, likewise, among the things I am intending are some things that will eventually be your actions. So how could we engage in parallel planning without irrationality or confusion? After all, you can't knowingly intend my actions, at least not in the kinds of case we are considering (there may be other cases; see Roth 2010).

As things stand, then, anyone seeking to avoid conceptually radical innovation in theorising about collective agency has an awkward choice. She has a choice between existing proposals which seem to fail to give genuinely sufficient conditions for alpha collective agency and a new proposal involving conditions that, apparently, are met only when the agents are irrational or confused. What to do?

7. Two Lines of Argument

There are at least two lines of argument that might be used to show that parallel planning can occur without irrationality or confusion, thereby defending the idea that reflection on parallel planning will yield insights about collective agency. One line—call it the *hard line*—would be to argue that there is a propositional attitude which resembles intention in some essential respects, and which is inferentially integrated with intention, but which you can have towards actions even where you know that some of these actions will be performed by others in exercising collective agency with you. Perhaps, for instance, we can rescue parallel planning from charges of irrationality by identifying certain ways in which intentions or closely related attitudes can be open-ended not only with respect to means but also with respect to agents. While it may be tempting to take the hard line, I shall not do so here.¹²

There is an alternative line of argument which might be used to show that parallel planning can occur without irrationality or confusion. This involves considering processes and representations more primitive than full-blown planning and intention. As we shall see, there appear to be representations which are like intention in some ways but which can be had concerning others' actions (as well as your own) without irrationality or confusion.

This second line of argument is fraught with difficulties. Philosophers have barely begun to consider how discoveries about the mechanisms which make collective agency possible might bear on theories of what it is.¹³ The next section offers reasons for holding that if we broaden our view to include processes and representations more primitive than full-blown planning and intention, there are mundane cases in which parallel planning (or something resembling it in respects to be specified) need involve neither irrationality nor confusion.

¹² Support for the hard line might be extracted from Laurence (2011). He defends the view that, in some cases, several agents' 'individual, first-person-singular actions are all subject to the special collective action sense of the question "Why?" and [...] the same answer holds in each case' (p. 289). While there are many differences between Laurence's view and the line I am considering, both appear committed to the idea that exercising collective agency sometimes involves this: that from the point of view of any individual agent, it is almost as if all the agents' actions are guided by a single piece of practical reasoning.

¹³ This is not to say that philosophers have not attempted this at all; see, for example, Tollefsen (2005) or Gallotti and Frith (2013).

8. Parallel Planning and Motor Representation¹⁴

Consider the motor processes that enable you to reach for, grasp and transfer objects in a coordinated and fluid way. Are these planning processes? They are distinct from the sort of planning that might involve getting your diary out—planning to paint a house, or to travel from London to Stuttgart, say. Nevertheless, such motor processes resemble planning activities both in that they involve computing means from representations of ends (Bekkering et al. 2000; Grafton and Hamilton 2007) and in that they involve meeting constraints relating actions which must occur at different times (Jeanerod 2006; Zhang and Rosenbaum 2007; Rosenbaum et al. 2012). If we think about planning in a narrow way, then, plausibly, these motor processes are not planning processes. But given the two points of resemblance just mentioned, we can coherently follow much of the scientific literature in adopting a broader notion of planning, one that encompasses both planning-like motor processes and full-blown, get-your-diary-out planning. Here I shall use ‘planning’ in this second, broader sense.

Motor representations are representations of the sort that characteristically feature in motor processes; they play a key role in monitoring and planning actions (e.g. Wolpert et al. 1995; Miall and Wolpert 1996). Motor representations can be distinguished from intentions and representations of other kinds by their format, much as visual representations can be distinguished by their format (Butterfill and Sinigaglia 2014). Motor representations are not limited to representations of bodily configurations and joint displacements. In fact, some motor representations resemble intentions in representing outcomes such as the grasping of an object or the movement of an object from one place to another. Relatedly, some motor representations resemble intentions in coordinating multiple component activities by virtue of their role as elements in hierarchical, plan-like structures, and coordinating these activities in a way that would normally facilitate the represented outcome’s occurrence (Hamilton and Grafton 2008; Pacherie 2008, pp. 189–90).

How are motor representations relevant to our difficulties with the notion of parallel planning? Motor representations lead a kind of double life. For motor representations occur not only when you perform an action but sometimes also when you observe an action. Indeed there are some striking similarities between the sorts of processes and representations usually involved in performing a particular action and those which typically occur when observing someone else perform that action. In some cases it is almost

¹⁴ This section draws on work in progress with Corrado Sinigaglia and, separately, Natalie Sebanz and Lincoln Colling. But it is not endorsed by these researchers, who would probably have avoided the mistakes I will doubtless have made.

as if the observer were planning the observed action, only to stop just short of performing it herself.¹⁵ When motor representations of outcomes trigger a planning-like process in action observation, this may allow the observer to predict others' actions (Flanagan and Johansson 2003; Ambrosini et al. 2011, 2012; Costantini et al. 2013). There is no question, then, that motor representations concerning others' actions, not just your own, can occur in you.¹⁶ Perhaps, then, there could be parallel planning involving motor processes and representations. If so, the double life of motor representations tells us that parallel planning does not always involve irrationality or confusion.

Of course, one might doubt that motor representations can coherently lead a double life. After all, this might appear to involve one and the same attitude having two directions of fit, world-to-mind insofar as it is involved in planning for action and mind-to-world insofar as it is involved in observing and predicting others' actions. Avoiding this apparent incoherence may demand some subtleties; but as this issue is one that arises independently of issues about collective agency, I shall put it aside here. The overwhelming evidence that motor representations do in fact lead a double life is a reason, not decisive but compelling, to hold that the core idea must somehow be coherent.

My suggestion, then, is that parallel planning is not always incoherent or confused if it sometimes involves motor representations rather than intentions. But is there any evidence that parallel planning involving motor representations ever occurs? Planning concerning another's actions sometimes occurs not only in observing her act but also in exercising collective agency with her (Kourtis et al. 2013; Meyer et al. 2011). Such planning can inform planning for your own actions, and even planning that involves meeting constraints on relations between your actions and hers (Vesper et al. 2013; Novembre et al. 2013; Loehr and Palmer 2011; Meyer et al. 2013). This is suggestive but compatible with two possibilities. It could be that there is a single planning processes concerning all agents' actions, just as parallel planning requires; but it might also be that, in each agent, there are two largely separate planning processes, one for each agent's actions. Where the latter possibility obtains, there is no parallel planning. Fortunately, there is some neurophysiological and behavioural evidence for the former possibil-

¹⁵ For reviews, see Jeannerod (2006); Rizzolatti and Sinigaglia (2008, 2010). If motor representations occur in action observation, then observing actions might sometimes facilitate performing compatible actions and interfere with performing incompatible actions. Both effects do indeed occur, as several studies have shown (Brass et al. 2000; Craighero et al. 2002; Kilner et al. 2003; Costantini et al. 2012).

¹⁶ Note that this does not imply that others' actions are ever represented as others' actions motorically. It may be that some or all motor representations are agent-neutral in the sense that their contents do not specify any particular agent or agents (Ramsey and Hamilton 2010).

ity. Sometimes when exercising collective agency, the agents have a single representation of the whole action, not only separate representations of each agent's part (Tsai et al. 2011; Loehr et al. 2013; Ménoret et al. 2013). So while we should be cautious given that the most relevant evidence is relatively recent, it is reasonable to conjecture, at least provisionally, that parallel planning may sometimes involve motor representations rather than intentions. This conjecture is one way (perhaps not the only way) of rescuing the view that we can give sufficient conditions for (alpha) collective agency by invoking parallel planning.

Or is it? Consider again the fourth contrast, the contrast of Ayesha and Beatrice's driving off the cliff with Thelma and Louise's (see Section 4 on page 7). Earlier I suggested that Ayesha and Beatrice's actions couldn't rationally be driven by parallel planning because this would involve each taking a practical perspective on both of their actions, thereby making it irrational for her to consider forcing the other to act in something like the way it would be irrational for her to consider forcing herself to act. But now we are considering a broader notion of planning, one that does not invariably involve the agent adopting a practical perspective, or any perspective at all. This may make it unclear whether we can give sufficient conditions for collective agency by invoking parallel planning.

To solve this problem we must consider how intentions and motor representations are related. Motor representations are not themselves objects of awareness, but they do shape agents' awareness of the kinds of actions they can perform. Further, they do this in such a way that it is possible, sometimes at least, for intending that an outcome obtain (that you grasp this mug, say) to trigger a motor representation of a matching¹⁷ outcome. When an intention is appropriately related to awareness of the kinds of action one can perform and triggers a motor representation, no further practical reasoning about how to act is needed from the agent and it would be irrational for her to so reason. So where two agents each intend that they, these two agents, move an object or drive a car over the cliff, and where these intentions are appropriately related to the agents' awareness of their abilities and trigger motor representations of matching outcomes, the only rational way for their actions to proceed is by parallel planning. This is why neither Ayesha's nor Beatrice's initial intention that they drive over the cliff could have been appropriately related to awareness of action possibilities and triggered a motor representation of a matching outcome. Had this been the case, their further deliberation about how to proceed would imply that they are irrational or confused; but, by stipulation, they are not. So the notion of parallel planning, even where it involves appeal to a broad notion of planning, can, after

¹⁷ Two outcomes *match* in a particular context just if, in that context, either the occurrence of the first outcome would normally constitute or cause, at least partially, the occurrence of the second outcome or vice versa.

all, be used to give sufficient conditions for alpha collective agency.

9. Conclusion

My question was, Which planning mechanisms enable agents to coordinate their actions, and what if anything do these tell us about the nature of collective agency? In this chapter I have contrasted two planning mechanisms (see Section 5 on page 9). The first, due to Bratman (1992; 2014), is interconnected planning involving a structure of intentions and knowledge states. In such interconnected planning, you have intentions concerning my intentions, and I likewise; your plans feature facts about mine, and conversely. So interconnected planning is meta-planning: each agent's plan concerns not only facts about her environment and the objects in it but also facts about the other agents' plans. The second mechanism is parallel planning as implemented by motor processes: in each agent there is a single plan concerning all of the agents' actions. In parallel planning, no agent need have intentions about any other agent's intentions or represent facts about any other agent's plans: rather, each plans the others' actions as well as her own. This allows agents to use ordinary planning mechanisms to coordinate with others (as illustrated in Section 5).

Reflection on the sorts of contrast case that are sometimes used to get a pre-theoretic fix on a notion of collective agency indicates that interconnected planning is not sufficient for collective agency unless much simpler conditions are also sufficient (see Section 4 on page 7). Or, if we allow that there are multiple kinds or degrees of collective agency, reflection on these contrast cases indicates that interconnected planning is not sufficient for alpha collective agency, that is, for the form or degree of collective agency that we aimed to characterise. It seems that interconnected planning can't be what collective agency at bottom consists in because agents can have interconnected plans while thinking of each other's intentions only as opportunities to exploit and constraints to work around, and so without conceiving of themselves as exercising collective agency. It may be that, contra Bratman's proposal (outlined in Section 3 on page 5), no amount of forming intentions about others' intentions and acquiring knowledge of such intentions is sufficient, all by itself, for alpha collective agency.

This motivates considering the possibility that we can understand something about collective agency, or some forms of it at least, by invoking parallel planning. While it was perhaps initially tempting to suppose that parallel planning always involves irrationality or confusion, reflection on discoveries about motor representation reveals that parallel planning can occur in good cases too (see Section 8 on page 16). But what grounds might there be to think that invoking parallel planning will succeed where invoking in-

terconnected planning has failed? Parallel planning involves explicitly or implicitly adopting a perspective that allows each agent to make sense of the idea that her own and others' actions are all elements in a single plan or all parts of the exercise of a single ability (see Section 6 on page 12). So in parallel planning, the agents' actions appear to the agents themselves to have a kind of practical unity not unlike the kind of practical unity that multiple actions of a single agent sometimes have. And this is just what we need for collective agency. In some cases, it may be that parallel planning is what distinguishes exercises of (alpha) collective agency from parallel exercises of merely individual agency.

Does this imply that there is no role for interconnected planning at all? One requirement for parallel planning to support exercises of collective agency is that the agents involved have, or eventually end up with, non-accidentally matching plans. This matching can sometimes be achieved thanks to a combination of similarities in the agents' planning abilities, environmental constraints and experimentation. But there will, of course, be many situations in which these factors are insufficient to yield non-accidentally matching plans. Perhaps, then, interconnected planning matters in part because it enables agents to non-accidentally make matching plans in some situations. On this view, interconnected planning is not constitutive of collective agency but it does extend the range of cases in which agents can coordinate their actions in ways necessary to successfully exercise collective agency.

References

- Ambrosini, E., M. Costantini, and C. Sinigaglia (2011). Grasping with the eyes. *Journal of Neurophysiology*.
- Ambrosini, E., C. Sinigaglia, and M. Costantini (2012). Tie my hands, tie my eyes. *Journal of experimental psychology. Human perception and performance* 38(2), 263–266.
- Bekkering, H., A. Wohlschläger, and M. Gattis (2000). Imitation of gestures in children is goal-directed. *The Quarterly Journal of Experimental Psychology A* 53(1), 153–164.
- Brass, M., H. Bekkering, A. Wohlschläger, and W. Prinz (2000). Compatibility between observed and executed finger movements: Comparing symbolic, spatial, and imitative cues. *Brain and Cognition* 44(2), 124–143.
- Bratman, M. E. (1992). Shared cooperative activity. *The Philosophical Review* 101(2), 327–341.

- Bratman, M. E. (1997). I intend that we J. In R. Tuomela and G. Holmstrom-Hintikka (Eds.), *Contemporary Action Theory, Volume 2: Social Action*. Dordrecht: Kluwer. Reprinted in Bratman, M. (1999) *Faces of Intention*. Cambridge: Cambridge University Press (pp. 142-161).
- Bratman, M. E. (2014). *Shared Agency: A Planning Theory of Acting Together*. Oxford: Oxford University Press.
- Butterfill, S. A. and C. Sinigaglia (2014). Intention and motor representation in purposive action. *Philosophy and Phenomenological Research* 88(1), 119–145.
- Chisholm, R. (1966). Freedom and action. In K. Lehrer (Ed.), *Freedom and Determinism*, pp. 11–44. New York: Random House.
- Costantini, M., E. Ambrosini, P. Cardellicchio, and C. Sinigaglia (2013). How your hand drives my eyes. *Social Cognitive and Affective Neuroscience*.
- Costantini, M., E. Ambrosini, and C. Sinigaglia (2012). Does how I look at what you're doing depend on what I'm doing? *Acta Psychologica* 141(2), 199–204.
- Craighero, L., A. Bello, L. Fadiga, and G. Rizzolatti (2002). Hand action preparation influences the responses to hand pictures. *Neuropsychologia* 40(5), 492–502.
- Flanagan, J. R. and R. S. Johansson (2003). Action plans used in action observation. *Nature* 424(6950), 769–771.
- Frankfurt, H. (1971). Freedom of the will and the concept of a person. *The Journal of Philosophy* 68(1), 5–20.
- Gallotti, M. (2011). A naturalistic argument for the irreducibility of collective intentionality. *Philosophy of the Social Sciences*.
- Gallotti, M. and C. D. Frith (2013). Social cognition in the we-mode. *Trends in Cognitive Sciences* 17(4), 160–165.
- Gilbert, M. (2010). Collective action. In T. O'Connor and C. Sandis (Eds.), *A Companion to the Philosophy of Action*, pp. 67–73. Oxford: Blackwell.
- Gilbert, M. P. (1990). Walking together: A paradigmatic social phenomenon. *Midwest Studies in Philosophy* 15, 1–14.
- Gold, N. and R. Sugden (2007). Collective intentions and team agency. *Journal of Philosophy* 104(3), 109–137.

- Grafton, S. T. and A. F. Hamilton (2007). Evidence for a distributed hierarchy of action representation in the brain. *Human Movement Science* 26(4), 590–616.
- Hamilton, A. F. d. C. and S. T. Grafton (2008). Action outcomes are represented in human inferior frontoparietal cortex. *Cerebral Cortex* 18(5), 1160–1168.
- Jeannerod, M. (2006). *Motor Cognition: What Actions Tell the Self*. Oxford University Press.
- Khoury, C. (1992). *Thelma and Louise*. London: Pathé Entertainment.
- Kilner, J., Y. Paulignan, and S. Blakemore (2003). An interference effect of observed biological movement on action. *Current Biology* 13(6), 522–525.
- Kourtis, D., N. Sebanz, and G. Knoblich (2013). Predictive representation of other people’s actions in joint action planning: An EEG study. *Social Neuroscience* 8(1), 31–42.
- Kutz, C. (2000). Acting together. *Philosophy and Phenomenological Research* 61(1), 1–31.
- Laurence, B. (2011). An anscombian approach to collective action. In *Essays on Anscombe’s Intention*. Harvard University Press.
- Linnebo, Ø. (2005). Plural quantification. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Spring 2005 Edition)*.
- Loehr, J. D., D. Kourtis, C. Vesper, N. Sebanz, G. Knoblich, et al. (2013). Monitoring individual and joint action outcomes in duet music performance. *Journal of cognitive neuroscience forthcoming*.
- Loehr, J. D. and C. Palmer (2011). Temporal coordination between performing musicians. *The Quarterly Journal of Experimental Psychology* 64(11), 2153–2167.
- Ludwig, K. (2007). Collective intentional behavior from the standpoint of semantics. *Nous* 41(3), 355–393.
- Ménoret, M., L. Varnet, R. Fargier, A. Cheylus, A. Curie, V. des Portes, T. Nazir, and P. Y. (2013). Neural correlates of non-verbal social interactions: A dual-egg study. *Neuropsychologia in press*.
- Meyer, M., S. Hunnius, M. v. Elk, F. v. Ede, and H. Bekkering (2011). Joint action modulates motor system involvement during action observation in 3-year-olds. *Experimental Brain Research* 211(3-4), 581–592.

- Meyer, M., R. P. R. D. van der Wel, and S. Hunnius (2013). Higher-order action planning for individual and joint object manipulations. *Experimental Brain Research* 225(4), 579–588.
- Miall, R. and D. Wolpert (1996). Forward models for physiological motor control. *Neural Networks* 9(8), 1265–1279.
- Miller, S. (2001). *Social action: a teleological account*. Cambridge University Press.
- Novembre, G., L. F. Ticini, S. Schütz-Bosbach, and P. E. Keller (2013). Motor simulation and the coordination of self and other in real-time joint action. *Social Cognitive and Affective Neuroscience* forthcoming.
- Pacherie, E. (2008). The phenomenology of action: A conceptual framework. *Cognition* 107(1), 179–217.
- Pears, D. (1971). Two problems about reasons for actions. In A. M. R. Binkley, R. Bronaugh (Ed.), *Agent, Action and Reason*, pp. 128–153. Oxford: Oxford University Press.
- Pietroski, P. M. (1998). Actions, adjuncts, and agency. *Mind* 107(425), 73–111.
- Ramsey, R. and A. F. d. C. Hamilton (2010). Understanding actors and object-goals in the human brain. *Neuroimage* 50(3), 1142–1147.
- Rizzolatti, G. and C. Sinigaglia (2008). *Mirrors in the Brain: How Our Minds Share Actions, Emotions*. Oxford University Press.
- Rizzolatti, G. and C. Sinigaglia (2010). The functional role of the parieto-frontal mirror circuit: interpretations and misinterpretations. *Nature Reviews: Neuroscience* 11(4), 264–274.
- Rosenbaum, D. A., K. M. Chapman, M. Weigelt, D. J. Weiss, and R. van der Wel (2012). Cognition, action, and object manipulation. *Psychological Bulletin* 138(5), 924–946.
- Roth, A. S. (2010). Shared agency. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Fall 2001 Edition)*.
- Schmid, H. B. (2009). *Plural action: Essays in philosophy and social science*, Volume 58. Springer.
- Searle, J. R. (1990). Collective intentions and actions. In P. Cohen, J. Morgan, and M. Pollack (Eds.), *Intentions in Communication*, pp. 90–105. Cambridge: Cambridge University Press. Reprinted in Searle, J. R. (2002) *Consciousness and Language*. Cambridge: Cambridge University Press (pp. 90–105).

- Seemann, A. (2009). Why we did it: An anscombian account of collective action. *International Journal of Philosophical Studies* 17(5), 637–655.
- Smith, T. H. (2011). Playing one's part. *Review of Philosophy and Psychology*.
- Tarantino, Q. (1992). *Reservoir Dogs*. New York: Miramax Films.
- Tollefsen, D. (2005). Let's pretend: Children and joint action. *Philosophy of the Social Sciences* 35(75), 74–97.
- Tsai, J. C.-C., N. Sebanz, and G. Knoblich (2011, January). The GROOP effect: Groups mimic group actions. *Cognition* 118(1), 135–140.
- Tuomela, R. (2005). We-Intentions revisited. *Philosophical Studies* 125(3), 327–369.
- Tuomela, R. and K. Miller (1988). We-intentions. *Philosophical Studies* 53(3), 367–389.
- Vesper, C., R. P. R. D. van der Wel, G. Knoblich, and N. Sebanz (2013). Are you ready to jump? predictive mechanisms in interpersonal coordination. *Journal of Experimental Psychology: Human Perception and Performance* 39(1), 48–61.
- Watson, G. (1975). Free agency. *The Journal of Philosophy* 72(8), 205–220.
- Wolpert, D. M., Z. Ghahramani, and M. Jordan (1995). An internal model for sensorimotor integration. *Science* 269(5232), 1880–1882.
- Zhang, W. and D. A. Rosenbaum (2007). Planning for manual positioning: the end-state comfort effect for manual abduction–adduction. *Experimental Brain Research* 184(3), 383–389.