

Joint Action: A Minimalist Approach

Stephen A. Butterfill
<s.butterfill@warwick.ac.uk>

1. Introduction

There are phenomena, call them *joint actions*, paradigm cases of which are held to involve two people painting a house together (Bratman 1992), lifting a heavy sofa together (Velleman 1997), preparing a hollandaise sauce together (Searle 1990), going to Chicago together (Kutz 2000), and walking together (Gilbert 1990). In developmental psychology paradigm cases of joint action include two people tidying up the toys together (Behne, Carpenter and Tomasello 2005), cooperatively pulling handles in sequence to make a dog-puppet sing (Brownell, Ramani and Zerwas 2006), bouncing a block on a large trampoline together (Tomasello and Carpenter 2007), and pretending to row a boat together. Other paradigm cases from research in cognitive psychology include two people lifting a two-handled basket (Knoblich and Sebanz 2008), putting a stick through a ring (Ramenzoni et al. 2011), and swinging their legs in phase (Schmidt and Richardson 2008, p. 284). What feature or features distinguish joint actions such as these from events involving multiple agents who are merely acting in parallel?

This is a useful question to pursue because joint action raises a tangle of scientific and philosophical questions. Psychologically and neuroscientifically we want to know which mechanisms make it possible (Sebanz, Bekkering and Knoblich 2006; Vesper et al. 2010; Sacheli et al. 2015). Developmentally we want to know when joint action emerges, what it presupposes and whether it might somehow facilitate socio-cognitive, pragmatic or symbolic development (Hughes and Leekam 2004; Brownell, Ramani and Zerwas 2006; Moll and Tomasello 2007). Phenomenologically we want to characterise what (if

anything) is special about experiences of action and agency when collective agency is involved (Pacherie, forthcoming 2010). Metaphysically we want to know what kinds of entities and structures are implied by the existence of joint action (Searle 1994; Helm 2008). And normatively we want to know what kinds of commitments (if any) are entailed by joint action and how these commitments arise (Roth 2004; Gilbert 2013). For investigating any of these questions it would be useful to understand which feature or features distinguishes joint actions from events involving multiple agents who are merely acting in parallel.

Could coordination be this distinguishing feature? Compare two sisters cycling to school together in a way that sisters characteristically cycle together with two strangers who, some way behind the sisters, happen to be cycling the same route side-by-side. Both pairs of cyclists need to coordinate their actions in order to avoid colliding, but only the former is a joint action. So the bare fact that actions are coordinated, even very tightly coordinated in ways that require expertise, cannot be what distinguishes joint action from events involving multiple agents who are merely acting in parallel.

Another initially tempting idea is that common effects distinguish joint actions. When members of a flash mob in Café Central respond to a pre-arranged cue by noisily opening their newspapers, they perform a joint action with a common effect. But when someone not part of the mob just happens to noisily open her newspaper in response to the same cue, her action is not part of any joint action. Yet her action together with the actions of the flash mob members have a common effect in startling the people around them. So what distinguishes joint actions from events involving multiple agents who merely act in parallel can't be just that joint actions have common effects.

At this point it is natural to appeal to intention. Perhaps joint action occurs when there is an act-type, ϕ , such that each of several agents intends that they, these agents, ϕ together, and their actions are appropriately related to these intentions. Does the appeal to togetherness make this proposal circular? Not as long as we understand 'together' only in the sense in which the three legs of a tripod can support a flask together. Can people ever have intentions which concern not only their own actions but also each others'? Yes, at least if whether each person's intention persists depends on the others' intentions persisting (Bratman 1993).

Appealing to intention seems to take us further than the first two ideas (coordination and common effects). Consider the cycling sisters again. Cycling together in the way that sisters characteristically cycle together plausibly involves each sister intending that they, the two sisters cycle to school together. But nothing like this is characteristic of strangers who just happen to be

cycling side-by-side—neither is likely to have intentions concerning whether the other gets to school. So have we already identified what distinguishes joint action?

Not yet. Imagine two sisters who, getting off an aeroplane, tacitly agree to exact revenge on the unruly mob of drunken hens behind them by positioning themselves so as to block the aisle together. This is a joint action. Meanwhile on another plane, two strangers happen to be so configured that they are collectively blocking the aisle. The first passenger correctly anticipates that the other passenger, who is a complete stranger, will not be moving from her current position for some time. This creates an opportunity for the first passenger: she intends that they, she and the stranger, block the aisle together. And, as it happens, the second passenger's thoughts mirror the first's. So the feature under consideration as distinctive of joint action is present: each passenger is acting on her intention that they, the two passengers, block the aisle. But the contrast between this case and the sisters exacting revenge suggests that these passengers are not taking part in a joint action—at least, their's is not the kind of joint action associated with the paradigm cases mentioned at the start of this chapter. Apparently, then, our being involved in a joint action can't be a matter only of there being something such that we each intend that we, you and I, do it together. What are we missing?

It's just here that, in philosophy at least, things get a little wild. Attempts to provide the missing ingredient in characterising joint action include introducing novel kinds of intentions (Searle 1990) or modes (Gallotti and Frith 2013), novel kinds of agents (Helm 2008), and novel kinds of reasoning (Gold and Sugden 2007). Others suggest embedding intentions in special kinds of commitment (Gilbert 2013), or creating special nested structures of intention and common knowledge (Bratman 2014). Perhaps some or all of these innovations are in some way useful. But are they really needed just to understand how joint actions differ from events involving multiple agents who are merely acting in parallel?

The dominant assumption is that they are. To illustrate, consider Gilbert's position. According to her, all joint action involves shared intention, and our having a shared intention that we ϕ involves our being jointly committed to emulate a single body with an intention that it ϕ . In order to create the joint commitment necessary for us to have a shared intention, we must each openly express readiness to participate in this commitment. Further, it must be common knowledge among us that we each express such readiness. Her account thus implies that, in order for us to share a smile or carry a two-handled picnic basket together, each of us must know that the other is ready to form a joint commitment to emulate a single body with an intention to

share a smile or carry the basket (Gilbert 2013, p. 334). Few would agree with Gilbert that exactly this nesting of mental states and commitments is necessary for joint action. (This is no reflection on Gilbert—few philosophers would agree in any detail with anyone’s view on what is necessary for joint action.) But many do follow Gilbert in thinking that distinguishing the kind of joint action involved in the examples given at the start of this chapter requires either comparably complex nested structures or novel ingredients.

By contrast Bratman has recently observed, in effect, that introducing such complex structures or novel ingredients is not obviously needed just to distinguish joint action from events involving multiple agents who are merely acting in parallel (Bratman 2014, p. 105). For all anyone has yet shown, there may be a way of capturing what is distinctive of the kind of joint actions mentioned at the start of this chapter without invoking novel ingredients or structures. This chapter explores the possibility that there is, with the twofold aim of constructing a minimalist theoretical framework for understanding at least simple forms joint action and illuminating the nature of the intentions and commitments involved in the most sophisticated forms of joint action.

2. A Minimalist Approach: First Step

Taking a minimalist approach means finding a simplest possible starting point, adding ingredients only as needed, and avoiding as far as possible ingredients which would require the agents to have abilities additional to those already required.¹ What determines whether an additional ingredient is needed? The aim is to distinguish joint actions like those mentioned at the start of the chapter from events involving multiple agents who are merely acting in parallel.

As a promising starting point, consider a claim from Ludwig’s semantic analysis:

‘A *joint action* is an event with two or more agents’ (Ludwig 2007, p. 366).

To illustrate, suppose two hunters each attack a deer. Neither attack was individually fatal but together they were deadly. In this case the hunters are agents of the killing of the deer, so the event counts as a joint action on Ludwig’s proposal.

¹ An approach might be minimalist without being conceptually conservative in Bratman’s sense (see Bratman 2014, pp. 14–5), and conversely. Minimalism concerns what a theory demands of the agents whose joint actions it characterises; conceptual conservatism is about what a theory demands of theoreticians.

To fully understand Ludwig's proposal we need to understand what it is for an individual to be among the agents of an arbitrary event and not just an action. This can be done in terms of a notion of grounding which I adapt from a discussion of action by Pietroski (1998). Pietroski identified a simple and elegant way of generalising from the idea that an individual can be the agent of an action to the idea that an individual can be the agent of a larger event. (His account does require a minor correction, but this is not relevant here.) This can be generalised to allow for any number of agents. Let us stipulate that events D_1, \dots, D_n ground E , if: D_1, \dots, D_n and E occur; D_1, \dots, D_n are each (perhaps improper) parts of E ; and every event that is a proper part of E but does not overlap D_1, \dots, D_n is caused by some or all of D_1, \dots, D_n . Then let us say that for an individual to be *among the agents of an event* is for there to be actions A_1, \dots, A_n which ground this event, where the individual is an agent of some (one or more) of these actions. To illustrate, consider the hunters again. Let THE EPISODE be an event comprising only the hunter's actions, the deer's death and the events causally linking these. Since, for each hunter, there is a set of events including this hunter's attacking which ground THE EPISODE, we can conclude that THE EPISODE is a joint action on Ludwig's proposed definition.

This definition is too broad. To see why, first recall our premise that one requirement on any account of joint action is this: it should distinguish joint actions like those mentioned at the start of this chapter from events involving multiple agents who are merely acting in parallel. Now consider two ways of elaborating the story about the hunters. In one they are best friends who have set out together with the aim of killing this deer, and they are exhibiting many features associated with paradigm cases of joint action. In the other elaboration, the hunters are bitter rivals completely unaware of each other's presence. In fact, were either to have suspected the other was present, she would have abandoned the deer in order to target her rival. In both elaborations, Ludwig's proposal entails that THE EPISODE is a joint action. But whereas the 'best friends' elaboration resembles paradigm cases of joint action, the bitter rivals are merely acting in parallel. By itself, Ludwig's attractively simple proposal is insufficient.²

² Should we have considered the idea that a joint action is an *action* (rather than an event) with two or more agents? This question raises several issues beyond the scope of the present chapter. The short answer is no, because primitive actions (whether bodily movements or tryings) are 'all the actions there are' (Davidson 1971, p. 59), and in many paradigm cases of joint action there are clearly no primitive actions with multiple agents. In painting a house, walking together or lifting a two-handed basket we each move only our own bodies directly. The notion of a joint action as an action with two or more agents is therefore too narrow relative to our aim of theorising about a range of cases taken to

What is missing from this first attempt to capture joint action? Many joint actions are goal-directed in the sense that, among all their actual outcomes, there is an outcome to which they are directed. Perhaps we can make progress by integrating goal-directedness into our theoretical framework.

3. Goal Directed Joint Action

A *goal* of some actions is an outcome to which they are directed. A *goal-state* is a state of an agent which specifies an outcome and which is, or could be, related to the agent's actions in such a way that these actions are directed to the outcome represented. Intentions are perhaps the most familiar kind of goal-state; goals are things like the shooting of a deer or the carrying of a basket.

Note that, confusingly, the term 'goal' is sometimes used for what I am calling goal-states. Whatever terminology is used, it is essential to distinguish goals from goal-states. From the fact that an action is directed to a particular goal it does not follow that the agent of the action has a goal-state representing this goal (Bratman 1984). It doesn't even follow that the agent has any goal-states at all if, as some have argued, it is possible to understand what it is for at least some actions to be directed to goals without appeal to goal-states (e.g. Taylor 1964; Bennett 1976).

In almost all of the events offered as paradigm cases of joint action in philosophy and psychology, there is a single goal to which all the agents' actions are directed. To illustrate, return once more to the sisters cycling to school together. Cycling together in the way sisters characteristically cycle together plausibly involves there being a single goal to which the sisters' actions are all directed. Perhaps the goal is the arrival of the two sisters at their school. By contrast, there is plausibly no goal to which all the actions of the two strangers cycling side-by-side are directed.

Note that distinguishing the sisters from the strangers in this way depends on distinguishing *same goal* from *same type of goal*. The two strangers' actions may have goals of the same type; for instance, each stranger's actions may be directed to her own arrival somewhere. But this does not amount to there being a single goal to which the strangers' actions are all directed. After all, if one stranger falls into a hole and is taken to hospital in an ambulance, the other's actions may still succeed relative to the goal of reaching her destination.

be paradigmatic joint actions. (This is not to say that no actions have two or more agents; see Blomberg 2011.)

Consider a second attempt to characterise joint action, which builds on the earlier attempt by incorporating goals:

A joint action is an event with two or more agents where there is a single goal to which all the actions grounding that event are directed.

Although narrower than the earlier attempt (which was introduced in section 2), this attempt is still not narrow enough. To see why, consider the deer hunters again. The best friends' actions are clearly directed to a single goal, that of killing the deer. But so are the bitter rivals' actions. So this second attempt to characterise joint action still fails to distinguish the kind of joint action characteristic of best friends hunting together from the parallel but merely individual actions of the bitter rivals who end up hunting the same deer. More is needed.

4. Collective Goals

Some predicates can be interpreted either distributively or collectively (see Linnebo 2005 for an introduction). Consider the statement, 'The injections saved her life.' This could be true in virtue of her receiving several injections on different occasions, each of which saved her life. In this case, the injections saving her life is just a matter of each injection individually saving her life; this is the distributive interpretation. But the statement is also true if she was given two injections on a single occasion where each injection was necessary but not sufficient to save her life. In this case the injections saving her life is not, or not just, a matter of each injection individually saving her life; this is the collective interpretation. The difference between distributive and collective interpretations is clearly substantial, for on the distributive interpretation the statement can only be true if her life has been saved more than once, whereas the truth of the collective interpretation requires only one life-threatening situation.

Just as some injections can be collectively life-saving, so some actions can be collectively directed to a goal. For some actions to have a *collective goal* is for there to be a single outcome to which the actions are directed where this is not, or not just, a matter of each of the actions being individually directed to that outcome. To illustrate, there is a sense in which some of the actions of swarming bees are directed to finding a nest and this is not, or not just, a matter of each bee's actions being individually directed to finding a nest. So finding a nest is a collective goal of the bees' actions. Likewise, when two

people use a rope and pulleys to lift a heavy block between them, lifting the block will typically be a collective goal of their actions.

In virtue of what could any actions ever have a collective goal? One possibility involves coordination. To illustrate, return to the deer hunters who are best friends again. Knowing the difficulty of killing a deer, they coordinate so that when one of them startles it, the other is positioned along the deer's escape route. The coordination ensures that their actions being directed to the killing of the deer is not just a matter of each action's being directed to this outcome, and so entails that killing the deer is a collective goal of their actions. In general, for an outcome to be a collective goal of some actions it is sufficient that all the actions are coordinated as a means to bringing about this outcome

It is natural to assume that the hunters' coordination is a consequence of their intentions (at least if they are human rather than, say, lycine). This may make it tempting to assume that what ultimately determines which actions have which collective goals is not coordination but intention. But this temptation should be resisted. The coordination needed for multiple individuals' behaviours to have a collective goal can be provided by entirely non-psychological mechanisms, as popular findings about bees (Seeley 2010) and ants (Hölldobler and Wilson 2009, pp. 178-83, 206-21) indicate. It is also likely that not all coordination between humans involves intention (see, for example, Repp and Su 2013, §3, and Knoblich, Butterfill and Sebanz 2011 on emergent coordination). Where some actions are coordinated in order to bring about an outcome, the actions are collectively directed to that outcome in virtue of being so coordinated.

Staying with the deer hunters for a moment, note that appealing to collective goals enables us to distinguish the bitter rivals who are merely acting in parallel from the best friends who are performing a joint action. The actions of the bitter rivals are directed to a single outcome, that of killing a particular deer; but this is just a matter of each hunter's actions being directed to this outcome. So killing this deer is not a collective goal of the bitter rivals' actions. By contrast, the best friends' actions are (by stipulation) coordinated in a way that would normally increase the probability of their killing the deer. So their actions do have a collective goal.

Could appealing to collective goals eventually enable us to distinguish more generally between joint actions and events involving multiple agents who are merely acting in parallel? Consider invoking collective goals to narrow the previous attempt to characterise joint action:

A joint action is an event with two or more agents where the actions grounding that event have a collective goal.

This attempted characterisation is a step forward. But it does not enable us to distinguish all joint actions from events involving multiple agents who are merely acting in parallel. To see why, return to the contrast between two ways of blocking the aisle of an aeroplane (which was mentioned in section 1). First, two sisters tacitly agree to block the aisle by positioning themselves side-by-side; this is a paradigm case of joint action. Second, two strangers also succeed in blocking the aisle although neither guesses the other's intention and each knows of the other only that she is unlikely to move from her current position. The strangers appear not to be involved in a joint action because from either stranger's point of view the other is merely a conveniently stationary and sufficiently bulky object. Yet the strangers' actions are clearly coordinated: it is no coincidence that they are positioned as they are. Further, their actions are coordinated in order to block the aisle: each seeks to position herself relative to the other in such a way as to prevent passengers getting past. This implies that the strangers' actions have a collective goal. Once again our attempted characterisation of joint action appears to be too broad. What to do?

5. Some Intentions Specify Collective Goals

We have seen that distinguishing two ways of blocking the aisle, one involving sisters performing a joint action and the other involving strangers performing parallel but merely individual actions, cannot be captured either by a simple appeal to intention (see section 1) or by a simple appeal to collective goals (see section 4). But perhaps we can distinguish the joint action from its merely parallel counterpart by invoking intentions which specify collective goals.

How might intentions specify collective goals? An obvious possibility is for collective goals to feature in what agents intend. For instance, the sisters blocking the aisle might in principle each intend that they, the two sisters, perform actions which have the collective goal of blocking the aisle and succeed relative to this intention. If we were to suppose that intentions about collective goals are a characteristic feature of joint action, then we would have a simple way of distinguishing joint actions from parallel actions like those of the strangers blocking the aisle. After all, the strangers blocking the aisle cannot rationally intend that their actions have blocking the aisle as a collective goal (because each believes the other is not performing actions directed to this goal).

But are intentions about collective goals a characteristic feature of joint action? Whereas joint action is a pervasive feature of every day life, it would be surprising to discover that intentions about collective goals are

similarly pervasive. After all, having the intention about the collective goal appears to require understanding quite generally what collective goals are. This motivates considering whether intentions might specify collective goals other than by virtue of being part of what agents intend.

Start by thinking about ordinary, individual action. Consider events in which a vertical stroke is made with a pencil. On some occasions this is realised by an action directed to the goal of making a vertical stroke with a pencil. On other occasions it is realised by something not directed to any such goal; perhaps you are jolted while holding a pencil over some paper. It is a familiar idea that, often at least, someone who intends to make a stroke with a pencil has an intention that would not be fulfilled if, say, she was jolted while holding the pencil in such a way as to make a vertical stroke. Instead, fulfilling this particular intention requires performing an action directed to the goal of making a vertical stroke with a pencil. Should we infer that what she intends is not simply to make a vertical stroke but to perform an action directed to the goal of making a vertical stroke and to succeed relative to this goal? No. Either intending to make a vertical stroke just is intending the thing about the goal, or else the thing about the goal enters the satisfaction conditions of the intention without being part of what she intends. Either way, having an intention that locks onto a particular type of goal-directed action such as drawing a stroke with a pencil does not require having intentions (or any thoughts) about actions and goals generally.

A related point holds for collective goals. Consider events in which two or more agents move a fallen tree that is blocking a road. Some such events involve actions for which moving the tree is a collective goal. Other such events involve actions which are merely individually directed to moving the fallen tree (the tree is so big and the storm so intense that the several agents are unaware of each other until after having moved the tree). Suppose Ayesha intends that she and Beatrice move a fallen tree blocking their path, and that fulfilling this particular intention requires moving the tree to be a collective goal of Ayesha's and Beatrice's actions. As in the related case of ordinary, individual goals, this does not require that collective goals feature in what Ayesha intends, or not in a way that goes beyond her intending that they, she and Beatrice, move the tree. Having an intention that locks onto a particular type of event involving a collective goal, such as the moving of fallen tree, does not require having intentions about collective goals generally.

Why is this relevant? Earlier (in section 1) we briefly considered the simple idea that joint action occurs when there is an act-type, ϕ , such that each of several agents intends that they, these agents, ϕ , and their actions are appropriately related to these intentions. This simple idea seemed inadequate

for distinguishing joint actions from events involving multiple agents who are merely acting in parallel. Why? Because the strangers blocking the aisle of an aeroplane are not involved in a joint action but do each intend that they, the two strangers, block the aisle. But we are now in a position to improve on the simple idea. Note that the strangers' intentions do not require for their fulfilment that they, the strangers, perform actions with the collective goal of blocking the aisle. Indeed, by stipulation each stranger falsely believes that the other is not performing actions directed to blocking the aisle. So what each stranger believes is straightforwardly incompatible with blocking the aisle being a collective goal of their actions. This suggests that we can improve on the simple idea by requiring that the relevant intentions must require for their fulfilment actions with a corresponding collective goal. And we have just seen that imposing such a requirement would not entail tacitly imposing an implausible further requirement on abilities to think about collective goals generally.

Consider a further attempt to characterise joint action, one which builds on earlier efforts by requiring intentions that specify collective goals:

A joint action is an event with two or more agents where:

1. the actions grounding that event are appropriately related to intentions on the part of each agent that they, these agents, ϕ together; and
2. each intention requires for its fulfilment that all the actions have a collective goal concerning ϕ -ing.

Call this attempted characterisation the *Flat Intention View*. It improves on the earlier attempts insofar as, unlike them, it distinguishes all the joint actions so far considered from the counterpart events involving multiple agents who are merely acting in parallel.

The Flat Intention View is so-called because it relies on a single, unnested intention where some other approaches require intentions nested in intentions. To motivate invoking nested intentions, Bratman (1992, p. 333; 2014, p. 49) introduces a pair of contrasting cases in which two people intend that they, the two people, go to New York City. One case involves the sort of situation best friends planning a holiday might be in. The other involves two members of competing gangs. Each gangster intends that they, the two gangsters go to New York City by means of her 'throwing the other into the trunk of the car and driving to NYC' (Bratman 2014, p. 49). Bratman takes this contrast between how intentions to go to NYC typically unfold in friendly situations and how intentions to go there 'in the mafia sense' unfold to motivate the view that distinguishing these situations requires not just intentions

but intentions about intentions. But the Flat Intention View provides a way of distinguishing the friendly case from the mafia case without introducing higher-order intentions. If one gangster succeeds in bundling the other into the trunk of their car, the two gangsters will not perform actions with the collective goal of their going to NYC. For this reason, neither gangster can rationally and knowingly have both an intention whose fulfilment requires them to perform such actions and also an intention to bundle the other into the trunk. So the Flat Intention View excludes going to NYC ‘in the mafia sense’ without any need for intentions nested in intentions. While Bratman’s aims extend far beyond the issues considered in this chapter, the success of the Flat Intention View in distinguishing the ordinary case from the mafia case suggests that further arguments are needed to show that characterising joint action requires nested intentions.

The Flat Intention View can be contrasted with views which invoke intentions with novel kinds of subjects, namely plural subjects (see, for example Schmid 2009), novel kinds of attitudes such as ‘we-intentions’ as Searle (1990) characterises them, or novel kinds of commitments such as Gilbert (2013)’s joint commitments. By contrast, the Flat Intention View follows Bratman’s view in requiring neither a novel kind of subject nor a novel kind of attitude. On the Flat Intention View, the only special feature of the intentions associated with joint action is that their fulfilment conditions involve collective goals.

Does the Flat Intention View need supplementing with requirements to the effect that each agent knows or believes something about the other’s (or others’) intentions? A positive answer would be consistent with the arguments of this chapter, and there is no obvious obstacle to supplementing the Flat Intention View in some such way. But, as far as I know, philosophers have yet to show that requirements about knowledge or belief are necessary (compare Blomberg 2015). It may be true, of course, that successful joint action often requires knowledge of others’ intentions. But it may be possible to explain this fact (if it is a fact) by appeal to rational requirements on having intentions such as those specified by the Flat Intention View. After all, you cannot rationally intend that, say, we, you and I, make a pizza together unless many background requirements are met. These background requirements may include requirements on your beliefs or knowledge about my intentions. So while successfully performing a joint action may often require that each agent knows or believes something about the other’s (or others’) intentions, this could be a consequence of the knowledge or belief being a requirement on the rationality of having the intentions specified by the Flat Intention View. As things stand, then, it appears we have yet to see sufficient reasons

to complicate the Flat Intention View by adding requirements concerning knowledge of, or beliefs about, of others' intentions.

Should we conclude that our latest attempt, the Flat Intention View, enables us to distinguish generally between joint actions and events involving multiple agents who are merely acting in parallel? This seems implausible. There may be further contrasts between joint actions and merely parallel actions which the Flat Intention View fails to discriminate. This would motivate further narrowing the Flat Intention View by adding additional ingredients. But at this point there is a more pressing reason not to simply accept the Flat Intention View: it is too narrow. The Flat Intention View does not enable us to distinguish between joint actions and merely parallel actions when the agents lack intentions concerning the joint actions. This matters not just if joint actions can occur in agents such as bees or ants which probably lack intentions altogether. It also matters because, arguably, some forms of joint actions in humans need not involve intention. Since the aim of the minimalist approach is to distinguish joint actions from events involving multiple agents who are merely acting in parallel, the Flat Intention View cannot be where the story ends.

6. The Agents' Perspective

Recall our earlier attempt to characterise joint action using collective goals only and no intentions (from section 4):

A joint action is an event with two or more agents where the actions grounding that event have a collective goal.

This attempt is adequate for distinguishing joint actions from their merely parallel counterparts in many cases where intentions are not considered, such as the flash mob (see section 1) and the deer hunters (see section 2). It was only a contrast case involving intentions, namely that of blocking the aisle of an aeroplane (see section 4), which showed us that this attempt is inadequate for distinguishing all joint actions from events involving multiple agents who are merely acting in parallel. So the grounds for finding this earlier attempt to characterise joint action inadequate do not motivate the further claim that all joint action involves having intentions.

What then can we conclude from the failure of this earlier attempt to characterise joint action? The case of strangers blocking the aisle of the aeroplane involves features that create a tension. The fact that their actions are coordinated in such a way as to be collectively directed to the outcome of blocking the aisle indicates that they are performing a joint action. But

the fact that the truth of the beliefs informing their intentions would require their actions to lack a collective goal indicates that they are not performing a joint action. These conflicting indicators should not be weighed against each other because they involve different perspectives. Seen from the outside there appears to be a joint action, whereas from the perspective of either agent there does not.

It turns out that a minimalist approach to joint action needs to be pluralist. To capture joint action from the agents' perspective, we sometimes need to invoke intentions or other goal-states which specify collective goals. And to capture joint action while being neutral on the agents' perspective, we need to avoid invoking any such intentions or other goal-states. We must therefore recognise that there are multiple contrasts and multiple kinds of joint action.

This conclusion is unlikely to be controversial. Our question about what distinguishes joint actions from events involving multiple agents who are merely acting in parallel is analogous to one about ordinary, individual action. On ordinary, individual action, Davidson (1971) and others have asked, Which feature or features distinguish actions from events otherwise involving an agent? The variety of different agents, and the variety of control structures within relatively complex agents such as humans, suggests that answering this question will probably involve recognising that there are multiple kinds of action. The claim that there are multiple kinds of joint action is not more controversial.

Recognising that there are multiple kinds of joint action suggests a simple answer to Schweikard and Schmid (2013)'s challenge to views, such as Bratman's (2014) and the Flat Intention View (see section 5), on which facts about individual agents' intentions and other mental states are taken to explain how joint actions differ from events involving multiple agents who are merely acting in parallel. Schweikard and Schmid suggest that such views must presuppose the very distinction to be explained and ask, apparently rhetorically, '[H]ow can an individual refer to a joint activity without the jointness [...] already being in place?' On the Flat Intention View, the intentions that constitute one kind of 'jointness' refer to joint actions characterised without intention—these are the joint actions we attempted to capture with the proposal that a joint action is an event with two or more agents where the actions grounding that event have a collective goal.

7. Conclusion

What feature or features distinguish joint actions from events involving multiple agents who are merely acting in parallel? This chapter introduced a

minimalist approach to answering this question, one which involves finding a simplest possible starting point, adding ingredients only as needed and avoiding as far as possible adding ingredients which would require the agents to have abilities additional to those already required. This yielded attempted characterisations for two kinds of joint action. One characterisation involves collective goals only, and no intentions or other mental states:

A joint action is an event with two or more agents where the actions grounding that event have a collective goal.

The consideration that actions can have a collective goal even though each agent conceives of herself as merely exploiting a conveniently stationary and sufficiently bulky object (see section 4) motivated introducing a further attempt, labelled the Flat Intention View:

A joint action is an event with two or more agents where:

1. the actions grounding that event are appropriately related to intentions on the part of each agent that they, these agents, ϕ together; and
2. each intention requires for its fulfilment that all the actions have a collective goal concerning ϕ -ing.

These attempted characterisations raise more questions than they answer. For each attempted characterisation of joint action, what are the counter-examples to its adequacy and what additional ingredients are needed to refine it? Offering a pair of characterisations involves commitment to a form of pluralism about joint action (see section 6). But is it pluralist enough? The pair of characterisations assume there is a single distinction between those joint actions that do, and those that do not, essentially involve intention. Are further distinctions necessary? For instance, are there multiple kinds of joint action which essentially involve intention?

Even without answers to these basic questions, a minimalist approach to joint action is clearly useful in at least two related ways. First, the minimalist approach yields ingredients such as the notion of a collective goal that can be used in specifying the contents of intentions and other states. Some philosophers tacitly assume that characterising the kind or kinds of joint action which essentially involve intention is possible without reflection on kinds of joint action which do not essentially involve intention. For instance, Bratman (2014, p. 46) relies exclusively on activities which are ‘neutral with respect to shared intentionality’ in constructing an account of a kind of joint action that involves intention. The category of activities which are neutral with respect to shared intentionality is extremely broad. It includes actions

with collective goals as well as actions with no collective goal (such as those intended by the strangers blocking the aisle in section 4, and Bratman's gangsters who each intend that they go to NYC by means of her bundling the other into the trunk from section 5). On Bratman's construction, then, the agents of a joint action essentially involving intentions might as well have no conception of the possibility of joint actions not involving intentions; from their point of view, it might as well be that all joint actions essentially involve intentions to do things together. But, as we saw in section 5, there may be intentions whose fulfilment requires not merely activities which are neutral with respect to shared intentionality but, more demandingly, actions with collective goals. This suggests that kinds of joint action which essentially involve intention may be constructed on top of, and perhaps even emerge from, kinds of joint action which do not essentially involve intention.

The minimalist approach is also useful as a way of testing claims that particular ingredients are needed to characterise joint action and related notions. To illustrate, in section 5 we saw that reflection on the Flat Intention View creates difficulties for an attempt to argue that characterising a kind of joint action which essentially involves intentions requires not merely intentions but intentions about intentions. This line of argument generalises. Many researchers agree that all joint action requires shared intention: for them, the central problem in giving an account of joint action is how to characterise shared intention.³ But since the Flat Intention View enables us to distinguish many paradigm cases of joint action from events involving multiple agents who are merely acting in parallel, it is at least unclear that shared intention is needed in characterising even forms of joint action that essentially involve intention.

Is joint action fundamentally a matter of collective goals and the ordinary, individual intentions which specify collective goals? This seems unlikely. More ingredients are surely needed. And challenges to the further development of a minimalist approach surely lie ahead. Even so, as an alternative to currently dominant attempts to characterise joint action by appeal to complex nested structures or conceptually novel ingredients, a minimalist approach is promising.

³ See, for instance, Gilbert (2006, p. 5): 'I take a collective [joint] action to involve a collective [shared] intention'; and Alonso (2009, pp. 444-5): 'the key property of joint action lies in its internal component ... in the participants' having a "collective" or "shared" intention.' See also Tomasello (2008, p. 181) and Carpenter (2009, p. 381). Clearly these authors' claims should be restricted to the kind, or kinds, of joint action that essentially involve intention.

References

- Alonso, Facundo M. 2009. 'Shared Intention, Reliance, and Interpersonal Obligations'. *Ethics* 119, no. 3 (April): 444–475. doi:10.1086/599984.
- Behne, Tanya, Malinda Carpenter and Michael Tomasello. 2005. 'One-year-olds comprehend the communicative intentions behind gestures in a hiding game'. *Developmental Science* 8 (6): 492–499.
- Bennett, Jonathan. 1976. *Linguistic Behaviour*. Cambridge: Cambridge University Press.
- Blomberg, Olle. 2011. 'Socially Extended Intentions-in-Action'. *Review of Philosophy and Psychology* 2 (May): 335–353. doi:10.1007/s13164-011-0054-3.
- . 2015. 'Common Knowledge and Reductionism about Shared Agency'. *Australasian Journal of Philosophy* (July): 1–12. doi:10.1080/00048402.2015.1055581.
- Bratman, Michael E. 1984. 'Two Faces of Intention'. *The Philosophical Review* 93 (3): 375–405.
- . 1992. 'Shared Cooperative Activity'. *The Philosophical Review* 101 (2): 327–341.
- . 1993. 'Shared Intention'. *Ethics* 104:97–113.
- . 2014. *Shared Agency: A Planning Theory of Acting Together*. Oxford: Oxford University Press.
- Brownell, Celia A., Geetha B. Ramani and Stephanie Zerwas. 2006. 'Becoming a social partner with peers: cooperation and social understanding in one- and two-year-olds'. *Child Development* 77 (4): 803–21.
- Carpenter, Malinda. 2009. 'Just how joint is joint action in infancy?' *Topics in Cognitive Science* 1 (2): 380–392.
- Davidson, Donald. 1971. 'Agency'. In *Agent, Action, and Reason*, edited by Robert Binkley, Richard Bronaugh and Ausonia Marras. Reprinted in Davidson, D. (1980) *Essays on Actions and Events*. Oxford: Oxford University Press. Toronto: University of Toronto Press.
- Gallotti, Mattia, and Chris D. Frith. 2013. 'Social cognition in the we-mode'. *Trends in Cognitive Sciences* 17 (4): 160–165. doi:10.1016/j.tics.2013.02.002.

- Gilbert, Margaret P. 1990. 'Walking Together: A Paradigmatic Social Phenomenon'. *Midwest Studies in Philosophy* 15:1–14.
- . 2006. 'Rationality in Collective Action'. *Philosophy of the Social Sciences* 36 (1): 3–17.
- . 2013. *Joint Commitment: How We Make the Social World*. Oxford: Oxford University Press.
- Gold, Natalie, and Robert Sugden. 2007. 'Collective Intentions and Team Agency'. *Journal of Philosophy* 104 (3): 109–137.
- Helm, Bennett W. 2008. 'Plural Agents'. *Nous* 42 (1): 17–49. doi:10.1111/j.1468-0068.2007.00672.x.
- Hölldobler, B., and E.O. Wilson. 2009. *The superorganism: the beauty, elegance, and strangeness of insect societies*. W.W. Norton.
- Hughes, Claire, and Sue Leekam. 2004. 'What are the Links Between Theory of Mind and Social Relations? Review, Reflections and New Directions for Studies of Typical and Atypical Development'. *Social Development* 13 (4): 590–619.
- Knoblich, Guenther, Stephen Butterfill and Natalie Sebanz. 2011. 'Psychological research on joint action: Theory and data'. In *Psychology of Learning and Motivation*, edited by Brian Ross, vol. 51. Academic Press.
- Knoblich, Günther, and Natalie Sebanz. 2008. 'Evolving Intentions for Social Interaction: from Entrainment to Joint Action'. *Philosophical Transactions of the Royal Society B* 363:2021–2031.
- Kutz, Christopher. 2000. 'Acting Together'. *Philosophy and Phenomenological Research* 61 (1): 1–31.
- Linnebo, Øystein. 2005. 'Plural Quantification'. In *The Stanford Encyclopedia of Philosophy (Spring 2005 Edition)*, edited by Edward N. Zalta. Stanford, CA: CSLI.
- Ludwig, Kirk. 2007. 'Collective Intentional Behavior from the Standpoint of Semantics'. *Nous* 41 (3): 355–393. doi:10.1111/j.1468-0068.2007.00652.x.
- Moll, Henrike, and Michael Tomasello. 2007. 'Cooperation and human cognition: the Vygotskian Intelligence Hypothesis'. *Philosophical Transactions of the Royal Society B* 362 (1480): 639–648.

- Pacherie, Elisabeth. forthcoming 2010. 'The Phenomenology of Joint Action: Self-Agency vs. Joint-Agency'. In *Joint Action*, edited by Axel Seeman. MIT Press.
- Pietroski, Paul M. 1998. 'Actions, Adjuncts, and Agency'. *Mind*, New Series, 107 (425): 73–111.
- Ramenzoni, Verónica C., Tehran J. Davis, Michael A. Riley, Kevin Shockley and Aimee A. Baker. 2011. 'Joint action in a cooperative precision task: nested processes of intrapersonal and interpersonal coordination'. *Experimental Brain Research* 211 (3-4): 447–57. doi:10.1007/s00221-011-2653-8.
- Repp, Bruno H., and Yi-Huang Su. 2013. 'Sensorimotor synchronization: A review of recent research (2006–2012)'. *Psychonomic Bulletin & Review* 20, no. 3 (February): 403–452. doi:10.3758/s13423-012-0371-2.
- Roth, Abraham Sesshu. 2004. 'Shared Agency and Contralateral Commitments'. *The Philosophical Review* 113 (3): 359–410.
- Sacheli, Lucia M., Matteo Candidi, Vanessa Era and Salvatore M. Aglioti. 2015. 'Causative role of left aIPS in coding shared goals during human-avatar complementary joint actions'. *Nature Communications* 6 (July). doi:10.1038/ncomms8544.
- Schmid, Hans Bernhard. 2009. *Plural action: Essays in philosophy and social science*. Vol. 58. Dordrecht: Springer.
- Schmidt, Richard C., and Michael J. Richardson. 2008. 'Dynamics of Interpersonal Coordination'. Chap. XV in *Coordination: Neural, Behavioral and Social Dynamics*, edited by Armin Fuchs and Viktor K. Jirsa, 280–308. Berlin, Heidelberg: Springer.
- Schweikard, David P., and Hans Bernhard Schmid. 2013. 'Collective Intentionality'. In *The Stanford Encyclopedia of Philosophy*, Summer 2013, edited by Edward N. Zalta.
- Searle, John R. 1990. 'Collective Intentions and Actions'. In *Intentions in Communication*, edited by P. Cohen, J. Morgan and M.E. Pollack, 90–105. Reprinted in Searle, J. R. (2002) *Consciousness and Language*. Cambridge: Cambridge University Press (pp. 90–105). Cambridge: Cambridge University Press.
- . 1994. *The Construction of Social Reality*. New York: The Free Press.

- Sebanz, Natalie, Harold Bekkering and Guenther Knoblich. 2006. 'Joint Action: Bodies and Mind Moving Together'. *Trends in Cognitive Sciences* 10 (2): 70–76.
- Seeley, T.D. 2010. *Honeybee Democracy*. Princeton University Press. Princeton University Press.
- Taylor, Charles. 1964. *The Explanation of Behaviour*. London: Routledge.
- Tomasello, M. 2008. *Origins of human communication*. The MIT Press.
- Tomasello, Michael, and Malinda Carpenter. 2007. 'Shared Intentionality'. *Developmental Science* 10 (1): 121–5.
- Velleman, David. 1997. 'How To Share An Intention'. *Philosophy and Phenomenological Research* 57 (1): 29–50.
- Vesper, Cordula, Stephen Butterfill, Günther Knoblich and Natalie Sebanz. 2010. 'A minimal architecture for joint action'. *Neural Networks* 23 (8-9): 998–1003. doi:10.1016/j.neunet.2010.06.002.