

Towards a Mechanistically Neutral Account of Acting Jointly: The Notion of a Collective Goal

STEPHEN A. BUTTERFILL

Department of Philosophy, University of Warwick, UK
s.butterfill@warwick.ac.uk

CORRADO SINIGAGLIA 

Dipartimento di Filosofia, Università degli Studi di Milano, Italy
corrado.sinigaglia@unimi.it

1. Introduction

Many of the things we do are, or could be, done with others. Mundane examples favoured by philosophers include 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). These examples are supposed to be paradigm cases of a phenomenon, or class of phenomena, we shall call *acting jointly*, although a variety of labels have been used.¹ To delimit which cases are of interest, philosophers have also contrasted things they take to be cases of acting jointly with things they take to be cases of people merely acting in parallel with each other. For instance, when members of a flash mob in the Central Cafe respond to a pre-arranged cue by noisily opening their newspapers in order to create a salient marker for the opening of their performance, they are held to be acting jointly. But when someone not part of the mob just happens to noisily

¹ Labels include 'joint action' (Brooks 1981; Sebanz, Bekkering, and Knoblich 2006; Knoblich, Butterfill, and Sebanz 2011; Tollefsen 2005; Pettit and Schweikard 2006; Carpenter 2009; Pacherie 2010; Brownell 2011; Sacheli, Arcangeli, and Paulesu 2018; Meyer, Wel, and Hunnius 2013), 'social action' (Tuomela and Miller 1985), 'collective action' (Searle 1990; Gilbert 2010), 'joint activity' (Baier 1997), 'acting together' (Tuomela 2000), 'shared intentional activity' (Bratman 1997), 'plural action' (Schmid 2008), 'joint agency' (Pacherie 2013), 'small scale shared agency' (Bratman 2014), 'intentional joint action' (Blomberg 2016), 'collective intentional behaviour' (Ludwig 2016), and 'collective activity' (Longworth 2019). We leave open whether these are all labels for a single phenomenon or whether different researchers are targeting different things. As we use 'acting jointly', the term applies to everything any of these labels applies to.

open her newspaper in response to the same cue, she is held to be acting in parallel with the others but merely individually.² To give another example, two former members of the mob are held to be acting jointly when they later walk to the metro station together. But two people who merely happen to be walking to the metro station side by side among a crowd of people heading that way are held to be acting in parallel but merely individually. (This example is adapted from [Bratman \(2014, pp. 5–6\)](#), who borrowed it from [Gilbert \(1990\)](#).) These paradigm and contrast cases invite the question: what features distinguish acting jointly from acting in parallel but merely individually?

Although some philosophers may question some of these claims about paradigm cases and contrast cases,³ we shall follow those cited above in assuming, provisionally,⁴ that reflection on the examples and contrasts can ensure sufficient agreement on the topic that it makes sense to ask which features distinguish acting jointly from acting in parallel but merely individually. This assumption is consistent with neutrality both on whether there is just one distinction or several to be drawn, and also on which of the above claims about paradigm cases and contrast cases are correct.

A standard strategy for distinguishing acting jointly from acting in parallel but merely individually involves invoking states of the agents who are acting jointly, often dubbed ‘we-’, ‘shared’ or ‘collective intentions’.⁵ The idea, very roughly, is that acting jointly (or one kind of acting jointly if there are several) is distinguished from acting in parallel but merely individually by these states playing some particular role. The central task, on this strategy, is to characterise the states and their explanatory relation to acting jointly (compare

² See [Searle \(1990\)](#); in his example park visitors simultaneously run to a shelter, in one case as part of dancing together and in another case because of a storm. Compare [Pears \(1971\)](#) who uses contrast cases to argue that whether something is an ordinary, individual action depends on its antecedents.

³ For views which imply or may motivate dissent, see for example [Baier \(1997\)](#), [Chant \(2007\)](#), [Pettersson \(2007\)](#), and [Longworth \(2019, p. 13ff\)](#).

⁴ We propose a way to avoid relying on this assumption in §7 below.

⁵ This strategy has been pursued by a number of philosophers. One early example (although not the first) is [Tuomela and Miller \(1985\)](#); it may be that [Tuomela and Miller \(1988\)](#) and [Searle’s \(1990\)](#) response initiated contemporary debate. This is not to say that no philosophers have taken an alternative line. [Pettersson \(2007, p. 138\)](#), for instance, attempts to explicate the distinction between acting jointly and acting in parallel but merely individually ‘in terms of dispositions and causal agency’. See also [Chant \(2007\)](#) for another alternative line.

(Bratman 2014, p. 10)). There is, however, considerable debate about the nature of the states on which this distinction hinges. Some hold that the states in question involve a novel attitude (Searle 1990; Gallotti and Frith 2013). Others have explored the notion that the primary distinguishing feature of these states is not the kind of attitude involved but rather the kind of subject, which is plural (Helm 2008). Or they may differ from ordinary intentions in involving distinctive obligations or commitments to others (Gilbert 1992; Roth 2004).⁶ Or perhaps the most fundamental distinguishing mark of these states is the way they arise, namely through team reasoning (Gold and Sugden 2007; Pacherie 2012). Opposing all such views, Bratman (1992, 2014) argues that the distinctive states, which he calls ‘shared intentions’, can be realised by multiple ordinary individual intentions and other attitudes whose contents interlock in a distinctive way.⁷ Bratman’s approach has inspired a family of accounts along broadly these lines, including Asarnow (2020), Blomberg (2016), Ludwig (2007, 2016) and Tollefsen (2005).

How are we to determine when any two of these accounts should be regarded as competing attempts to characterise a single phenomenon and when they should be regarded as compatible attempts to characterise different phenomena? And how are we to single out, from among all of these accounts, those which are correct? The growing number and increasing diversity of accounts make urgent these twin problems. It may be that they can be solved. But we believe that there is sufficient doubt to motivate investigating an alternative to the standard strategy, one which makes available a distinctive way of singling out correct accounts. And our aim in this paper is to do just this.

⁶ Because Gilbert can be interpreted as characterising states of an agent in terms of commitments (compare Gilbert (2013, p. 10): ‘I take [...] acting together to involve collective intentions—understood in terms of joint commitment’), she does provide materials that can be used by a proponent of the standard strategy. However, it may be more accurate to interpret her as offering an alternative to the standard strategy, one which does not involve invoking any states of the agents at all but only certain normative facts (see further §7). Related points could be made about others cited in this paragraph (particularly Gold and Sugden (2007) and Bratman (2014)). What unites them is that they do provide materials for proponents of the standard strategy to identify states of the agents in distinguishing acting jointly.

⁷ In Bratman (1992), the ordinary individual intentions and other attitudes were offered as jointly sufficient and individually necessary conditions; the retreat to sufficient conditions only occurs in Bratman (1997, pp. 143–44): ‘for all that I have said, shared intention might be multiply realizable’. Because Bratman distinguishes two tasks in characterising acting jointly—specifying a functional role and identifying states which realise it—his work provides materials both for proponents of the standard strategy and also for proponents of an alternative strategy. We return to this in §7.

In investigating an alternative strategy, we are not promising to adjudicate among accounts which follow the standard strategy. However, we will show how those accounts can be recast in line with the alternative strategy and their main insights adjudicated. To anticipate, one consequence of this will be that some accounts which appeared incompatible from the point of view of the standard strategy turn out not to be incompatible when recast.

2. An alternative strategy

Proponents of what we call the standard strategy differ in important ways but have one thing in common: in attempting to distinguish acting jointly from acting in parallel but merely individually, they are invoking states of the agents. On the standard strategy, there is no way of drawing this distinction without being committed to a claim about which states, or structures of states, cause the actions agents perform when they act jointly.

To see that this is not the only possible strategy, consider philosophical accounts of ordinary, individual action. We might ask, following Davidson (1971), what distinguishes those events which are actions from events which merely happen to an agent? In answering this question many have, like Davidson, focused on intentional actions only. Some philosophers answer the question by saying—to put it very roughly and incompletely—that intentional actions are those caused in a certain way by beliefs, desires, intentions or other kinds of state.⁸ This is an example of what we shall call a *mechanistically committed* answer: that is, an answer which involves making commitments concerning which states, or structures of states, cause intentional actions. Although this is the dominant approach, and one which has provided a framework for much philosophy of action (Brand (1984) is a nice illustration), it is not the only way philosophers have attempted to answer the question. An intentional action may be characterised, very roughly and incompletely, as an action that happens because its agent has certain reasons for bringing an outcome about, or at least for attempting to do so. This is an example of what

⁸ Examples include Bach (1978) and Dretske (1988). Although Bratman (1987) is sometimes said to be a proponent of this view (Schlosser 2019), we take Bratman to be primarily concerned with developing a theory of intention—and although this involves investigating relations between intentional action and intentions, it does not obviously require commitment to characterising the former in terms of the latter.

we shall call a *mechanistically neutral* answer: that is, one which does not involve making commitments concerning which states, or structures of states, cause intentional actions.⁹

Proponents of a mechanistically neutral characterisation need not deny that intentional actions are caused by intentions or other states of agents. They merely insist on separating two questions:

- (1) What distinguishes intentional actions from things which merely happen to an agent (and from nonintentional actions, if there are any)?
- (2) Which states cause intentional actions?

All sides can agree that fully understanding action requires answering both questions (among others).¹⁰ But whereas a mechanistically committed characterisation answers the first question in a way that involves answering, partially or wholly, the second, a mechanistically neutral characterisation answers the first question in a way that does not.¹¹ What makes an approach mechanistically neutral is not the pattern of answers given to the questions but the recognition that the answers are to an interesting degree theoretically independent.

To see when a mechanistically neutral characterisation could be useful, consider Bratman's position as an example. He allows that actions can be intentional 'even though [the agent] has no distinctive attitude of intending' (Bratman 1987, p. 132), and even though the agent lacks the capacity to form intentions altogether (Bratman 2000, p. 51). This view follows from two claims: first, intentions are distinct from any combination of beliefs and desires; and second, beliefs and

⁹ But is this second answer really mechanistically neutral? According to Alvarez (2016), philosophers sometimes conceive of 'motivating and explanatory reasons ... as mental states of agents'. Providing a mechanistically neutral characterisation of intentional action requires a mechanistically neutral characterisation of reasons, too. If, as we suppose, such a characterisation exists, then there are mechanistically neutral attempts to characterise intentional action.

¹⁰ Of course, there are philosophers who might deny that the second question bears on any philosophical questions about action (Ginet (1990), for example). Proponents of a mechanistically neutral characterisation are free to oppose such philosophers on the grounds that the answers to the two questions are mutually constraining (see §7 below). They may therefore oppose both causal theorists (whose strategy is to answer the first question by answering the second) and non-causal theorists (who deny that the second question is relevant).

¹¹ Note that the possibility of characterising A in terms which do not mention B does not in general imply that it is possible for there to be As without corresponding Bs. Proponents of a mechanistically neutral approach may therefore accept that intentional actions are caused by intentions and could not be caused in some other way (to borrow, with our thanks, some words from a critic).

desires alone may, in certain cases, determine what an agent intentionally does. Of course, this second claim would make no sense if we answered the first question above by saying that intentional actions are things caused by intentions. But Bratman's position clearly makes sense if we rely on a mechanistically neutral characterisation of intentional action. If we anchor the notion of intentional action by saying that intentional actions are things which happen because an agent has reasons, we can then coherently postulate variety in the states and processes which cause intentional actions.

Until now we have considered approaches which focus on intentional actions only. But this is not the only coherent approach. Another possibility is to first characterise a notion of purposive action (invoking goals) and then elaborate on this in characterising intentional action (invoking reasons concerning the goals). In this case there is also a divide between mechanistically committed and mechanistically neutral characterisations. A purposive action is an action directed to one or more goals. Some take goals to be states of agents (Austin and Vancouver 1996, p. 338). Any characterisation relying on such a notion of goal is mechanistically committed: purposive actions are characterised in terms of states which cause them. By contrast, others have used 'goal' as a label for those outcomes to which an action is directed (for example, Wilson, Shpall, and Pineros Glasscock (2016)) and have offered, or relied on, mechanistically neutral characterisations of purposive action (Bennett 1976, p. 42; Butterfill 2001; Gergely et al. 1995; Csibra and Gergely 2007; Wright 1976).

How does any of this bear on our question about which features distinguish acting jointly from acting in parallel but merely individually? Consideration of ordinary, individual action suggests two things: first, that it may also be possible to provide a mechanistically neutral characterisation of acting jointly; and, second, that it is coherent not to focus exclusively on a notion of acting jointly where acting jointly is intentional. We can therefore separate two questions:

- (1) What distinguishes acting jointly from acting in parallel but merely individually?
- (2) Which states cause the actions agents perform when they act jointly?

We shall extend our terminology and distinguish *mechanistically committed* from *mechanistically neutral* answers to the first question

according to whether they do, or do not, involve making commitments concerning the second question.

Giving a mechanistically neutral answer is consistent with denying that the second question bears on the first at all, perhaps because you see the first as a normative question that is quite separate from any concern about causes. But, as we explain in §7, giving a mechanistically neutral answer is also consistent with allowing that the questions are mutually constraining.

In what follows, we explore the prospects for an alternative strategy which deviates from the standard strategy in aiming to provide a mechanistically neutral characterisation. The initial focus is not on joint counterparts of intentional action but on the notion of acting jointly (allowing that intentional cases of acting jointly may eventually turn out to be in some respect more fundamental). We shall argue that the alternative yields a characterisation of acting jointly which, although probably incomplete, meets criteria standardly used in evaluating other accounts. We shall also argue that the alternative strategy offers an advantage in solving the twin problems from §1: the problem of singling out correct accounts, and the problem of determining when two accounts are competing attempts to characterise a single phenomenon rather than compatible attempts to characterise different phenomena.

3. Collective versus distributive

This and the following sections sketch a mechanistically neutral answer to the question: what distinguishes acting jointly from acting in parallel but merely individually? We focus initially on the purposive aspect of acting jointly, and will only later (in §8) turn to its intentional aspect. Just as the notion of a goal is central to any account of ordinary, individual purposive action, so our central notion will be that of a collective goal. In order to introduce this notion (in §5), we first need to clarify how we will use the term ‘collective’ (in this section) and what it means for an action to have a goal (in §4).

Consider these sentences:

- (1) The tiny leaves fell from the tree.
- (2) The tiny leaves blocked the drain.

The first sentence is naturally read *distributively*; that is, as specifying something that each leaf did individually. Perhaps first one leaf fell, then another fell. But the second sentence is naturally read *collectively*.

No one leaf blocked the drain; rather the blocking was something that the leaves accomplished together. For the sentence to be true on this collective reading, the tiny leaves' blocking the drain cannot be, or cannot only be, a matter of each leaf blocking the drain.¹²

Now consider an example involving actions and their outcomes:

(3) Those thoughtless actions blocked the drain.

This sentence can be read in at least two ways, distributively or collectively. We can read it distributively as concerning a sequence of actions done over a period of time, each of which blocked the drain. In this case, the truth of the sentence is just a matter of the same type of outcome, namely blocking the drain, being an outcome of each action. Alternatively, we can read it as concerning several actions which have this outcome collectively—perhaps a number of people dropped cigarette butts into the drain more or less simultaneously. In this case the outcome, blocking the drain, is not necessarily an outcome of any of the individual actions, but it is an outcome of all of them taken together (or, depending on one's views about the semantics of plural quantification, an outcome of one thing that somehow bundles together the actions). This is the collective reading.

Note that the difference is not merely terminological. To see this, consider how many times the drain must have been blocked. On the distributive reading it was blocked at least as many times as there were actions. On the collective reading it was not necessarily blocked more than once. So the difference between collective and distributive is not just a matter of words: it concerns how the actions and outcomes are related.

4. Goals

Our aim (in the next section) will be to show that there are collective readings not only of sentences about the actual outcomes of actions but also of sentences about the outcomes to which actions are directed—that is, about the goals of actions.¹³ To this end we first need to say what it is for an action to be directed to an outcome.

¹² This informal contrast between collective and distributive readings is linked to a debate about the logic of plural quantification; see [Linnebo \(2005\)](#) for an overview of that debate.

¹³ As we noted earlier (in §2), the term 'goal' has been used, coherently, both to label states of agents and also to label those outcomes to which an action is directed. We shall (by stipulation) use 'goal' to label outcomes.

The challenge is to cash out the metaphor of directedness. One familiar response involves intention. To a first approximation, the directedness of an action to an outcome consists in an agent acting on an intention where the intention plays a role in coordinating (and maybe in planning) the agent's actions in a way that would normally facilitate the outcome's occurrence. This is one coherent way of thinking about directedness. But it may not be the whole story about the goals of actions. For some have argued that actions can be directed to outcomes in virtue of states other than intentions, including desires (for example, [Bratman \(2000\)](#)) and motor representations (for example, [Butterfill and Sinigaglia \(2014\)](#)). Note that regardless of which states of an agent (intentions or others) are invoked, we can use the same basic pattern: directedness is grounded by states which coordinate the action in a way that would normally facilitate the outcome's occurrence.¹⁴

Is it possible to characterise directedness without appealing to any states of an agent at all? [Bennett \(1976, p. 61\)](#) suggests that the directedness of an action to an outcome is a matter of the agent acting because she is 'so structured and situated' as to do things which increase the probability of the outcome occurring. We take Bennett to be introducing two ideas. First, directedness is not always only a matter of states of the agent but can also involve her environment and history. Second, directedness can be characterised independently of any particular mechanism. Borrowing these ideas we can draw a parallel with the core idea about intention mentioned above: the directedness of an action to an outcome consists in there being a state, structure or situation which plays a role in coordinating the agent's actions in a way that would normally facilitate the outcome's occurrence.¹⁵ The state, structure or situation may be intention, habit, biological function or other behaviour-organizing circumstance connecting the agent's actions to the outcome (or any combination of these). The key to characterising directedness is not any specific state, structure or situation but the role of those in linking actions to outcomes.

¹⁴ [Dickinson \(2016, p. 177\)](#) offers a more sophisticated analysis of directedness which is also neutral on which states are involved.

¹⁵ Perhaps directedness is better characterised in terms of an agent's relation to an outcome (as in [Bennett \(1976\)](#)). We explain below why we do not offer this as our primary candidate (our position is neutral).

This proposal about directedness is, to put it politely, theoretically modest. It captures nothing but the bare structure of how actions and goals relate. (And even that incompletely, for goals can be partially ordered by the means–ends relation and actions by the part–whole relation; directedness needs to be understood, more generally, as a relation of structures of actions to structures of outcomes.) But theoretical modesty is an advantage in one respect: the proposal leaves open for discovery questions about what it is in virtue of which actions are directed to outcomes.

This openness is valuable because we can make discoveries about the goals of particular actions without yet knowing anything about the kind of state, structure or situation in virtue of which the actions are directed to the goal. For example, consider that ant behaviour is routinely and uncontroversially characterized in terms of goals:

The ants protect their fungal cultivar from pathogens and parasites, provide the fungus with a constant source of nutrients, and aid in its growth and dispersal. [...] To promote the initial degradation of plant biomass this material is masticated, mixed with ant fecal droplets, and inoculated with fungal mycelia. (Scott et al. 2010, e9922)

Facts about the specific goals of the ants' actions are important discoveries. But for the most part, these discoveries were made independently of much insight into what grounds the directedness of the actions to the goals. Indeed, discoveries about goals are often foundational for understanding mechanisms. To illustrate with another species, knowing that the goal of certain actions is hunting enabled the discovery that portid spiders use information about routes to prey even after that information is no longer available in their environment (Jackson and Cross 2011, pp. 118–21), suggesting that representation may be involved.

One quirk of our proposal is that we have characterised directedness as linking actions, not agents, to outcomes. Our concern here is to avoid a commitment. Consider:

The goal of her action is to compost the leaf cuttings.

and:

Her goal is to compost the leaf cuttings.

To some, the first may sound like an awkward paraphrase of the second. But others might deny that the first implies the second—

perhaps on the grounds that the first can be true of ants but ants cannot have goals; or perhaps on the grounds that actions can be subagential in something like the way that (on some views) mental states can be subpersonal. To avoid such controversies, we characterise directedness as a relation between an action and an outcome. However, what follows could be adapted to work with directedness as a relation between agents and outcomes.

5. Collective goals

Having clarified our use of ‘collective’ (in §3) and characterised goals (in the previous section) we are in a position to see how the distinction between collective and distributive readings appears to apply to sentences concerning the goals of actions. Consider the sentence:

(4) Those actions had the goal of blocking the drain.

Whereas the previous sentence (in §3) was about causal relations between actions and outcomes, this sentence concerns teleological relations. We claim that, like the previous sentence, this sentence has both distributive and collective readings. On the distributive reading, each of the actions had the goal of blocking the drain. This would fit a sequence of events in which someone attempted to block the drain first by covering it with a metal sheet and then, after this failed to block it, removing the metal sheet and pouring cement powder into it. On the collective reading, by contrast, the actions’ having the goal of blocking the drain was not, or not only, a matter of each of the actions having that goal. This would fit an episode in which someone maliciously and patiently blocks a drain by dropping just one cigarette butt into it each day, counting on their accumulation to do its work.

The collective reading also appears to be possible in many mundane cases. Imagine kneading some dough. This involves a sequence of folding and stretching actions, and perhaps adding some flour as you go. The point of all this is to get the dough into a condition such that it will later rise in a particular way. So we might say, of the sequence, that those actions had the goal of getting the dough into a state where it will rise nicely. If we are describing an ordinary baking activity, this statement will be true on a collective reading.

So far we have considered actions with just one agent. But we can also read sentences about goals collectively when there are two or more agents involved. After all, the examples of blocking the drain

and kneading the dough could just as well have involved two agents rather than just one.

Where a sentence about some actions being directed to a single outcome is true on the collective reading, we stipulate, as a matter of terminology, that the actions are *collectively directed* to that outcome and that the outcome is a *collective goal* of the actions.

We have explicated the notion of a collective goal only negatively by saying that it is not, or not only, a matter of each action individually being directed to the outcome. What is it a matter of? One answer to this question might involve using one of the existing theories about states of agents to further explicate the notion of a collective goal. This would be incompatible with our aim of being mechanistically neutral, of course. But there are also other kinds of structure that can ground collective goals. In the case of honey bees foraging (Leadbeater and Chittka 2005) or ants farming fungus (see §4), what underpins collective goals may be behavioural patterns, scent marks and other signals which ensure coordination; and of course much the same may be true in some situations involving humans too. In characterising collective goals we therefore avoid invoking any particular states or structures, just as we did in characterising goals (in §4). Indeed, our proposal about goals can be extended to collective goals in a natural way by switching from intra-agential to inter-agential coordination: the collective directedness of some agents' actions to an outcome consists in there being a state, structure or situation which plays a role in coordinating all the actions in a way that would normally facilitate the outcome's occurrence.

Nothing that follows depends on the details of this proposal. We require only that the collective readings of sentences about goals serve to single out, within a limited but useful range of cases, which things are collective goals.

6. The Collective Goal Account

Can collective goals assist in providing a mechanistically neutral account of how acting jointly differs from acting in parallel but merely individually? Consider the least subtle attempt to give a positive answer, which we will call *The Collective Goal Account*:

When the actions of two or more agents have a collective goal, the agents are, in performing those actions, acting jointly; otherwise they are acting in parallel but merely individually.

This account is neutral on which states of the agents make it the case that their actions have a collective goal. It is consistent with, but neutral on, the further claim that it is only ever in virtue of the agents' shared intentions that actions have collective goals.

How does this account fare with respect to the examples and contrast cases we mentioned at the start of this paper? In giving the examples standardly offered which we mentioned in opening this paper—painting a house together, lifting a heavy sofa together, and the rest (see §1)—philosophers are distinguishing which thing is occurring by specifying a goal (the lifting of the sofa, for instance). And these examples do indeed appear to be cases where the specified goal is typically a collective goal. It can hardly be very controversial, therefore, that one feature of standard examples is the presence of a collective goal.¹⁶

Next consider the two contrast cases we mentioned at the start of this paper (in §1). The first contrasted the actions of members of a flash mob who noisily crack open their newspapers in order to create a salient marker for the start of their performance with the actions of the flash mob plus those of an onlooker who happens to crack open her newspaper simultaneously in response to the same cue. The actions of the flash mob would typically have a collective goal, namely the creation of the salient marker. By contrast, there is no single token outcome to which the onlooker's actions plus the flash mob members' actions are directed.¹⁷ It follows, of course, that there is no collective goal. So invoking the notion of a collective goal suffices to distinguish this contrast case.

The second contrast case was a contrast between two members of the flash mob later walking to the metro station together, who (by stipulation) are acting jointly, and two people who happen to be walking to the station side by side among a crowd heading that way. The success of the flash mob members' activity requires at least that both walk to the station. And this is just the kind of case in which the two agents' arrival at the metro station would be a collective goal

¹⁶ This is not entirely uncontroversial, of course. If Baier (1997) is right, some cases of acting jointly do not involve collective goals.

¹⁷ This point is easily overlooked, perhaps because there is clearly one type of outcome to which all the agents' actions are directed, namely the opening of some newspaper or other. But note that the onlooker's action may be entirely successful (she may crack open her own newspaper) while another agent's fails (she fumbles and drops hers). This makes it clear that the outcomes are distinct, although they are of the same type.

of their actions. By contrast, when two people just happen to be walking to the station side by side, one person's actions may succeed even though the other accidentally falls into the sewer through a hole and never reaches the station. It follows that invoking the notion of a collective goal is sufficient to distinguish this contrast case too.

Reflection on both contrast cases shows, however, that invoking the notion of a collective goal is not necessary. There is a simpler way to distinguish these contrast cases. Indeed, the presence or absence of a single outcome to which all the agents' actions are directed is already sufficient to distinguish the contrast cases. Our appeal to collective goals therefore lacks motivation. If the examples and contrast cases can be dealt with just by invoking the existence of a single outcome to which all agents' actions are directed, why bother with the additional complexity of collective goals?

An initially tempting idea is to introduce further contrast cases. For instance:

Case 1: A bear has been spotted near each of two villages on either side of a mountain. In each village there is a hunter who sets out to hunt the bear. The villagers are entirely unaware of each other. Despite tracking the bear from nearly opposite directions, the bear comes into view for both hunters at the same time. Each takes a shot at it. Neither shot is individually fatal, but sadly their combined effect kills the bear.

Case 2: Two friends from a city go hunting a bear together. They hide in a tree waiting for the bear to come into view. When it does, each takes a shot at it. Neither shot is individually fatal, but sadly their combined effect kills the bear.

As each villager justifiably regards herself as having succeeded in her project, we can be confident that the actions of the villagers are directed to a single outcome, namely the death of the bear. This implies that there is a single outcome to which all their actions are directed. So we cannot capture this contrast by appealing only to the idea that the agents' actions are all directed to a single outcome; this is a feature of both cases. Nevertheless, there does appear to be a contrast with respect to acting jointly between the villagers' actions and those of the city dwellers. We can capture this contrast by appeal to collective goals. We might say, 'The villagers' actions had the goal of killing the bear'. This sentence is true on the distributive reading, as each of them performed actions which were indeed directed to the

goal of killing the bear. But the sentence is untrue on the collective reading, for there is nothing more to their actions having the goal of killing the bear than that each villager individually performed actions directed to this outcome. By contrast, actions performed by the two city dwellers are coordinated in a way that would normally facilitate the bear's killing and therefore collectively directed to this outcome (see §5). So the Collective Goal Account is able to distinguish acting jointly in this case.

We doubt that introducing further examples and contrast cases will be sufficient, however. When we have given talks or shared drafts, there are always researchers in the audience whose intuitions differ from what we expect on just about any contrast case we have mentioned. Some awkward dissenters want to say, for instance, that the village dwellers are acting jointly no less than the city dwellers are.¹⁸ Others may dissent on the grounds that the presence of a collective goal is insufficient for the city dwellers to be acting jointly.¹⁹ Several lengthy conversations combined with reflection on the diversity of philosophers' accounts of acting jointly have gradually shifted us away from thinking of the awkward dissenters as people who just haven't tuned in. Their dissent is informative. Simply providing examples and contrast cases does not reliably give an audience a single signal to tune in to. The contrast cases on which there is little dissent do not enable us to lock on to whatever phenomena those responsible for the leading accounts of acting jointly are interested in; and the contrast cases with the potential to do this are, to a sufficiently diverse audience, themselves too controversial to put much weight on.

As far as examples and contrast cases go, we find no grounds for preferring this account over a competitor (or conversely). In the next section we shall therefore introduce another way of defending the Collective Goal Account.

¹⁸ We have also encountered awkward dissenters on other influential cases in the literature including Bratman's going to New York 'in the mafia sense' (Bratman 1992, p. 333) and Blomberg's no-common-knowledge cases (Blomberg 2016). The Collective Goal Account broadly agrees with Bratman on the former (as the victim, once locked in the trunk, does not perform any action directed to the goal) and with Blomberg on the latter (as it does not necessarily require common knowledge).

¹⁹ Blomberg (2015) provides one means of constructing contrast cases which appear to support such dissent.

7. How to defend a mechanistically neutral account of acting jointly

Our problem is how to single out correct from incorrect attempts to distinguish acting jointly from acting in parallel but merely individually. The problem arises because it is possible to construct multiple accounts—whether mechanistically committed (see §1) or mechanistically neutral (see §6)—where each account will be found by at least some people to accord with how they intuitively make the distinction. In this section we shall characterise a strategy on which this problem is actually an advantage.

As mentioned in the introduction (§1), the mechanistically neutral strategy can be found in the literature. Consider the approach of [Sebanz, Bekkering and Knoblich \(2006\)](#), who offer this influential characterisation:

Joint action can be regarded as any form of social interaction whereby two or more individuals coordinate their actions in space and time to bring about a change in the environment. ([Sebanz, Bekkering and Knoblich 2006](#), p. 70)

Although this is offered as a ‘working definition’ and is not supposed to provide deep insight (as it invokes social interaction, which is hardly less difficult to pin down than joint action), it has proven to be useful. We suggest its usefulness is due at least in part to its clearly separating the thing to be explained from proposed explanations of it. Opposing groups of researchers adopt the working definition but propose different theories about the mechanisms of joint action. The working definition provides the common ground necessary for evaluating the various proposals and understanding them as broadly compatible attempts to fill in details about the mechanisms which make joint action possible (see [Knoblich, Butterfill and Sebanz \(2011\)](#) for a review).

The working definition above was not intended to, and does not, distinguish acting jointly from acting in parallel but merely individually. This is no objection to [Sebanz, Bekkering and Knoblich \(2006\)](#), of course. But it does mean that we cannot use it for our aim of understanding this distinction. This is why we have considered the Collective Goal Account (§6).

The Collective Goal Account can be used in the same way as [Sebanz, Bekkering and Knoblich’s \(2006\)](#) working definition. That is, we can consider it as one attempt to characterise some aspect of

a phenomenon. Whether we eventually accept this attempt should depend on what we discover about mechanisms underpinning collective goals. If it turns out that there is at least one mechanism, this will increase our confidence that the Collective Goal Account is correct.

Indeed, there is already one study which explicitly sets out to establish a mechanism underpinning collective goals (della Gatta et al. 2017). These authors conclude that agents' actions can have collective goals in virtue of motor representations. This is not an isolated finding. Other studies might be interpreted as indirectly supporting the same conclusion about motor representations causing actions agents perform when their actions have collective goals (these include Baus et al. (2014); Clarke et al. (2019); Kourtis et al. (2014); Loehr and Vesper (2015); Ménoret et al. (2014); Meyer, Wel and Hunnius (2013); Novembre et al. (2014); Ramenzoni, Sebanz and Knoblich (2014); Schmitz et al. (2017); Sacheli, Arcangeli and Paulesu (2018); Sacheli et al. (2021)). Of course, fully defending—or decisively rejecting—the idea that collective goals can be used in distinguishing acting jointly from acting in parallel but merely individually would require careful analysis of these findings, and perhaps further experimental discoveries. While that is beyond the scope of this paper, we can already say that the Collective Goal Account receives some support from discoveries about mechanisms.

Should we therefore accept the Collective Goal Account? Not yet. Some may object that some essential feature is missing, such as cooperation (Searle 1990) or commitment (Gilbert 2013).

How could we discover whether the objections are justified? We need, first, to construct mechanistically neutral accounts of cooperation and commitment; then, second, to discover which states and processes causally explain the actions one or another kind of agent performs when she performs actions which are cooperative or committed in the relevant ways; and, third, to understand whether such states and processes are distinct from, or bound up with, those which enable two or more agents' actions to have collective goals. If a feature such as cooperation or commitment turns out to be enabled by the same states and processes that enable collective goals, then we have grounds for rejecting or revising the Collective Goal Account, and for regarding this feature as an essential feature of (at least one kind of) acting jointly. But if such a feature is enabled by clearly dissociable states and processes, we would need alternative grounds for regarding it as an essential feature. Given that all this is unknown at present, we cannot accept either the objections or the Collective Goal Account.

Suppose that the Collective Goal Account were vindicated against the objections. It would not follow that the account is uniquely correct, nor that notions such as cooperation and commitment have no role in characterising acting jointly. Other mechanistically neutral characterisations of acting jointly may be equally well supported. In this case we might eventually conclude that there are multiple correct ways of distinguishing acting jointly from acting in parallel but merely individually. This form of pluralism about acting jointly would be unproblematic because it is disciplined. It would be analogous to discovering that dissociable systems map onto different aspects of memory (for example, [Jacoby 1991](#)).

The approach we are describing is not limited to defending the Collective Goal Account. Here is the general idea. Start by being maximally permissive in considering any proposed mechanistically neutral account of acting jointly, requiring only that the account is a conceivably successful attempt to capture cases reasonably taken to be paradigms of acting jointly.²⁰ Next attempt to discover which states and processes might enable the features identified by these mechanistically neutral accounts to be exhibited by actual agents. Where no such states or processes can be discovered, the mechanistically neutral account may be disregarded. Where multiple accounts specify features which are all enabled by the same, or overlapping, states and processes, there is a case for treating those mechanistically neutral accounts as fragments of some larger, unified account. And where distinctive states and processes enable features specified by just one mechanistically neutral account, this indicates that the account captures one distinction between acting jointly and acting in parallel but merely individually.

This way of defending mechanistically neutral accounts may be relevant to a conflict between proponents of leading characterisations of acting jointly. Consider two well-known attempts to specify shared intention and its explanatory relation to action. [Gilbert \(2013, p. 10\)](#) does this in terms of a special kind of mutual obligation which she labels *joint commitment*, [Bratman \(2014, p. 7\)](#) in terms of interpersonal

²⁰ In saying that a case is ‘reasonably taken to be’ a paradigm of acting jointly, we mean only that one or more people have taken it to be one and would continue to do so on reflection. We do not mean to imply that it is actually a case of acting jointly. And as we intend the phrase to be understood, a ‘conceivably successful attempt’ need not capture every case that is reasonably taken to be paradigmatic: what is required is just that there would appear to be insufficient reason to regard uncaptured cases as actual cases of acting jointly if the account were true.

coordination of planning. Both present themselves as offering incompatible accounts of a single phenomenon (Gilbert 2014; Bratman 2014, chap. 5).

Suppose we wanted to address this conflict using the above way of defending mechanistically neutral accounts. We could use elements from each protagonist to construct a mechanistically neutral adaptation of their account.²¹ In Gilbert's case, this would be an account on which, roughly, the directedness of actions to outcomes involves an element of commitment. Where some agents are jointly committed to bringing an outcome about, their actions are directed to this outcome insofar as they are guided by this commitment. As long as joint commitment is understood in a way that does not involve any particular states of the agents, this could yield a mechanistically neutral characterisation. In Bratman's case, instead of using ideas about interpersonal coordination of planning to characterise shared intention, we can treat them as characterising shared agency directly. For instance, we might consider that shared agency involves two or more agents' actions being directed to an outcome where the directedness of their action to this outcome consists not only in their actions being coordinated in a way that would normally facilitate the outcome's occurrence but also in there being a state, structure or situation which plays a role in coordinating their plans and structuring their bargaining and deliberation.

These mechanistically neutral adaptations may seem barely different from Gilbert's and Bratman's original versions. Yet the difference really is substantial because it affects whether the accounts are incompatible and how correct accounts can be singled out. As adapted, there are no immediate grounds to regard Gilbert's account as incompatible with an account which, like Bratman's, centres on planning abilities; nor conversely. After all, as in the case of the Collective Goal Account, a mechanistically neutral adaptation of Gilbert would be vindicated (or not) through discovering mechanisms which actually underpin joint commitment and support acting jointly; and likewise for Bratman on interpersonal coordination of planning. Any such vindication is consistent, of course, with entirely different accounts also being vindicated. There is a range of possibilities. Mechanisms

²¹ We characterised mechanistically neutral in terms of acting jointly (in §2). As Gilbert and Bratman take the primary explanandum to be cases where two or more agents' acting jointly is intentional, we are relying on a more general notion of *mechanistically neutral* here. We discuss intentional cases of acting jointly in §8.

underpinning joint commitments could coincide with mechanisms underpinning other mechanistically neutral accounts—perhaps, for instance, there is no interpersonal coordination of planning without joint commitments.²² This might motivate combining what had initially appeared to be distinct mechanistically neutral accounts. Alternatively, it may be that distinct mechanisms are involved, which would lend support to the view that a mechanistically neutral adaptation of Gilbert or of Bratman identifies one among several forms of acting jointly.²³

Adopting the mechanistically neutral strategy turns diversity in people's feelings about which cases are paradigms of acting jointly into an advantage. It provides a way of answering questions about when two accounts are incompatible attempts to characterise a single phenomenon rather than compatible attempts to characterise different phenomena. To single out which accounts of acting jointly succeed, we need to investigate mechanisms underpinning their constructs.

8. From purposive to intentional

As promised in §2, our strategy for characterising how acting jointly differs from acting in parallel but merely individually deviates from the standard strategy in two respects. Not only is it mechanistically neutral: it also starts by focusing on acting jointly without any prior assumption that this must be understood as a joint counterpart of ordinary, individual intentional action. This led us to the notion of a collective goal and a corresponding attempt to characterise acting jointly. Invoking collective goals is conceivably sufficient for capturing a sense in which acting jointly can be purposive where the purposivity is not, or not only, a matter of each agent's actions being purposive (or so we argued in §6). But of course none of this is sufficient to capture what many philosophers take the primary target phenomena to be: cases where two or more agents' acting jointly is intentional.

²² Bratman (2014, chaps. 4–5) shows that it is possible in principle to have interpersonal coordination of planning without joint commitment by constructing possible mechanisms. But of course the mechanisms (if any) which actually underpin planning and joint commitment may coincide.

²³ Things could be more complicated, of course. Multiple dissociable mechanisms may be associated with joint commitment, and some but not all of these may coincide with mechanisms underpinning interpersonal coordination of planning.

How might we provide a mechanistically neutral characterisation of a joint counterpart of intentional action?

We noted earlier that an intentional action may be characterised, very roughly and incompletely, as an action that happens because its agent has certain reasons for bringing an outcome about, or at least for attempting to do so. One striking feature of attempts to characterise intentional action in this way is that they all rely, explicitly or implicitly, on a particular way of individuating actions. Actions are individuated by the outcomes to which they are directed. The reasons in virtue of which an action is intentional are reasons which an agent has for bringing about the outcome. Note that while this approach involves no commitment to the existence of actions which are not intentional, it does involve commitment to the possibility of individuating some actions independently of features in virtue of which they are intentional. That is, we are committed to a conceptual distinction between purposive and intentional action.

A parallel account of acting jointly is possible. Intentional cases of acting jointly may be characterised, very roughly and incompletely, as actions that happen because their agents have certain reasons for bringing an outcome about as a collective goal of their actions, or at least for attempting to do so.²⁴ Given that this is a coherent, and natural, extension of a characterisation of ordinary, individual action, we conclude that it is probably possible to provide at least one mechanistically neutral characterisation of acting jointly where acting jointly is intentional.

This characterisation of acting jointly contrasts with mechanistically committed characterisations. According to Ludwig (2015, p. 12), who is formulating a claim common to a range of views, ‘collective intentional activity in general should be seen as the result of shared intentions being satisfied’. By contrast, on our mechanistically neutral proposal, it should be seen as the result of agents having reasons for bringing an outcome about as a collective goal of their actions. Of course, it may turn out that these are extensionally equivalent because no agents have such reasons without having corresponding shared intentions. The disagreement concerns not whether shared intentions are necessary but whether they (or any other states of agents) need be invoked in characterising how acting jointly differs from acting in parallel but merely individually.

²⁴ One way of advancing beyond this rough characterisation is suggested by Laurence (2011).

9. Conclusion

Our question was about which features distinguish acting jointly from acting in parallel but merely individually. We observed that twin problems face the standard, mechanistically committed strategy for answering this question. These problems are how to determine when two accounts should be regarded as competing attempts to characterise a single phenomenon; and how to single out, from a growing number and increasing diversity of accounts, those which are correct (§1). Taking these problems to motivate considering an alternative, we set out to introduce and pursue a mechanistically neutral strategy for characterising acting jointly.

To develop a mechanistically neutral account, we invoked uncontroversial parts of a logical distinction between collective and distributive prediction (§3) and we outlined a mechanistically neutral account of goals (in §4) in order to introduce the notion of a collective goal (§5). According to what we call the Collective Goal Account, when the actions of two or more agents have a collective goal, the agents are, in performing those actions, acting jointly (§6). We also outlined an attempt to capture intentional aspects of acting jointly (§8).

Our thesis is not that the Collective Goal Account is correct. We introduced a method for defending mechanistically neutral accounts borrowed from existing literature (§7). This method partially vindicates the Collective Goal Account but leaves open the possibility that it is one among several significant ways of distinguishing acting jointly from acting in parallel but merely individually. Indeed, we showed that mechanistically neutral adaptations of leading accounts might in principle be vindicated in the same way. Further, our vindication allows that the Collective Goal Account may ultimately need to incorporate mechanistically neutral characterisations of cooperation, commitment, coordination and more besides. If this is right, the Collective Goal Account is at best an incomplete first draft (albeit one which is no worse than the leading accounts in handling contrast cases)—but one that provides a platform on which proponents of planning and joint commitment can build.

Many will reject mechanistically neutral characterisations of acting jointly, not on the grounds that acting jointly necessarily involves some kind of shared intention (this is consistent with a mechanistically neutral characterisation, after all), but on the grounds that, as they see things, there is a conceptual or constitutive connection between the two. (Such a view might be motivated by reflection on [Pacherie](#)

(2013).) We have offered no argument against such opponents; nor do we aim to do so. The issue should be decided according to which strategy yields most progress in understanding what distinguishes acting jointly. The mechanistically committed strategy has dominated discussion to date but faces twin problems that may be challenging to overcome. By contrast, the mechanistically neutral strategy avoids these problems and is, we submit, a promising new contender.²⁵

References

- Alvarez, Maria. 2016. "Reasons for Action: Justification, Motivation, Explanation." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University.
- Asarnow, Samuel. 2020. "Shared Agency Without Shared Intention." *The Philosophical Quarterly* forthcoming. <https://doi.org/10.1093/pq/pqaa012>.
- Austin, James, and Jeffrey Vancouver. 1996. "Goal Constructs in Psychology: Structure, Process, and Content." *Psychological Bulletin* 120 (3): 338–75.
- Bach, Kent. 1978. "A Representational Theory of Action." *Philosophical Studies* 34 (4): 361–79. <https://doi.org/10.1007/BF00364703>.
- Baier, Annette C. 1997. "Doing Things With Others: The Mental Commons." In *Commonality and Particularity in Ethics*, edited by Lilli Alanen, Sarah Heinamaa, and Thomas Wallgren, 15–44. Palgrave Macmillan. https://doi.org/10.1007/978-1-349-25602-0_2.

²⁵ Heartfelt thanks to our critical and generous anonymous referees for the many ways in which their comments have improved this work and will guide our next steps; and to The Editors, whose guidance materially improved our work.

We have benefitted from discussion with many people while preparing this work. For extended discussions and their generous support over a long time we would especially like to thank Olle Blomberg, Chiara Brozzo, Naomi Eilan, Peter Fossey, Mattia Gallotti, Bart Geurts, Guenther Knoblich, Hemdat Lerman, Guy Longworth, Judith Martens, John Michael, Elisabeth Pacherie, Johannes Roessler, Tobias Schlicht, Natalie Sebanz, Axel Seemann, Thomas Smith, Matthew Soteriou, Anna Strasser, Hong Yu Wong, and Cordula Vesper. Thank you!

CS was supported by the Department of Philosophy 'Piero Martinetti' of the University of Milan with the Project 'Departments of Excellence 2018-2022' awarded by the Italian Ministry of Education, University and Research (MIUR) and by the PRIN 2017 project 'The cognitive neuroscience of interpersonal coordination and cooperation: a motor approach in humans and non-human primates' (Cod. Prog. 201794KEER).

- Baus, Cristina, Natalie Sebanz, Vania de la Fuente, Francesca Martina Branzi, Clara Martin, and Albert Costa. 2014. "On predicting others' words: Electrophysiological evidence of prediction in speech production." *Cognition* 133 (2): 395–407. <https://doi.org/10.1016/j.cognition.2014.07.006>.
- Bennett, Jonathan. 1976. *Linguistic Behaviour*. Cambridge: Cambridge University Press.
- Blomberg, Olle. 2015. "Shared Goals and Development." *The Philosophical Quarterly* 65 (258): 94–101. <https://doi.org/10.1093/pq/pqu059>.
- . 2016. "Common Knowledge and Reductionism about Shared Agency." *Australasian Journal of Philosophy* 94 (2): 315–26. <https://doi.org/10.1080/00048402.2015.1055581>.
- Brand, Myles. 1984. *Intending and Acting: Toward a Naturalized Action Theory*. Cambridge, Mass.: MIT Press.
- Bratman, Michael E. 1987. *Intentions, Plans, and Practical Reasoning*. Cambridge, MA: Harvard University Press.
- . 1992. "Shared Cooperative Activity." *The Philosophical Review* 101 (2): 327–41.
- . 1997. "I Intend That We J." In *Contemporary Action Theory, Volume 2: Social Action*, edited by Raimo Tuomela and Ghita Holmstrom-Hintikka. Dordrecht: Kluwer.
- . 2000. "Valuing and the Will." *Noûs* 34 (supplement 14): 249–65. <https://doi.org/10.1111/0029-4624.34.s14.13>.
- . 2014. *Shared Agency: A Planning Theory of Acting Together*. Oxford: Oxford University Press. <http://o-dx.doi.org.pugwash.lib.warwick.ac.uk/10.1093/acprof:oso/9780199897933.001.0001>.
- Brooks, D. H. M. 1981. "Joint Action." *Mind*, New series, 90 (357): 113–19. <http://www.jstor.org/stable/2253670>.
- Brownell, Celia A. 2011. "Early Developments in Joint Action." *Review of Philosophy and Psychology* 2: 193–211. <https://doi.org/10.1007/s13164-011-0056-1>.
- Butterfill, Stephen A. 2001. "Two Kinds of Purposive Action." *European Journal of Philosophy* 9 (2): 141–65.
- Butterfill, Stephen A., and Corrado Sinigaglia. 2014. "Intention and Motor Representation in Purposive Action." *Philosophy and Phenomenological Research* 88 (1): 119–45. <https://doi.org/10.1111/j.1933-1592.2012.00604.x>.
- Carpenter, Malinda. 2009. "Just How Joint Is Joint Action in Infancy?" *Topics in Cognitive Science* 1 (2): 380–92.

- Chant, Sara Rachel. 2007. "Unintentional Collective Action." *Philosophical Explorations: An International Journal for the Philosophy of Mind and Action* 10 (3): 245. <https://doi.org/10.1080/13869790701535246>.
- Clarke, Sam, Luke McEllin, Anna Francová, Marcell Székely, Stephen A. Butterfill, and John Michael. 2019. "Joint Action Goals Reduce Visuomotor Interference Effects from a Partner's Incongruent Actions." *Scientific Reports* 9 (1): 1–9. <https://doi.org/10.1038/s41598-019-52124-6>.
- Csibra, Gergely, and György Gergely. 2007. "Obsessed with Goals': Functions and Mechanisms of Teleological Interpretation of Actions in Humans." *Acta Psychologica* 124 (1): 60–78.
- Davidson, Donald. 1971. "Agency." In *Agent, Action, and Reason*, edited by Robert Binkley, Richard Bronaugh, and Ausonia Marras, 3–25. Toronto: University of Toronto Press.
- della Gatta, Francesco, Francesca Garbarini, Marco Rabuffetti, Luca Viganò, Stephen A. Butterfill, and Corrado Sinigaglia. 2017. "Drawn Together: When Motor Representations Ground Joint Actions." *Cognition* 165: 53–60.
- Dickinson, Anthony. 2016. "Instrumental Conditioning Revisited: Updating Dual-Process Theory." In *Associative Learning and Cognition*, edited by J. B. Trobalon and V. D. Chamizo, 51: 177–95. Edicions Universitat Barcelona.
- Dretske, Fred. 1988. *Explaining Behavior*. Cambridge, Mass.: MIT Press.
- Gallotti, Mattia, and Chris D. Frith. 2013. "Social Cognition in the We-Mode." *Trends in Cognitive Sciences* 17 (4): 160–65. <https://doi.org/10.1016/j.tics.2013.02.002>.
- Gergely, György, Z. Nadasky, Gergely Csibra, and S. Biro. 1995. "Taking the Intentional Stance at 12 Months of Age." *Cognition* 56: 165–93.
- Gilbert, Margaret P. 1990. "Walking Together: A Paradigmatic Social Phenomenon." *Midwest Studies in Philosophy* 15: 1–14.
- . 1992. *On Social Facts*. Princeton, NJ: Princeton University Press.
- . 2010. "Collective Action." In *A Companion to the Philosophy of Action*, edited by Timothy O'Connor and Constantine Sandis, 67–73. Oxford: Blackwell.
- . 2013. *Joint Commitment: How We Make the Social World*. Oxford: Oxford University Press. <http://o-dx.doi.org.pugwash.lib.warwick.ac.uk/10.1093/acprof:oso/9780199970148.001.0001>.

- . 2014. “The Nature of Agreements: A Solution to Some Puzzles about Claim-Rights and Joint Intention₁.” In *Rational and Social Agency: The Philosophy of Michael Bratman*, edited by Manuel Vargas and Gideon Yaffe, 215–56. Oxford; New York: Oxford University Press.
- Ginet, Carl. 1990. *On Action*. Cambridge University Press.
- Gold, Natalie, and Robert Sugden. 2007. “Collective Intentions and Team Agency.” *Journal of Philosophy* 104 (3): 109–37.
- Helm, Bennett W. 2008. “Plural Agents.” *Nous* 42 (1): 17–49. <https://doi.org/10.1111/j.1468-0068.2007.00672.x>.
- Jackson, Robert R., and Fiona R. Cross. 2011. “Spider Cognition.” In *Advances in Insect Physiology*, edited by Jérôme Casas, 41:115–74. Spider Physiology and Behaviour. Academic Press. <https://doi.org/10.1016/B978-0-12-415919-8.00003-3>.
- Jacoby, Larry L. 1991. “A Process Dissociation Framework: Separating Automatic from Intentional Uses of Memory.” *Journal of Memory and Language* 30 (5): 513–41. [https://doi.org/10.1016/0749-596X\(91\)90025-F](https://doi.org/10.1016/0749-596X(91)90025-F).
- Knoblich, Günther, Stephen A. Butterfill, and Natalie Sebanz. 2011. “Psychological Research on Joint Action: Theory and Data.” In *Psychology of Learning and Motivation*, edited by Brian Ross, 51: 59–101. San Diego, CA: Academic Press.
- Kourtis, Dimitrios, Günther Knoblich, Mateusz Woźniak, and Natalie Sebanz. 2014. “Attention Allocation and Task Representation during Joint Action Planning.” *Journal of Cognitive Neuroscience* 26 (10): 2275–86. https://doi.org/10.1162/jocn_a_00634.
- Kutz, Christopher. 2000. “Acting Together.” *Philosophy and Phenomenological Research* 61 (1): 1–31.
- Laurence, Ben. 2011. “An Anscombian Approach to Collective Action.” In *Essays on Anscombe’s Intention*. Cambridge, MA: Harvard University Press.
- Leadbeater, Ellouise, and Lars Chittka. 2005. “A New Mode of Information Transfer in Foraging Bumblebees?” *Current Biology* 15 (12): R447–48.
- Linnebo, \tf="P4BBD39"\char"5Bystein. 2005. “Plural Quantification.” In *The Stanford Encyclopedia of Philosophy (Spring 2005 Edition)*, edited by Edward N. Zalta. Stanford, CA: Metaphysics Research Lab, Stanford University.
- Loehr, Janeen D., and Cordula Vesper. 2015. “The Sound of You and Me: Novices Represent Shared Goals in Joint Action.” *The*

- Quarterly Journal of Experimental Psychology* 0 (ja): 1–30. <https://doi.org/10.1080/17470218.2015.1061029>.
- Longworth, Guy. 2019. “Sharing Non-Observational Knowledge.” *Inquiry* 0 (0): 1–21. <https://doi.org/10.1080/0020174X.2019.1680430>.
- Ludwig, Kirk. 2007. “Collective Intentional Behavior from the Standpoint of Semantics.” *Nous* 41 (3): 355–93. <https://doi.org/10.1111/j.1468-0068.2007.00652.x>.
- . 2015. “Shared Agency in Modest Sociality.” *Journal of Social Ontology* 1 (1): 7–15. <http://www.degruyter.com/view/j/jso.2015.1.issue-1/jso-2014-0046/jso-2014-0046.xml>.
- . 2016. *From Individual to Plural Agency: Collective Action*. Oxford University Press.
- Ménoret, Mathilde, L. Varnet, R. Fargier, A. Cheylus, A. Curie, V. des Portes, T. A. Nazir, and Y. Paulignan. 2014. “Neural Correlates of Non-Verbal Social Interactions: A Dual-EEG Study.” *Neuropsychologia* 55: 75–97.
- Meyer, Marlene, Robrecht P. R. D. van der Wel, and Sabine Hunnius. 2013. “Higher-Order Action Planning for Individual and Joint Object Manipulations.” *Experimental Brain Research* 225 (4): 579–88. <https://doi.org/10.1007/s00221-012-3398-8>.
- Novembre, G., L. F. Ticini, S. Schutz-Bosbach, and P. E. Keller. 2014. “Motor Simulation and the Coordination of Self and Other in Real-Time Joint Action.” *Social Cognitive and Affective Neuroscience* 9 (8): 1062–68. <https://doi.org/10.1093/scan/nst086>.
- Pacherie, Elisabeth. 2010. “The Phenomenology of Joint Action: Self-Agency Vs. Joint-Agency.” In *Joint Action*, edited by Axel Seemann. MIT Press.
- . 2012. “Intentional Joint Agency: Shared Intention Lite.” *Synthese* forthcoming.
- . 2013. “Intentional Joint Agency: Shared Intention Lite.” *Synthese* 190 (10): 1817–39. <https://doi.org/10.1007/s11229-013-0263-7>.
- Pears, David. 1971. “Two Problems about Reasons for Actions.” In *Agent, Action and Reason*, edited by A. Marras, R. Binkley, R. Bronaugh, 128–53. Oxford: Oxford University Press.
- Petersson, Björn. 2007. “Collectivity and Circularity.” *Journal of Philosophy* 104 (3): 138–56.
- Pettit, Philip, and David Schweikard. 2006. “Joint Actions and Group Agents.” *Philosophy of the Social Sciences* 36 (1): 18–39. <https://doi.org/10.1177/0048393105284169>.

- Ramenzoni, Verónica C., Natalie Sebanz, and Günther Knoblich. 2014. "Scaling up Perception-action Links: Evidence from Synchronization with Individual and Joint Action." *Journal of Experimental Psychology: Human Perception and Performance* 40 (4): 1551–65. <https://doi.org/10.1037/a0036925>.
- Roth, Abraham Sesshu. 2004. "Shared Agency and Contralateral Commitments." *The Philosophical Review* 113 (3): 359–410.
- Sacheli, Lucia Maria, Elisa Arcangeli, and Eraldo Paulesu. 2018. "Evidence for a Dyadic Motor Plan in Joint Action." *Scientific Reports* 8 (1): 5027. <https://doi.org/10.1038/s41598-018-23275-9>.
- Sacheli, Lucia Maria, Margherita Adelaide Musco, Elisa Zazzera, and Eraldo Paulesu. 2021. "Mechanisms for Mutual Support in Motor Interactions." *Scientific Reports* 11 (1): 3060. <https://doi.org/10.1038/s41598-021-82138-y>.
- Schlosser, Markus. 2019. "Agency." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2019. Metaphysics Research Lab, Stanford University.
- Schmid, Hans Bernhard. 2008. "Plural Action." *Philosophy of the Social Sciences* 38 (1): 25–54. <https://doi.org/10.1177/0048393107310877>.
- Schmitz, Laura, Cordula Vesper, Natalie Sebanz, and Günther Knoblich. 2017. "Co-Representation of Others' Task Constraints in Joint Action." *Journal of Experimental Psychology: Human Perception and Performance* 43 (8): 1480–93. <https://doi.org/10.1037/xhp0000403>.
- Scott, Jarrod J., Kevin J. Budsberg, Garret Suen, Devin L. Wixon, Teri C. Balsler, and Cameron R. Currie. 2010. "Microbial Community Structure of Leaf-Cutter Ant Fungus Gardens and Refuse Dumps." *PLOS ONE* 5 (3): e9922. <https://doi.org/10.1371/journal.pone.0009922>.
- Searle, John R. 1990. "Collective Intentions and Actions." In *Intentions in Communication*, edited by P. Cohen, J. Morgan, and M. E. Pollack, 90–105. Cambridge: Cambridge University Press.
- Sebanz, Natalie, Harold Bekkering, and Günther Knoblich. 2006. "Joint Action: Bodies and Mind Moving Together." *Trends in Cognitive Sciences* 10 (2): 70–76.
- Tollefsen, Deborah. 2005. "Let's Pretend: Children and Joint Action." *Philosophy of the Social Sciences* 35 (75): 74–97.
- Tuomela, Raimo. 2000. *Cooperation: A Philosophical Study*. Dordrecht: Springer.

- Tuomela, Raimo, and Kaarlo Miller. 1985. "We-Intentions and Social Action." *Analyse & Kritik* 7 (1): 26–43. <https://doi.org/10.1515/auk-1985-0102>.
- . 1988. "We-Intentions." *Philosophical Studies* 53 (3): 367–89. <https://doi.org/10.1007/BF00353512>.
- Velleman, David. 1997. "How to Share an Intention." *Philosophy and Phenomenological Research* 57 (1): 29–50.
- Wilson, George, Samuel Shpall, and Juan S. Pinerros Glasscock. 2016. "Action." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2016. Metaphysics Research Lab, Stanford University.
- Wright, Larry. 1976. *Teleological Explanations*. Berkeley: University of California Press.