PHILOSOPHICAL PSYCHOLOGY
2025, VOL. 38, NO. 5, 1909–1931
https://doi.org/10.1080/09515089.2025.2490242

Routledge
Taylor & Francis Group

Three strategies for shared intention: plural, aggregate and reductive

Stephen A. Butterfill

Department of Philosophy, University of Warwick, Coventry, UK

ABSTRACT

When deciding on a strategy for explicating shared intention, we all face two fundamental questions. First, can an intention or any other mental state have more than one subject? A positive answer to this allows the plural subject strategy: shared intention is a matter of there being one mental state with two or more subjects. Mental states are shared in the same sense that siblings share a parent; no simpler view exists. A negative answer blocks the plural subject strategy. This motivates asking the second fundamental question. Are there aggregate subjects and, if so, can they have intentions? The aggregate strategy depends on a positive answer to this question: the idea is that shared intention is a matter of there being aggregate subjects of mental states, that is subjects of mental states with proper parts that include subjects of mental states. By contrast, a negative answer to this question limits us to the reductive strategy: shared intention is a structure of ordinary, individual subjects' emotions, intentions and other mental states. I contribute a limited review of the three strategies. I also defend a novel thesis. Whereas these strategies are often presented as conflicting attempts to characterize a single set of phenomena, my thesis is that for each strategy there are phenomena which can be correctly characterized only by following that strategy. Instead of attempting to find one true strategy, we may need to seek ways to combine insights from different strategies.

ARTICLE HISTORY
Received 2 October 2023
Accepted 28 March 2025

KEYWORDS
Shared intentionality; joint action; intention; plural subject

1. Introduction

Start with the idea that shared intention, whatever that is, is something which makes things we do together the genuinely joint activities they are. We manifest shared intention in walking together, playing a piano duet, or painting a house together. Philosophers, in attempting to elucidate ideas about shared intention, have followed three distinct strategies. One involves plural subjects, one aggregate subjects,[1] and one a reduction of

CONTACT Stephen A. Butterfill s.butterfill@warwick.ac.uk Department of Philosophy, University of Warwick, Coventry, CV4 7AL, UK

the apparently plural or aggregate to the merely individual.[2] Whereas these strategies are usually presented as conflicting attempts to characterize a single set of phenomena, my thesis is that for each strategy there are phenomena which can be correctly characterized only by following that strategy.

This thesis is not entirely novel. Both List and Pettit (2011) and Bratman (2022a), for instance, have developed views on which, roughly speaking, a reductive strategy is applied to small-scale, informal interactions whereas the aggregate strategy is applied to corporate or institutional agents.[3] My thesis is that, similarly, a combination of strategies may be needed even when restricting attention to small-scale interactions involving two or three people.

What consequences would follow from this thesis if true – if, that is, each of the three strategies is needed to characterize some phenomena associated with shared intention? This would make things easier in one way: it is not necessary for proponents of one strategy to refute the theories invented by those following a different strategy. This would be a good outcome because, despite much effort, few if any existing attempts to refute particular theories of shared intention are widely regarded as successful. But it would also make things harder in another way. For if there is not one set of things that the different theories are aiming to fully capture, then it will not be possible to specify the target of a theory simply by reference to shared intention. We will need a deeper understanding of what each particular theory is aiming to explain.

Superficially, the need for a deeper understanding of targets of explanation might not seem like much of a challenge. After all, we already know that some theorists primarily target semantics (Ludwig, 2007), others "the deep structure of our thought about acting together" (Gilbert, 2022, p. 2), and yet others "the explanatory structures that directly underlie [. . .] cases of acting together" (Bratman, 2022b, p. 7). There are, then, already things written about what theories of shared intention aim to explain. Yet little attention has been given to the possibility of co-existence. The usual assumption is that the various strategies lead to theories which do contradict each other; or else, more rarely, that multiple strategies can be combined in pursuit of a single explanatory target.

We already know that it is coherent to pursue a combination of strategies because some have already done this. Bratman (2014, 127ff), for example, considers the possibility that his reductive strategy yields a theory on which there are also what he calls "group agents" (for which I shall use the term "aggregate agents"; see section 5). This is a relatively conservative way of combining strategies. Not only is there a single explanatory target, but the group agents are merely epiphenomena of the reductive strategy. The novel feature of what follows is the observation that different explanatory targets

require different strategies, which, as we will see, appears to create a challenge to understanding co-existence.

If, as the following argues, multiple different strategies for shared intention really are needed, then we all face a problem. The problem is that we, as researchers, need, but lack, a common understanding of what theories of shared intention are theories of.

2. Living out a theory

My thesis is that, for each strategy for explicating shared intention, there are phenomena which can be correctly characterized only by following that strategy. The argument for this thesis (in Section 8) will hinge on the idea that it is possible, in some cases, to make a theory true by living it out. To avoid surprises later, this section introduces that idea.

While laws of mechanics apply no less to us agents than to anything else, it would usually be futile, perhaps even incoherent, to attempt to move according to these laws. By contrast, one of the roles of mental state attribution is to provide norms which individuals can measure themselves against and aim to live by (McGeer, 2007; Zawidzki, 2013).

To illustrate, consider the norm of agglomeration: it is a mistake to knowingly have several intentions if it would be a mistake to knowingly have one large intention agglomerating the several intentions (Bratman, 1987). Whether this is actually a mistake is controversial – several philosophers have defended views of intention which are incompatible with it (Setiya, 2022, section 4). But regardless of that, it is possible some people might, however mistakenly, take agglomeration as an ideal by which to live. They check their intentions against the norm and criticize each other for failing to implement it. It is equally possible that another group of people, having considered the matter deeply, intentionally disregard the norm of agglomeration. In their view, adhering to this norm would be a mistake.

The possibility that some ordinary agents might adopt or reject the norm of agglomeration in practice raises a question. Are ordinary agents' views ever relevant to whether the norm is correct? Imagine we were to say, crudely, that whether the norm holds is just a matter of whether people take it to hold. Such a view faces myriad challenges. One is to accommodate the fact that ordinary agents are wrong about norms, at least occasionally (as the present author can attest). Another challenge is to avoid a regress. This is not a line I propose to develop. Alternatively, one might take a hard line and insist that ordinary agents' views are irrelevant to whether the norm of agglomeration holds. Taking this line is complicated by the fact that ordinary agents' views shape at least some of their thoughts and actions. Their views are not idle speculations about themselves but can form ideals which they attempt to live out. Someone who takes the hard line cannot therefore

claim to be explaining how ordinary agents think or act. Minimally, then, anyone pursuing this line would have to identify which phenomena their position is supposed to explain. A further challenge is that philosophers' methods involve intuition, imagination and reasoning about consistency. These methods are good for identifying possible ways things could be. But where there are multiple theoretically coherent positions on which fully-informed ordinary agents reasonably differ, these methods are not likely to yield insight into how things actually are.

Resolving the issue of how, if at all, ordinary agents' views are relevant to the correctness of the norm of agglomeration is beyond anything I can offer here (or anywhere else). But for our purposes, what matters is a relatively uncontroversial point. There is a difference between, on the one hand, things which are merely described and predicted by a set of attributions and, on the other hand, agents who are attempting to live out a set of attributions together with some norms governing them. You might have a view about how combinations of yeast, sugar and heat can be used to influence how dough rises, but the dough itself has no perspective. And even if the dough achieved self-awareness, that would matter only insofar as its self-awareness influenced variables you care about. When ascribing attitudes and norms to agents, by contrast, philosophers are not required to adopt the outsiders' perspective – they can also take into account the agents' own perspective.

Characterizing intentions and norms from some agents' own perspective is a familiar and coherent philosophical project. Where some agents are attempting to live out a theory, it is reasonable to accept, in the absence of overriding reasons such as incoherence, ignorance or inertness, that the theory could be true of them.

Perhaps this will seem too hypothetical to be worth taking seriously. We have no idea which, if any, ordinary agents aim to live by the norm of agglomeration and which, if any, aim not to. My sense, however, is that philosophical theories are not supposed to depend on any such facts. They are, after all, usually developed independently of any investigation into what ordinary agents think.[4]

Not that it has to be hypothetical. Consider the familiar distinction between revealed and stated preferences. To illustrate, in investigating the value people place on a life, we could observe how much less people pay to own a house near a known source of carcinogenic pollution. Or, alternatively, we could give them a questionnaire asking how much they would pay to eliminate the same risk. A wide range of research has examined how revealed and stated preferences diverge (Carson et al., 1996; Alberini, 2019, for example). There is also research on the factors that 'often create a wedge between revealed and normative preferences' (Beshears et al., 2008, p. 1788).

Divergences between revealed and stated preferences matter in practice because they create difficult questions for policy makers on how much to invest in preventing deaths. For our, more theoretical purposes, the divergence illustrates how understanding agents' actions requires taking into account their own perspectives. Agents' stated preferences are views about how a model applies to them. Because agents sometimes aim to live out these views, they are not inert commentary. The problem is not fundamentally that this complicates predicting behavior. It is that where agents aim to conform to a model, incorrect predictions do not have the same significance.

The distinction between revealed and stated preferences illustrates the dual role of theories of attitudes and norms, in prediction and in offering ideals which people can attempt to live out. There is a difference between, on the one hand, things which are merely described and predicted by a set of attributions and, on the other hand, agents who are are attempting to live out a set of attributions together with some laws governing them. In the latter case, the agents' aiming to live out a theory is a reason not to reject that theory insofar our aims include understanding the agents' own perspective.

The argument that follows is an attempt to apply this general point to philosophical theories of shared intention. My aim is to show that for each of the three strategies for shared intention – plural, aggregate and reductive – there is at least one consistent theory following this strategy which it is possible for people to intentionally live out. This thesis is probably either too odd or too obvious to be interesting in its own right, but I will suggest that it has consequences which complicate existing attempts to understand shared intention. As these consequences depend on the aims of a theory of shared intention, I start with the aims.

3. Background on shared intention

Why do we need a notion of shared intention at all? Because it is what distinguishes genuinely joint activities from things people do in parallel but merely individually (Bratman, 2022b; Gilbert, 1990). This is, of course, at most a partial answer. The hope is that investigating a notion such as shared intention will enable us, eventually, to "discover the nature of social groups in general" (Gilbert, 1990, p. 2) and to understand the conceptual, metaphysical and normative aspects of basic forms of sociality (Bratman, 2014, p. 3). But one route to these lofty goals is to focus on distinguishing genuinely joint from merely parallel activities in mundane cases involving two or three people.

Ayesha and Ahmed have spent the morning in the kitchen washing the dishes together. This is a paradigm case of joint activity. We can contrast Ayesha and Ahmed's activities with those of two anti-social people who act

in parallel but merely individually. The idea is that these other, anti-social people wash the dishes side-by-side, but their actions are merely performed in parallel and so do not involve any joint activity. What distinguishes Ayesha and Ahmed's activities from these other people's?

A temptingly simple idea is to appeal to coordination. Could Ayesha and Ahmed's activities be distinguished by virtue of being coordinated? The obstacle is that acting in parallel can also involve coordination. The actions of the other, anti-social people, who are merely acting in parallel, may nevertheless need to be tightly coordinated because space in their communal kitchen is limited. They may also politely anticipate each other's movements and work around them. Mere coordination, then, cannot distinguish joint activity.

The failure of this and other simple ideas hints that distinguishing joint activities from their parallel but merely individual counterparts is a deep and difficult problem. This problem is a variant of one about ordinary, individual action. The "Problem of Action" is to distinguish a person's actions from things that merely happen to them (Frankfurt, 1978). If ordinary, individual intention is key to solving that problem, perhaps some joint counterpart of intention is the key to solving the problem of joint action. This motivates using the term *shared intention* to label whatever the normative or psychological structure is needed to distinguish joint activities from things people do in parallel but merely individually.

Introducing a label for the problem does not take us very far toward a solution. The problem now becomes to say what shared intention is. One possibility is the plural subject strategy: shared intention is a matter of there being one mental state with two or more subjects.

4. The plural subject strategy

A plural subject is two or more people who are each among the subjects of a single intention or other mental state.

The difficulty of understanding the plural subject strategy is mainly that it is so simple. Intentions and other mental states involve subjects, attitudes and contents. The content is what distinguishes two intentions from each other – the intention to cook dinner from the intention to go for a walk, say, differ in content. Attitude is what distinguishes intentions from other mental states – the intention to go for a walk differs in attitude from the desire to go for a walk. And the subject is what distinguishes your mental states from mine. The idea of the plural subject strategy is just that intentions can have more than one subject. You and I can share an intention in the same sense that siblings share a parent. Your intention to walk may also be my intention to walk: you and I are equally subjects of this intention.

It is helpful, in thinking about the contrast between individual and plural subjects, to draw on a related distinction between distributive and collective interpretations of sentences. Consider these sentences:

(1) The fans left the stadium.
(2) The fans completely blocked the road.

The first sentence is naturally read distributively: it is a matter of each fan individually leaving. But the second sentence is naturally read collectively. As the road is very wide, not even the largest individual fan did much at all to block the road. But because so many fans were milling around in the road, it was impossible to traverse it. So understood, the second sentence's truth is not, or not only, a matter of each fan individually blocking the road. This is a collective reading. The distinction seems applicable to sentences about intention:

(3) The twins intended to win the race

If we imagine a race that can only have one winner, a 100 meter sprint, say, then it is natural to read this sentence distributively. Its truth is just a matter of each twin intending to win the race. But if we imagine the twins running in a three-legged race together, it seems possible to read the sentence collectively. On this reading, there is one intention whose subject is the twins. They are, to put it colorfully, of one mind.

We can describe the twins as a *plural subject*. But note that the plural subject is nothing other than the twins themselves. We must avoid confusion on this point in order to distinguish plural subjects from aggregate subjects, which are fundamentally different (more on this in section 6).

The plural subject strategy requires no novel conceptual, metaphysical or normative ingredients over and above those already required in a theory of ordinary, individual action – it can be implemented in ways that respect Bratman's continuity thesis.[5] Just as the truth of the collective reading of (2) does not require anything other than the fans to mill in the road, so the truth of statements about intention collectively read requires nothing other than intentions and their subjects.[6]

But are there really plural subjects of intention? One possibility is that there merely seem to be, and are not actually, collective readings of sentences about intentions like (3) above. There are a range of objections along these lines. Most extreme is the claim that all apparently collective predication is really disguised distributive predication. A more limited objection is that statements about actions and intentions merely seem to have collective readings. Ludwig (2016, chap. 9) offers a detailed discussion along these lines. His conclusion is carefully nuanced:

> we do not need to accept genuine plural [...] agents into our ontology in order to accept what we say about [...] collective action, at least insofar as we express this using plural subject terms. (Ludwig, 2016, p. 168)

Ludwig might be right that semantic considerations do not force us to accept that plural subjects exist. Given his further premise that plural subjects should be avoided if possible, this would be a compelling argument against the existence of coherent collective readings. But, importantly, Ludwig finds nothing forcing us to reject their existence either. So as long as there is either no general presumption against plural subjects or else sufficient reason to suppose that they are necessary, it is not incoherent to imagine statements about intention have true collective readings. Minimally, collective readings are a helpful tool for clarifying what the plural subject strategy is.

But is it really coherent to suppose that intentions might have more than one subject? One might object that to have an intention it is necessary to have a mind; and that having a mind minimally involves having a range of mental states, and perhaps even being self-aware.[7] Schmid (2013), who has perhaps the best-developed version of the plural subject strategy, accepts this constraint but argues that there are no good grounds for supposing that it could not be met. Likewise, Helm (2008) argues that there are plural subjects with a range of "emotions and desires in the right sort of rational structure" (Helm, 2008, p. 29). Of course this would mean that plural subjects are unlikely to be involved in spontaneous interactions between strangers, as when I am struggling to propel my heavy push chair on to the bus and you helpfully seize the front and we lift together. Plural subjects on views like Schmid's or Helm's would require vastly more intimate, long-term connections between individuals.

If we follow Schmid or Helm, it is possible to wonder how there could be plural subjects. And the sense of mystery one might have about this could, perhaps, motivate rejecting the entire strategy in favor of apparently less mysterious alternatives. But this would be an error. Any of the strategies can be developed in ways that will seem mysterious to at least some philosophers. But one of my aims is to draw attention to the existence of straightforward, nonmysterious ways of developing each strategy.

An alternative, potentially less mysterious plural subject view might be based on rejecting the claim that having an intention entails having a range of mental states. This view could be inspired by reflection that humans are prone to attribute mental states on the slightest of pretexts to things which, as they know, lack not only minds but even physical bodies (Heider & Simmel, 1944).[8] Whether or not things cannot actually have intentions without having minds, no such constraint appears to apply to at least one significant strand of everyday thinking. Perhaps, then, the plural subject

strategy is needed for capturing ways in which some people sometimes think.

We can take this one step further. Here is a way in which you and I could become the plural subject of an intention. We each somehow become convinced, however mistakenly, that you and I are the plural subjects of an intention to cook dinner. This thought might influence our behavior: thinking, perhaps mistakenly, that having this intention means we are subject to various norms, we might aim to act in ways that conform to them. We are, by our lights, acting as if we had this intention. We could also be taking for granted that we were plural subjects of a range of other beliefs, desires and mental states, and perhaps explicitly attributing some as our activity unfolds. And others, if they became convinced, perhaps mistakenly, that we had this intention, might also think and act accordingly. In this way, what began as merely a mistake became real enough to shape the social world through being adopted as a normative ideal.[9]

In this section I have introduced the plural subject strategy and offered a preliminary and superficial case for its theoretical coherence. Following this strategy can lead to various quite different theories. On some theories, the existence of plural subjects involves long-term, intimate connections capable of supporting a shared mental life. On other possible theories, plural subjects can be temporary phenomena arising from the specific needs of a moment.

This falls short of showing that the plural subject strategy is successful. For it to succeed, minimally there must be cases in which the existence of plural subjects is actually what distinguishes joint activities from things people do in parallel but merely individually. (At least that is what one quite prominent approach requires, as we saw in section 3.) I have tried to indicate the difficulties involved in showing that no such cases could exist. But of course we have not seen positive grounds to suppose that there are now, or have ever been, any intentions (or other mental states) which do have plural subjects. Following the approach of section 2, the possibility that there is a theoretically coherent plural subject theory which some people could aim to live out is reason to accept that such a theory could capture an aspect of shared intention.

For what it is worth, my own sense is that it would be quite hard to establish, in practice, that a particular situation did involve a plural subject, and that the difficulty of doing so may have been underestimated. To illustrate, Schmid (2008) claims that plural agents feature in common sense thinking. This claim is hard to evaluate because other philosophers would probably reject it (Ludwig, 2016, perhaps). Certainly philosophers seem vulnerable to making wrong assumptions about common sense thinking in other cases (e.g., Nagel et al., 2013; Starmans & Friedman, 2012, 2013). One as yet unresolved challenge is

to develop an operationalization which would enable us to distinguish someone thinking and acting in terms of plural subjects from someone operating with a different conception of shared intention. Whereas we have multiple methods for identifying stated preferences (including contingent valuation and choice modeling), we currently lack any hint about how we might identify attitudes toward plural subjects. This challenge is made harder by the need to distinguish plural subjects from aggregate subjects.

5. The aggregate subject strategy

An aggregate subject is a subject with proper parts which are themselves subjects in their own right. If you have intentions and some proper parts of you also have intentions of their own, then you are an aggregate subject.

Although our interest is in subjects of intention, and of mental states generally, a non-mental illustration may be helpful. The Portuguese man o'war, *Physalia physalis*, is an animal composed of polyps which are themselves born as animals in their own right. Why is the man o'war an aggregate subject rather than a plural subject? Because it is numerically distinct from the animals which compose it. These may change over its lifetime. By contrast, in the case of a plural subject, there is nothing that could continue to exist if one of the individuals ceases to exist. A plural subject is not a thing at all: it is just some individuals. An aggregate subject, even one which right now consists of nothing but some individuals, is nevertheless a thing that is logically distinct from the individuals which comprise it.

For a non-mental illustration which features both plural and aggregate subjects, consider:

(4) The protestors formed a barrier which blocked the entrance.

Forming the barrier is something the protestors do collectively. They (and no other thing) are the plural subject of the forming. But in talking about the barrier we have introduced an aggregate entity. Although it is composed of the protestors and nothing else, it is numerically distinct from them. We know the barrier is distinct from the protestors because one of the protestors might abandon the barrier and be replaced by a new protestor.

Not everything true of an aggregate subject is true of a corresponding plural subject, even when, as in the protestors' case, the only parts of the aggregate subject are the plural subject. For example, the barrier may be capable of surviving an assault which would destroy the plural subject. And, conversely, it is true that the plural subject formed the barrier but false that the aggregate subject did so. The mental case is similar. To be an aggregate subject, a thing must have its own intentions (or other mental states) which

are at least potentially distinct from those of its parts (Björnsson & Hess, 2017, p. 274).

How could there be aggregate subjects of intention? As List and Pettit put it:

> Let a collection of individuals form and act on a single, robustly rational body of attitudes [. . .] and it will be an agent.

Individuals sometimes act in this way only because their interests are so closely aligned, as when a variety of finance professionals all rush to exploit a tax loophole so that state finances appear ravaged by a many-handed beast. [10] Alternatively, individuals may authorize a representative to speak for them as a group. Such cases are unlikely to be theoretically interesting given our aim of investigating shared intention more broadly (List & Pettit, 2011, 7ff). Instead we should focus on cases where an aggregate agent has what Sugden calls *autonomy*:

> An aggregate subject has autonomy if there is "the possibility that every member of the group has an individual preference for y over x (say, each prefers wine bars to pubs) while the group acts on an objective that ranks x above y." (Sugden, 2000)

The challenge, then, is to explain how there can be aggregate subjects which are autonomous from the subjects which compose them.

One approach to meeting this challenge borrows from decision theory. It is possible to use decision theory as an "elucidation of the notions of subjective probability [roughly, belief] and subjective desirability or utility [roughly, desire]" (Jeffrey, 1983, p. xi). Whether or not there are other ways of elucidating attitudes, decision theory provides one coherent way of thinking about them. But decision theory is also agnostic about what subjects are. As long as a thing's behavior fits a certain pattern, one that is specified by axioms linking attitudes to actions, the thing can coherently be attributed preferences. This is why decision theory and its derivatives can be applied not only in describing humans but also bacteria, business organizations and countries (Dixit et al., 2014, chap. 10). Being agnostic about what subjects are makes it a useful tool for constructing a theory of aggregate agents.

In essence, the construction goes like this.[11] Two or more individuals take themselves, rightly or wrongly, to be components of an aggregate agent. These individuals each ascribe preferences, and perhaps other attitudes, to the aggregate agent, and they all ascribe the same attitudes. They then use these preferences to rank combinations of individual actions, and each individual selects an action from a highest-ranking combination. Given the usual axioms about preferences being transitive and so on (Steele & Stefánsson, 2020), and given some background assumptions about the individuals' knowledge of their situation, it will be possible to use decision

theory to model the situation as if there were an aggregate agent. And if we follow Jeffrey in taking decision theory as elucidating preference and other attitudes, we can infer that there actually is an aggregate agent. Further, because the preferences and other attitudes ascribed by the individuals need not be their own, the aggregate agent has autonomy in the above sense.

How might the aggregate subject strategy provide a notion of shared intention? And how might it enable us to distinguish joint actions from things people do in parallel but merely individually (see section 3)? One possibility is to stipulate that the intentions arrived at by individuals through the process of determining how the aggregate agent will act comprise a shared intention (Gold & Sugden, 2007; alternative views are offered by; Bardsley, 2007; Pacherie, 2013). On this view, one way for an activity to be genuinely joint is for it to issue from reasoning about the preferences of an aggregate subject where each reasoner is a part of that aggregate subject.[12]

None of this shows, of course, that there actually are aggregate subjects of preference or intention. Even assuming there are, we cannot yet say whether their existence is actually what distinguishes joint activities from things people do in parallel but merely individually. My aim in this section was merely to defend the theoretical possibility of aggregate subjects.

6. The reductive strategy

If you seek to characterize shared intention entirely in terms of ordinary, individual subjects and their ordinary, individual attitudes then you are pursuing the reductive strategy.

Contemporary interest in the reductive strategy starts with Sellars (1963, p. 203)'s observation that statements to the effect that we intend that we cook dinner are "clearly not the logical sum of" statements about each of us individually intending that we cook dinner (Tuomela & Miller, 1988). Apparently, then, our having a shared intention that we cook dinner together cannot consist simply in our each intending this. A natural question is whether there is any combination of ordinary, individual intentions, knowledge states or other mental attitudes our having which could be necessary or sufficient for us to have a shared intention that we cook dinner.

The most extensively developed and widely discussed attempt to provide sufficient conditions for shared intention is Bratman (2014)'s. The full account is complex but the core idea, put roughly, is this. For us to have a shared intention that we cook dinner, it suffices that we each intend that we cook dinner, that we intend to do so by way and because of these intentions, and that this is all common knowledge among us.

Proponents of the reductive strategy have succeeded in providing sets of necessary or sufficient conditions which have intuitive pull for some

and against which none of the published counterexamples have been widely accepted as successful. This is remarkable given that the model for Bratman and several others is Grice's analysis of meaning (Bratman, 1992, footnote 13 to p. 334), which met a different fate. A "flood" of counterexamples to Grice's analysis led to extensive revisions, to which further counterexamples were developed (Searle, 2007, p. 11). Confidence in the project's eventual success was shaken when, in a dramatic change of direction, Schiffer, who was formerly a leading proponent of the Gricean analysis, argued that the whole project was based on a mistake.[13] While Grice's analysis has continued to inspire various projects (Moore, 2017, for example), we are no closer to a successful reductive analysis. By contrast, generally accepted counterexamples to reductive sets of necessary or sufficient conditions for shared intention appear to be rare.

A diversity of views about the features of shared intentions can be found in the reductive strategy. The various conditions proposed imply conflicting views about whether having a shared intention invariably involves dispositions to help (for: Bratman, 2014, pp. 56–57; against: Bratman, 1992; Ludwig, 2007), common knowledge (for: Bratman, 1993; against; Blomberg, 2016), and corresponding individual intentions on each subject's part (for: Bratman, 1992; against: Sellars, 1963). There are also further issues on which theorists could disagree, including on whether shared intention invariably involves contralateral commitment, cooperation (Salomone-Sehr, 2022) or nonobservational knowledge (Roessler, 2024).

This gives rise to a challenge to the reductive strategy. For the diversity makes it unclear when proponents of the reductive strategy can coherently be interpreted as offering competing attempts to characterize a single thing and when as offering compatible attempts to characterize different things.[14]

The reductive strategy allows us to construct many theories, each internally theoretically coherent but inconsistent with other reductive theories that share an explanandum. For any combination of views about the features of shared intention, it would be possible to construct a coherent reductive theory. If we simplify and regard the six features mentioned above as binary, this yields 64 reductive theories. The scarcity of counterexamples cuts two ways.

7. The limits of a metatheoretical principle

We have seen that the three strategies each yield theoretically coherent positions, and, further, that at least one of these strategies alone yields many theoretically coherent positions. Can we decide between the positions by invoking metatheoretical principles? This idea has been carefully developed by Bratman:

> If we can get a plausible model of modest sociality without appealing to a fundamental discontinuity in the step from individual planning agency to such sociality, then there is a presumption against an appeal to such a discontinuity in our theorizing. (Bratman, 2014, p. 36)

If true, this principle provides a good reason to prefer the reductive strategy over theories like that of Searle (1990) and perhaps also that of Gilbert (2013).[15] This is because each of those theorists postulates a fundamental discontinuity. In Searle's case, this is a novel kind of attitude, the "we-intention", which differs from ordinary intention along the same dimension as desire differs from intention.[16] For her part, Gilbert postulates a novel kind of commitment and associated nonmoral norms. The novel kinds of intention and commitment are fundamental discontinuities.

Although Bratman's metatheoretical principle rules against some theories, it does not exclude many of those derived from the reductive strategy. Nor does it exclude the plural and aggregate subject strategies outright. After all, the whole point of plural subjects is that they are nothing but some subjects (Boolos, 1984). And the bare idea of an aggregate subject is no more a fundamental discontinuity than is a barrier composed of protestors. Further, as we have seen (in sections 4 and 5), both plural subject and the aggregate subject strategies can be implemented without appealing to fundamental discontinuities.

Apparently, then, Bratman's metatheoretical principle is limited. There are theories from each of the three strategies – plural, aggregate and reductive – between which it fails to discriminate.

This is why I have presented the strategies in an unusual order. The usual way is to start with the reductive strategy and then possibly to justify adopting one of the others by some failure of that strategy (see, for example, Helm, 2008). In my view that is a mistake. There is no consensus on attempts to demonstrate failure of the reductive strategy generally. Quite the opposite: after three decades there is not yet even a successful published counterexample to the most widely discussed reductive theory (Bratman, 1992, 2022b). But, equally, the mere absence of successful objections to a theory is not enough to establish its truth. The plural and aggregate subject strategies are significant not because the reductive strategy can be shown to fail but because they also yield theoretically coherent, as yet unfalsified theories.

These are deep waters. Some researchers hold that plural subjects should be avoided if possible, even narrowly logical ones (see Ludwig, 2016 cited in section 4). Some even suggest that "all plural locutions should be paraphrased away".[17] Were this true, the plural subject strategy might be ruled out for reasons not specifically psychological. As this illustrates, the considerations offered here fall far short of demonstrating that it would be

impossible to find general principles which do discriminate among the three strategies for shared intention.

But there is also a positive argument for my thesis that all three strategies are needed.

8. Do we need multiple strategies for shared intention?

It is possible for people to intentionally live out one or another theory of shared intention: to think and act as if that theory were true of them. This indicates that no one theory alone could be sufficient to fully characterize shared intention. Or so I will argue in this section.

For each of the three strategies, there are recipes you and I could explicitly aim to follow. To illustrate, suppose the time for us to face the growing pile of dirty dishes in our kitchen has finally come. Having both been inspired by Schmid (2008), we might regard ourselves as the plural subject of an intention to wash the dishes and act accordingly (see section 4). Or perhaps what comes to mind is instead an idea about ascribing preferences to an aggregate subject and doing our parts to fulfill them (see section 7). Or maybe we have both just been reading Bratman (2014) and are impressed that we could benefit by forming and making explicit the intentions he identifies, thereby meeting his sufficient conditions for shared intention (see section 6). This being new to us, we even decide to write everything down so that we can track the attitudes and actions. Things go well and we continue to use the recipe for shared intention in our future activities. Over time our use of the chosen recipe becomes so familiar that we hardly need to think about it at all.

The possibility of our aiming to follow a recipe associated with any one of the three strategies for shared intention, first explicitly and then with greater skill, suggests that no one strategy can claim to be uniquely correct. Instead, capturing the full range of phenomena involving shared intention will require theories associated with several different strategies.

The extent to which we actually succeed in following a recipe may be quite limited, much as our stated preferences alone may explain only a small part of our behavior (see section 2). What matters for our purposes, however, is just that the aim of living out the theory is not entirely inert. It should influence some of our thoughts and actions. This is what makes it reasonable to accept, in the absence of overriding reasons, that the theory could be true of us.

It may be helpful to consider possible responses to this position. One response starts with the observation that following a recipe together may involve us having a shared intention to do so. Probably, then, not all shared intention is a consequence of our intentionally living out a theory of shared intention. We might therefore be motivated to search for phylogenically or

ontogenically foundational forms of shared intention (see, for example, Pacherie, 2013; Rakoczy & Tomasello, 2007; Tollefsen, 2005). Perhaps it would even be possible, eventually, to relate strategies and theories to different stages and needs. This is a radical response which breaks from the most extensively developed, best defended theories currently available, which give no such importance to evolutionary or developmental considerations.

An alternative response would aim to distinguish recipes that humans actually follow, noting that these may be fewer than those which could in principle be used.[18] There are at least two potential sources of inspiration for this response. One is narrowly philosophical attempts to establish that aggregate or plural subjects of shared intention are either practically indispensable or required for certain explanatory purposes (Roth, 2014 presents both kinds of argument, for example.). The other source of inspiration for this response could be taken from Gomez-Lavin and Rachar (2019), who offer findings which they interpret as showing that everyday thinking involves distinctive features of Gilbert (2013)'s account of shared intention. Just here we encounter a dilemma. Narrowly philosophical arguments may establish that humans do follow one recipe but appear unlikely to show that they do not also follow other recipes. On the other hand, taking inspiration from experimental research involves a radical departure from the kinds of consideration usually taken to motivate a theory of shared intention. This is clear from responses to Gomez-Lavin and Rachar (2019)'s work, which include Löhr (2022) who challenges their interpretation on methodological grounds and Michael and Butterfill (2022) who offer apparently contrasting findings. Those authors' interest in discovering how people actually think about joint activities has no counterpart in the work of the leading philosophers.

A bolder and more orthodox response might be to allow that we could coherently follow any of the recipes but deny that all of them yield shared intention (or, more ambitiously, even that any do). The challenge for proponents of this response is to identify grounds for rejecting the view that following the recipes yields genuine shared intention. They would need to enable all of us, as researchers, to know which things a theory of shared intention should explain independently of our knowing which theory is true. As things stand now, philosophers typically use particular examples of joint activities to introduce the topic (see section 1, and Bratman, 2014, pp. 5–6, for example). The usual assumption is that the examples are sufficient to identify "shared activity of the sort we are trying to understand" (Bratman, 2014, p. 6). But intentionally following one of the recipes is, of course, one way of walking together, playing a piano duet, or painting a house together. So if examples of joint activities provide us as researchers with a common understanding of the things to be explained, that common

understanding supports the view that more than one strategy's recipes are needed to capture them. Opponents of this view need further theory-independent ways of identifying what is to be explained.

Overall, it seems plausible that at least three different strategies for shared intention are needed. This is because each of the plural, aggregate and reductive strategies is associated with a recipe people could intentionally follow and thereby manifest phenomena for characterizing which the corresponding strategy is needed.

9. Conclusion

I have explored three strategies for elucidating ideas about shared intention. The plural, aggregate and reductive strategies are often regarded as competing attempts to characterize a single target. Proponents of the plural and aggregate subject strategies typically object that the reductive strategies fail, while proponents of reductive strategies aim to show that the other strategies are not needed (or, if they are needed, that they can be tacked on to a reductive strategy; see 1). Despite much effort, no such arguments currently enable us to determine which strategy is correct. In contrast, I have appealed to the possibility of intentionally living out different theories to argue that for each strategy there are phenomena which can be correctly characterized only by following that strategy (section 6). There may also be many theories derived from the reductive strategy for which the same is true (section 6).

If correct, this conclusion marks a collective success. Whereas it was initially assumed that at most one strategy would work, the careful development of multiple theories by their various proponents suggests there are multiple theoretically coherent possibilities, each having intuitive appeal to some.

Despite the success, this conclusion is not a tenable stopping point. Whereas progress surely requires that we can discover grounds to reject some theories, the conclusion that we need multiple strategies seems to imply that anything goes in constructing theories of shared intention.

What to do? The conclusion that we need multiple strategies rests on the premise that we need at least one theory of shared intention. One way to avoid it might be to reject this premise (see footnote 2 on page for discussions which may motivate considering this option). But while questioning the premise may lead to fresh insights, the idea that we might completely do away with theories of shared intention seems unpromising insofar as some existing theories have been fruitfully applied beyond philosophy.[19]

A more hopeful response to the untenable conclusion would be to take inspiration from other domains where multiple apparently incompatible approaches have been discovered. We might draw a very inexact parallel

with the twin possibilities of using sets to replace plural quantification and of using plural quantification to construct sets (Linnebo, 2022, sec. 4.3). Perhaps – so the hopeful response – we can show that the strategies for shared intention yield theories which are in some sense equivalent ways of elucidating a single set of ideas about shared intention.[20] Or perhaps the informal nature of the theories and their diversity means that this response is just too hopeful.

A simpler response is also available. Several researchers have pointed to things which stand in need of explanation and which, they suggest, might be explained using a theory of shared intention. These include behavioral and neuroscientific findings (Gallotti & Frith, 2013), patterns in cognitive development (Tomasello & Carpenter, 2007), and decision making (Sugden, 2000). In some cases this has led to debate on which things theories of shared intention are supposed to explain (see, for example, Bratman, 2014 on Gold & Sugden, 2007). One way to make further progress would be, in offering a theory about shared intention, to specify things which stand in need of explanation in a way that can be understood independently of the theory's truth or falsity; and to formulate the theory in such a way that makes it possible to determine, eventually, whether it does actually explain those things.

To illustrate how this might go, consider a relatively easy family of questions. How do various groups of individuals represent the activities of some agents acting together in particular situations? Instead of interpreting existing theories as claims about how shared intention is, we can also interpret (or usefully misinterpret) them as theories about how people represent situations involving shared intention. Generating predictions from existing theories is difficult but there are signs that this might be possible.[21] Relative to this project – that of discovering how individuals represent joint activities – the existence of many theories is not a bad thing. After all, there may well be differences between species, between infants, children and adults, and between cultures. Further, a single individual may adopt different models in different situations. Diversity in the theories of shared intention may enable us to discover genuine diversity in the things to be explained.

In conclusion, we researchers need, but lack, a common understanding of what theories of shared intention are theories of. It has been fruitful to construct different models of how aspects of shared intention might be. The next step is to find out which models explain which things.

Notes

1. Where I use "aggregate", others have used "collective" (Björnsson & Hess, 2017, p. 274) and "group" List and Pettit (2011, p. 74). Although more familiar, I have

avoided these terms because they seem to me to risk inviting confusing aggregate subjects and plural subjects.

2. This is not an exhaustive division of strategies. It would also be possible to reject the idea that we need a notion of shared intention at all. Although beyond the scope of this essay, there are interesting discussions which may motivate considering this view in Baier (1997), Chant (2007), Petersson (2007), and Longworth (2019, 13ff).

   It is also not the only way of dividing strategies. Schweikard and Schmid (2021, sec. 3) offer a division into "content-, mode-, and subject-accounts of collective intentionality." These cut across the division into plural, aggregate and reductive. Both the plural subject and the aggregate subject strategies yield "subject-accounts" while the reductive strategy would yield content-accounts.
3. See List and Pettit (2011, p. 33) on "joint intentions" and Bratman (2022a, 135ff) on "institutional intentions".
4. Gilbert (2009, p. 175) does offer a theory which appears to involve stipulations about ordinary agents' views. But this is not supported by investigation, nor is there any explicit suggestion that an investigation would support the theory. Others have attempted to investigate aspects of how well Gilbert's theory captures ordinary agents' views (Gomez-Lavin & Rachar, 2019; Michael & Butterfill, 2022). But those researchers are careful to distinguish the aims of their investigations from supporting, or refuting, Gilbert's philosophical position. The leading philosophical theories of attitudes and norms are about ways people might reasonably be. It is all entirely hypothetical.
5. See Bratman (2014, p. 8): "This is the continuity thesis. As we might try saying: once God created individual planning agents and placed those agents in a world in which they have relevant knowledge of each other's minds, nothing fundamentally new-conceptually, metaphysically, or normatively-needs to be added for there to be modest sociality."
6. Here I am assuming Ontological Innocence, which is a controversial claim (Linnebo, 2022, sec. 5).
7. Note that this line of objection could be pursued independently of whether intentions are mental states. On some views, intentions are not mental states (Russell, 2018; Thompson, 2008) but having them might nevertheless require having a mind.
8. Alternative motivation for intentions without much in the way of other mental states is offered by Bratman's characterization of what he calls "social-procedural-rule-based institutional Intentions" (Bratman, 2022a, pp. 147–50).
9. I am grateful to an anonymous reviewer for observing that Rovane (1998, 137ff) offers an extended discussion of the possibility that very committed people might, over time, achieve such a high degree of rational unity that "it would not be possible to engage just one of its human constituents separately" (p. 141). (The scenario I am imagining is merely one in which being a plural subject is taken as a normative ideal.)
10. See https://correctiv.org/top-stories/2021/10/21/cumex-files-2/
11. The construction is borrowed from Bacharach (2006) and Sugden (2000), p. I am not claiming that their views require aggregate subjects, only that some of their ideas can be (mis?)used to develop the aggregate subject strategy.
12. Whereas Gold and Sugden (2007) appear to defend their view as the only kind of shared intention, Pacherie (2013) explicitly offers a view on which the aggregate subject strategy and the reductive strategy each characterize forms of shared intention. Also, as none of these researchers present their views as involving aggregate subjects, my suggestion is only that we can use their ideas in pursuing the aggregate subject strategy.

13. See Schiffer (1987, p. 265): "if one were to make a list of all the things philosophers have in mind when they talk of 'theories of meaning or intentional content,' then I would claim that there are no true theories satisfying the descriptions on that list. The questions being asked [. . .] that would require positive theories as answers all have false presuppositions."
14. Individual theorists have expressed both views. For instance, Ludwig (2007) positions his view as characterizing something distinct from Bratman's, while Pacherie (2013) positions her view as a revision of Bratman's. My question is whether interpreting their views contrary to their statements would be theoretically coherent.
15. There is room for uncertainty about whether Gilbert's theory meets the requirement about no fundamental discontinuities. Smith (2015, 55ff) argues that it does, at least "to the extent that Bratman's" does. Bratman (2015, p. 75) objects to this claim on the grounds that "[t]he capacity to participate in the creation of [. . .] plural commitments does [. . .] go beyond capacities that are involved in individual agency."
16. Searle does not use the term "we-intention", which was notably used by Tuomela and Miller (1988) (Those authors credit Sellars (1963), although he does not use exactly that term.). Following Gilbert (2007, p. 33), it has become common to use this term in discussing Searle.
17. Linnebo (2022, sec. 5) identifies this as "the traditional view in analytic philosophy" (which Linnebo does not endorse).
18. I am grateful to an anonymous referee for suggesting this possibility.
19. See, for example, Tomasello and Carpenter (2007), Rakoczy and Tomasello (2007), Moll and Tomasello (2007) and Gräfenhain et al. (2009). Note that these researchers smoosh together incompatible philosophical theories when introducing notions of shared intention. The insights being applied are common to many theories and do not depend on the correctness of any one theory.
20. For an illustration of how this might begin, see Bratman (2014), chap. 6) who investigates how a reductive approach may enable the construction of plural and aggregate subjects.
21. See Gomez-Lavin and Rachar (2019). Although their study is not directly concerned with shared intention, their approach may illuminate how their participants represent joint activities.

Disclosure statement

No potential conflict of interest was reported by the author(s).

References

Alberini, A. (2019). Revealed versus stated preferences: What have we learned about valuation and behavior? *Review of Environmental Economics and Policy*, *13*(2), 283–298. https://doi.org/10.1093/reep/rez010

Bacharach, M. (2006). *Beyond individual choice*. Princeton University Press.

Baier, A. C. (1997). Doing things with others: The mental commons. In L. Alanen, S. Heinamaa, & T. Wallgren (Eds.), *Commonality and particularity in ethics* (pp. 15–44). Palgrave Macmillan.

Bardsley, N. (2007). On collective intentions: Collective action in economics and philosophy. *Synthese*, *157*(2), 141–159. https://doi.org/10.1007/s11229-006-9034-z

Beshears, J., Choi, J. J., Laibson, D., & Madrian, B. C. (2008). How are preferences revealed? *Journal of Public Economics*, *92*(8), 1787–1794. https://doi.org/10.1016/j.jpubeco.2008.04.010

Björnsson, G., & Hess, K. (2017). Corporate crocodile tears?: On the reactive attitudes of corporate agents. *Philosophy and Phenomenological Research*, *94*(2), 273–298. https://doi.org/10.1111/phpr.12260

Blomberg, O. (2016). Common knowledge and reductionism about shared agency. *Australasian Journal of Philosophy*, *94*(2), 315–326. https://doi.org/10.1080/00048402.2015.1055581

Boolos, G. (1984). To Be is to Be a value of a variable (or to Be some values of some variables). *The Journal of Philosophy*, *81*(8), 430–449. https://doi.org/10.2307/2026308

Bratman, M. E. (1987). *Intentions, plans, and practical reasoning*. Harvard University Press.

Bratman, M. E. (1992). Shared cooperative activity. *The Philosophical Review*, *101*(2), 327–341. https://doi.org/10.2307/2185537

Bratman, M. E. (1993). Shared intention. *Ethics*, *104*(1), 97–113. https://doi.org/10.1086/293577

Bratman, M. E. (2014). *Shared agency: A planning theory of acting together*. Oxford University Press.

Bratman, M. E. (2015). Shared agency: Replies to Ludwig, Pacherie, Petersson, Roth, and Smith. *Journal of Social Ontology*, *1*(1), 59–76. https://doi.org/10.1515/jso-2014-0043

Bratman, M. E. (2022a). *Shared and institutional agency: Toward a planning theory of human practical organization*. Oxford University Press.

Bratman, M. E. (2022b). A planning theory of acting together. *Journal of the American Philosophical Association*, *8*(3), 391–398. https://doi.org/10.1017/apa.2021.17

Carson, R. T., Flores, N. E., Martin, K. M., & Wright, J. L. (1996). Contingent valuation and revealed preference methodologies: Comparing the estimates for quasi-public goods. *Land Economics*, *72*(1), 80–99. https://doi.org/10.2307/3147159

Chant, S. R. (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

Dixit, A., Skeath, S., & Reiley, D. (2014). *Games of strategy*. W. W. Norton Company.

Frankfurt, H. G. (1978). The problem of action. *American Philosophical Quarterly*, *15*(2), 157–162.

Gallotti, M., & Frith, C. D. (2013). Social cognition in the we-mode. *Trends in Cognitive Sciences*, *17*(4), 160–165. https://doi.org/10.1016/j.tics.2013.02.002

Gilbert, M. P. (1990). Walking together: A paradigmatic social phenomenon. *Midwest Studies in Philosophy*, *15*, 1–14. https://doi.org/10.1111/j.1475-4975.1990.tb00202.x

Gilbert, M. P. (2007). Searle and collective intentions. In S. L. Tsohatzidis (Ed.), *Intentional acts and institutional facts* (pp. 31–48). Springer.

Gilbert, M. P. (2009). Shared intention and personal intentions. *Philosophical Studies*, *144*(1), 167–187. https://doi.org/10.1007/s11098-009-9372-z

Gilbert, M. P. (2013). *Joint commitment: How we make the social world*. Oxford University Press.

Gilbert, M. P. (2022). A simple theory of acting together. *Journal of the American Philosophical Association*, *X*(X), 1–10.

Gold, N., & Sugden, R. (2007). Collective intentions and team agency. *The Journal of Philosophy*, *104*(3), 109–137. https://doi.org/10.5840/jphil2007104328

Gomez-Lavin, J., & Rachar, M. (2019). Normativity in joint action. *Mind & Language*, *34*(1), 97–120. https://doi.org/10.1111/mila.12195

Gräfenhain, M., Behne, T., Carpenter, M., & Tomasello, M. (2009). Young children's understanding of joint commitments. *Developmental Psychology*, *45*(5), 1430–1443. https://doi.org/10.1037/a0016122

Heider, F., & Simmel, M. (1944). An experimental study of apparent behaviour. *The American Journal of Psychology*, *57*(2), 243–259. https://doi.org/10.2307/1416950

Helm, B. W. (2008). Plural agents. *Noûs*, *42*(1), 17–49. https://doi.org/10.1111/j.1468-0068.2007.00672.x

Jeffrey, R. C. (1983). *The logic of decision* (2nd ed.). University of Chicago Press.

Linnebo, Ø. (2022). Plural quantification. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/cgi-bin/encyclopedia/archinfo.cgi?entry=plural-quant

List, C., & Pettit, P. (2011). *Group agency: The possibility, design, and status of corporate agents*. Oxford University Press.

Löhr, G. (2022). Recent experimental philosophy on joint action: Do we need a new normativism about collective action? *Philosophical Quarterly*, *72*(3), 754–762. https://doi.org/10.1093/pq/pqab070

Longworth, G. (2019). Sharing non-observational knowledge. *Inquiry*, 1–21.

Ludwig, K. (2007). Collective intentional behavior from the standpoint of semantics. *Noûs*, *41*(3), 355–393. https://doi.org/10.1111/j.1468-0068.2007.00652.x

Ludwig, K. (2016). *From individual to plural agency: Collective action*. Oxford University Press.

McGeer, V. (2007). The regulative dimension of folk psychology. In D. D. Hutto & M. Ratcliffe (Eds.), *Folk psychology Re-assessed* (pp. 137–156). Springer.

Michael, J., & Butterfill, S. (2022). Intuitions about joint commitment. *Philosophical Psychology*, 1–16.

Moll, H., & Tomasello, M. (2007). Cooperation and human cognition: The vygotskian intelligence hypothesis. *Philosophical Transactions of the Royal Society B*, *362*(1480), 639–648. https://doi.org/10.1098/rstb.2006.2000

Moore, R. (2017). Gricean communication and cognitive development. *Philosophical Quarterly* forthcoming. pqw049. https://doi.org/10.1093/pq/pqw049

Nagel, J., Mar, R., & Juan, V. S. (2013). Authentic gettier cases: A reply to starmans and Friedman. *Cognition*, *129*(3), 666–669. https://doi.org/10.1016/j.cognition.2013.08.016

Pacherie, E. (2013). Intentional joint agency: Shared intention lite. *Synthese*, *190*(10), 1817–1839. https://doi.org/10.1007/s11229-013-0263-7

Petersson, B. (2007). Collectivity and circularity. *The Journal of Philosophy*, *104*(3), 138–156. https://doi.org/10.5840/jphil2007104329

Rakoczy, H., & Tomasello, M. (2007). The ontogeny of social ontology: Steps to shared intentionality and status functions. In S. L. Tsohatzidis (Ed.), *Intentional acts and intentional facts* (pp. 113–137). Springer.

Roessler, J. (2024). Plural practical knowledge. *Inquiry*, *67*(4), 1–20. https://doi.org/10.1080/0020174X.2020.1787221

Roth, A. S. (2014). Indispensability, the discursive dilemma, and groups with minds of their own. In G. Preyer, F. Hindriks, & S. R. Chant (Eds.), *From individual to collective intentionality: New essays* (pp. 137–162). Oxford University Press.

Rovane, C. A. (1998). *The bounds of agency: An essay in revisionary metaphysics*. Princeton University Press.

Russell, D. (2018). Intention as action under development: Why intention is not a mental state. *Canadian Journal of Philosophy*, *48*(5), 742–761. https://doi.org/10.1080/00455091.2017.1414524

Salomone-Sehr, J. (2022). Cooperation: With or without shared intentions. *Ethics*, *132*(2), 414–444. https://doi.org/10.1086/716877

Schiffer, S. R. (1987). *Remnants of meaning*. MIT Press.

Schmid, H. B. (2008). Plural action. *Philosophy of the Social Sciences*, *38*(1), 25–54. https://doi.org/10.1177/0048393107310877

Schmid, H. B. (2013). Plural self-awareness. *Phenomenology and the Cognitive Sciences*, *13*(1), 1–18. https://doi.org/10.1007/s11097-013-9317-z

Schweikard, D. P., & Schmid, H. B. (2021). Collective intentionality. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University.

Searle, J. R. (1990). Collective intentions and actions. In P. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in communication* (pp. 90–105). Cambridge University Press.

Searle, J. R. (2007). Grice on meaning: 50 Years Later. *Teorema: Revista Internacional de Filosofía*, *26*(2), 9–18.

Sellars, W. (1963). Imperatives, intentions, and the logic of 'ought'. In H.-N. Castaneda & G. Nakhnikian (Eds.), *Morality and the language of conduct* (pp. 159–218). Wayne State University Press.

Setiya, K. (2022). Intention. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University.

Smith, T. H. (2015). Shared agency on Gilbert and deep continuity. *Journal of Social Ontology*, *1*(1), 49–57. https://doi.org/10.1515/jso-2014-0045

Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, *124*(3), 272–283. https://doi.org/10.1016/j.cognition.2012.05.017

Starmans, C., & Friedman, O. (2013). Taking 'know' for an answer: A reply to Nagel, San Juan, and Mar. *Cognition*, *129*(3), 662–665. https://doi.org/10.1016/j.cognition.2013.05.009

Steele, K., & Stefánsson, H. O. (2020). Decision theory. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University.

Sugden, R. (2000). Team preferences. *Economics & Philosophy*, *16*(2), 175–204. https://doi.org/10.1017/S0266267100000213

Thompson, M. (2008). *Life and action: Elementary structures of practice and practical thought*. Harvard University Press.

Tollefsen, D. (2005). Let's Pretend! *Philosophy of the Social Sciences*, *35*(1), 74–97. https://doi.org/10.1177/0048393104271925

Tomasello, M., & Carpenter, M. (2007). Shared intentionality. *Developmental Science*, *10*(1), 121–125. https://doi.org/10.1111/j.1467-7687.2007.00573.x

Tuomela, R., & Miller, K. (1988). We-Intentions. *Philosophical Studies*, *53*(3), 367–389. https://doi.org/10.1007/BF00353512

Zawidzki, T. W. (2013). *Mindshaping*. MIT Press.