

9

Flavors of “Togetherness”

Experimental Philosophy and Theories of Joint Action

Deborah Tollefsen, University of Memphis

Roger Kreuz, University of Memphis

Rick Dale, University of California, Merced

Philosophical Accounts of Joint Action

The philosopher’s interest in joint action stems, in part, from a long-standing interest in understanding the distinction between intentional and non-intentional (non-voluntary) actions, between things that happen to us and things that we do. On a causal theory of action, actions are to be distinguished from mere happenings in virtue of the fact that they are caused by antecedent mental events such as beliefs, desires, and intentions. The fact that our actions are often “joint” (or “collaborative”), and intentionally so, raises the question of what antecedent mental states or events cause joint actions. How are the things we do together intentionally (soccer games) different from things which are the result of individual intentional actions (traffic jams), but for which we would not want to say “We did it together”?

There is a growing body of literature in philosophy that suggests that what is unique about doing things *together* is that it involves a sort of irreducible “jointness” that can be characterized as “we-intention.” Just as individual actions are shaped and informed by individual intentions, joint actions are shaped and informed by joint or shared intentions. Thought experiments are often used to solicit intuitions regarding the

existence of shared intentions. To offer the reader some examples of the use of intuitions in this literature, we extensively quote from the works of two prominent theorists in this area. First, consider the following example used by John Searle to motivate this intuition:

Imagine that a group of people are sitting on the grass in various places in a park. Imagine that it suddenly starts to rain and they all get up and run to a common, centrally located shelter. Each person has the intention expressed by the sentence “I am running to the shelter.” But for each person, we may suppose that his or her intention is entirely independent of the intentions and behavior of others. In this case there is no collective behavior; there is just a sequence of individual acts that happen to converge on a common goal. Now imagine a case where a group of people in a park converge on a common point as a piece of collective behavior. Imagine that they are part of an outdoor ballet where the choreography calls for the entire corps de ballet to converge on a common point. What exactly is the difference? . . . Intuitively, in the collective case the individual intentionality, expressed by “I am doing act A,” is derivative from the collective intentionality “We are doing act A.”

(Searle, 1990, p. 402)

Likewise, we find Margaret Gilbert motivating the existence of shared intention with the following example:

Suppose someone tells me “Ralph and Alice are shopping for clothes.” A variety of situations could be at issue. First, Alice could be in one store buying a dress, while Ralph is in a store five miles distant, buying a suit. Or Alice and Ralph may both be in the same store, each one buying what he or she individually needs. Her shopping is being carried out quite independently of his. In both of these cases, it might be said that Ralph and Alice will not be shopping for clothes *together*. In these cases it would be natural to say: ‘Ralph and Alice are both shopping for clothes.’ Finally, they may be doing something we really would think of as shopping for clothes together. He may sit near by while she looks at dresses, may make suggestions and comments, and may be consulted carefully before the final purchase is made. . . . The main point I want to make in relation to the possible situations covered. . . is that the third case, as described, has a very different flavour from either of the first two cases.

(Gilbert, 1989, p. 154)

These examples are meant to solicit the intuition that doing something together often means more than simply each individual performing an action at the same time and/or within the same vicinity. According to Searle, Gilbert, and many others, there is a “jointness,” a “we-ness,” a “togetherness” present that accounts for the different flavor of certain social interactions. How to understand the nature of this jointness or the

nature of we-intentions that underlie joint action is a topic of much dispute. There are a variety of theories on offer. Our aim here is not to question any particular theory of joint action but to consider the methodology that is universally adopted in this literature—the reliance on “common intuitions” to support analyses of joint action and shared intention. Although philosophers rely on judgments of particular cases in order to build their theories, they do not actually consider the judgments of ordinary people when building these theories. The question remains then whether such theories get at the heart of *our* concepts of jointness, or simply at the heart of the particular philosopher’s concept of jointness. Do ordinary folk actually make distinctions between doing things and doing things together? And do they make these distinctions on the basis of information about the action type or on the basis of some other information? We believe that experimental philosophy provides a useful method for gauging ordinary people’s intuitions about the concept of joint action and that doing so provides some empirical support for philosophical theories that appeal to primitive concepts of togetherness, “we-ness” or “jointness.”

Experimental Philosophy and Collective Intentionality

Philosophers interested in joint action are not the only ones whose methodology relies heavily on intuitions. Indeed, most philosophical theories are motivated by appeal to intuitions described as “common.” In recent years there has been a movement in philosophy that is seen by some to significantly challenge this approach. Experimental philosophy refers to a movement that seeks to contribute to fundamental philosophical questions by running empirical studies of the psychological processes underlying people’s intuitions about central philosophical concepts and issues. The idea is to move beyond a mere appeal to the “common” intuitions or to the intuitions of some academic philosophers. By doing so, theories can actually take into account the ways in which real people think, discriminate between cases, and make judgments concerning the nature of personal identity, moral agency, knowledge, justification, and a host of other notions that are of interest to the philosopher.

One particular strand of experimental philosophy, the strand we are most attracted to, attempts to preserve a role for common intuitions. However, the preservation of common intuitions is in the service of a very different sort of methodology from that used by much of twentieth-century Anglo-American philosophy. Anglo-American philosophy has, at least since the turn of century, been engaged in conceptual analysis which involves a consideration of a number of cases which are then supposed to allow for further refinement of the meaning of certain concepts and inevitably result in a proposed list of necessary and sufficient conditions (involving yet more concepts) for the concept of interest. This has led to a cottage industry of “But what about this!” articles offering counterexamples which challenge either the sufficiency or necessity of the conditions and send the philosopher back to the drawing board.

Conceptual analysis has taken a bit of a beating over the past few decades and there has been a great number of new methodological turns taken—the method of genealogy, for instance. But despite this, conceptual analysis is still prominent in some domains of philosophy.¹

This is particularly so in the field of collective intentionality, where concepts like joint action and shared intention are analyzed in terms of necessary and sufficient conditions that appeal to other concepts thought to be more fundamental. Consider Michael Bratman’s (1999) analysis of shared intention:

We intend to wash the dishes if and only if:

1. a. I intend that we wash the dishes.
b. You intend that we wash the dishes.
2. I intend that we wash the dishes in accordance with and because of 1a and 1b; you intend likewise.
3. 1 and 2 are common knowledge between us.

¹ It might be argued that Bratman, Gilbert, and others are not giving a conceptual analysis but the conditions under which joint action or shared intention exist—an ontic analysis. Gilbert, at least, specifically associates her work with conceptual analysis in her early work, *On Social Facts* (1989). Regardless of what their aim is, they are attempting to provide necessary and sufficient conditions and the sufficiency and necessity of these conditions is tested by thought experiments and judged according to the intuitions such thought experiments supposedly solicit.

As a first approximation, this complex of intentional attitudes seems a plausible candidate. But consider a case in which we each intend to wash the dishes together and we each do so in part because of the other's intention. However, I intend to wash the dishes with Palmolive and you intend to wash them with Joy. All of this is common knowledge and we will not compromise. Is there a collective intention present? It seems not. In this case we do not have our subplans coordinated in the appropriate way. Recall that one of the jobs that shared intention has is to coordinate our individual plans and goals. In the example above, our individual subplans are in conflict and this would prevent us from achieving our goal of getting the dishes washed.

Bratman avoids this counterexample by adding a clause about participants' subplans. It is not necessary that our subplans match, but they must mesh. So, if my subplan is to wash the dishes with Palmolive, and your subplan is to wash them with hot water, and I have no preference about the water temperature, then our subplans mesh though they don't match exactly. But if we have subplans to wash the dishes with completely different types of dish detergent then our subplans do not mesh. Bratman reformulates the account in the following way:

We intend to J if and only if:

1. (a) I intend that we J and (b) you intend that we J
2. I intend that we J in accordance with and because of 1a and 1b, and meshing subplans of 1a and 1b; you intend the same.
3. 1 and 2 are common knowledge.

Various counterexamples have been raised to question the adequacy of this account. Consider, for instance, competitive contexts. Although we play a game of chess together, it isn't clear that our subplans must mesh. Indeed, it would seem that in many cases they do not mesh because we share distinct goals vis-à-vis winning (Bratman, 1999). Consider, also, joint actions that occur on the fly, so to speak. If a group of people spontaneously run to rescue a person from a burning car, it seems implausible that it involves the sort of complex of attitudes Bratman suggests.

Margaret Gilbert has argued that at the core of our notion of sociality (or doing things together in the robust sense) is the concept of a plural subject of belief or action. A plural subject is an entity, or as Gilbert puts it, "a special kind of thing, a 'synthesis *sui generis*'" (1996, 268) formed

when individuals bond or unite in a particular way. This “special kind of thing” can be the subject to which intentional action and psychological attributes are attributed. We can formulate the conceptually necessary and sufficient conditions for the existence of plural subjects in the following way:

Individuals $A_1 \dots A_n \dots$ form a plural subject of X-ing (for some action X or psychological attribute X) if and only if $A_1 \dots A_n$ form a joint commitment to X-ing as a body.

A joint commitment to act as a body is a commitment made by a collection of individuals to perform some present or future action as would a single individual. Joint commitments are formed when each of a number of people expresses his or her willingness to participate in the relevant joint commitment with the others. Each person understands that only when all of the relevant people have agreed to participate in the joint commitment will the joint commitment be formed. Once everyone has agreed, a pool of wills is formed and individuals are then jointly committed. Once the joint commitment is established, each individual is individually obligated to do his or her part to make it the case that he or she acts as part of a body.

Consider a case in which a soccer team agrees to play a scrimmage match against another league team. The members of the team do not each individually agree to play a soccer match. If they did, it would lead to a proliferation of soccer matches. Each member, however, agrees to make it the case that they together play soccer and express their willingness to do so on the condition that every other member do the same. This expression of willingness need not be simultaneous. The members may express their willingness over time. Nor do they need to express their willingness verbally. In many cases, silence is an adequate expression of intention. They must, however, in order for the joint commitment to come into existence, communicate in some way and at some point in time their intention to do their part in playing the scrimmage game as a body (or a team) with others.

Gilbert’s account has also been subject to a number of counterexamples. Bratman (1999), for instance, has argued against the sufficiency of such an account. Imagine a mobster hit-man who points a gun at the head of his rival and tells him that they are going to take “a little ride” *together*. No doubt the rival will express his willingness to go, if only to

delay the inevitable. And, of course, the mobster has already expressed his willingness. But surely such an expression of willingness to go for “a little ride” together does not constitute a joint action on the part of the mobster and his rival. We would not ordinarily say that they were traveling together or doing anything *together*. Would we? Bratman simply assumes the answer to this is no.

Our discussion so far reflects, in many ways, the analysis and criticism that takes place in the area of collective intentionality. In many circumstances, the discussion revolves around particular examples that are generated to test analyses—washing the dishes, playing a game, constructing a house, mobster hit-men going for a ride—and these are supposed to reveal “common” intuitions and support or attack the necessity and sufficiency of certain features. In all cases, however, the appeal to common intuitions is done from the philosopher’s arm-chair. Indeed, though the philosopher will often reference “common” intuitions, there has been no attempt to find out if the postulated intuitions are actually the intuitions of the average person, or simply those of the philosopher him- or herself.² In addition, there have been no explorations of the underlying psychological mechanisms that would constitute or influence such intuitions.

The strand of experimental philosophy that we adhere to takes seriously the idea that common intuitions are relevant to philosophy though they serve a different purpose than those found in traditional philosophical analyses. Knobe and Nichols (2007) provide the following description of the methodology.

Typically, one starts out with a fairly superficial characterization of certain patterns in people’s intuitions [patterns revealed through experiments]. Maybe something like this:

People are more inclined to regard an agent as morally responsible when the case is described in vivid and concrete detail than they are when the case is described more abstractly.

² Although those working in collective intentionality have not appreciated the role of the intuitions of the common folk, there have been a number of recent articles in which experimental philosophers have explored the practice of attributing mental states to groups (Phelan, Arico, and Nichols, forthcoming; Knobe and Prinz, 2008). The focus of these experiments has been on people’s willingness (or lack of willingness in some cases) to attribute mental state attributions to groups. We have drawn inspiration from these articles but our focus here is not on the attribution of mental states to groups but rather on the way that individuals conceptualize joint actions.

The goal, however, is to provide some deeper explanation of why the intuitions come out this way. For example:

People are more inclined to regard an agent as morally responsible when they have strong affective reaction to his or her transgressions.

And ultimately, the hope is that one will be able to arrive at a more fundamental understanding of people’s thinking in the relevant domain. Maybe something like this:

People’s intuitions about moral responsibility are shaped by the interaction of two different systems—one that employs an abstract theory, another that relies on more immediate affective reactions.

(Knobe and Nichols, 2007, p. 5)

Experimental philosophy is often accused of making philosophy into a popularity contest: poll the common people and whatever they say about the concept of knowledge, belief, moral responsibility, and so on is the correct account of that concept. Such a characterization is important to address, but it does miss an important point about experimental philosophy: it is not pitting common intuitions against the philosopher’s intuitions. Rather, the data reveal something about the source of common intuitions and the context in which they arise. The philosopher’s intuitions and analyses can be richly informed by the common origin and flexibility of relevant concepts, and when integrated together they form a new foundation for advancing philosophical theory. Experimental philosophy, at least of the type under consideration here, does not want to replace the philosopher with the “man on the street.” Rather the experimental philosopher is attempting to give philosophers more data to consider when building their theories.

It is not that actual percentages themselves are supposed to directly impact our philosophical inquiries. Rather, the idea is that these experimental results can have a kind of indirect impact. First, we use experimental results to develop a theory about the underlying psychological processes that generate people’s intuitions; then we use our theory about the psychological processes to determine whether or not those intuitions are warranted.

(Knobe and Nichols, 2007, p. 6)

In what follows we provide some of our data on people’s judgments regarding examples of joint action (examples found within the literature) and reflect on what it might contribute to philosophical inquiries about the nature of joint action.

Study 1: “Togetherness” as a Conceptual Feature

As we have already noted, numerous examples of individual or joint tasks have been used in the philosophical debate on collective intentionality. Our goal in this chapter is to use experimental methods to test whether naïve participants perceive the “jointness” of these canonical cases. This problem can be cast as a question of *conceptual structure*, explored in numerous domains in cognitive psychology (e.g., Rips, Shoben, and Smith, 1973). In this area of research, we measure participants’ intuitions about some subset of canonical tasks to assess conceptual or semantic dimensions that underlie these intuitions. For example, if participants compared a set of tasks such as playing poker, dancing a tango, and reading a book (alone), then participants would be expected, over many trials, to evaluate poker and tango as more similar along a target cognitive dimension, namely that they are joint actions. In order to explore people’s intuitions about joint actions, we employed a paradigm that has been used by experimental psychologists for several decades. Multidimensional scaling (MDS) provides a way of identifying underlying dimensions that affect how people think about concepts.

MDS was developed by Shepard (1962) and Kruskal (1964) as a data analysis method based on a geometric model. Concepts are represented as points in a semantic space, and the distance between points reflects the similarity of concepts that they represent. Specifically, MDS attempts to find a best spatial representation that fits all pairs of concepts simultaneously. Although any number of dimensions can be used, solutions are typically mapped onto a two- or three-dimensional space.

As an example, consider thinking about the relative distance between a number of cities. Most people would judge New York and Denver as relatively far apart, and New York and Boston as relatively close together. When research participants are asked to make their judgments on a distance scale, and are asked to rate all *pairs* of a number of cities, the resulting data can be plotted as a two-dimensional solution which captures the best fit of all the rated pairs. When this is done, a rough map of the United States emerges, and the two underlying dimensions can be identified as “east-westness” and “north-southness.” Psychologists have used this method to explore more complex concepts, such as the concept “animal.” When asked to provide ratings for all pairs of a

number of exemplars, a two-dimensional solution capturing “size” and “ferocity” emerges (Henley, 1969).

Although the interpretation of the underlying dimensions can be rather subjective (see, for example, Jaworska and Chupetlovska-Anastasova, 2009), it provides a useful starting point in determining whether people spontaneously think about actions in terms of “jointness.” One of the benefits of this method is that it reveals ways people are thinking that are not easily articulated by the subjects themselves. We find this a particularly interesting method for use in the area of collective intentionality. The terms “joint commitment,” “we-intention,” “shared intention” are technical terms. It is unlikely that the ordinary person would say that these concepts play a role in their discrimination of certain types of actions or in their own social interactions. This makes the philosopher’s burden even heavier as they may be accused of creating concepts rather than exploring or explaining them. MDS, however, helps to reveal the underlying dimensions that inform, at an unconscious level, how people think about concepts. As we shall see, the data we collected using MDS suggests that there is a “unique flavor” of togetherness that subjects seem to be picking up on.

We asked a group of research participants to provide similarity ratings for ten activities that people engage in together. As a first experiment, we wished the first set of instructions to point directly to our target dimension, expecting that such instructions should readily bring about a multidimensional scaling solution with a dimension of “jointness” as a conceptual feature. We included three pairs of related activities that are typically joint or solo actions (“having a conversation” and “talking to oneself”; “singing a duet” and “singing a solo”; and “playing poker” and “playing solitaire”). We also included two other activities that are typically thought of as joint actions (“dancing the tango” and “having sex”), and two activities typically thought of as solo in nature (“driving a car” and “reading a book”).

The participants in the study were recruited via Amazon’s Mechanical Turk (www.mturk.com), a crowdsourcing service that allows “requesters” to recruit “workers” for a variety of online tasks. Participants are paid for their responses, with Amazon functioning as the intermediary between the requester and the workers. Data collected in this way tend to be as reliable as data collected face-to-face from college undergraduates (Munro et al., 2010).

We had 20 participants³ make 45 similarity judgments between all pairs of ten activities described above. To emphasize the concept of “jointness,” the web page viewed by the participants was entitled “Doing Things Together” and in the instructions, they were asked to “rate the similarity of activities that people engage in together.” Ratings were made on a seven-point scale, with endpoints labeled “not at all similar” and “very similar.” The participants were encouraged to look over the first few items before they began to get a sense of the kinds of activities they would be comparing. They were also told that there were no “right” or “wrong” answers, and that we were simply interested in their intuitions.

The data were analyzed using the ALSCAL algorithm in SPSS, with ordinal measurement and Euclidian distance. With this multidimensional scaling method, a two-dimensional solution of participant judgments explained 85.8% of the variance in their responses. If people spontaneously use “jointness” in making judgments about the similarity of activities, then we would expect to see that the stimuli lie along one of the two MDS dimensions in a corresponding manner. Dimension 1 shows just such a distribution. The five joint activities (having sex, dancing the tango, singing a duet, having a conversation, and paying poker) all fall at or below the midpoint on the left side of the solution. Four of the five solo activities (playing solitaire, reading a book, talking to oneself, singing a solo) fall above the midpoint on the right side of the solution, and the fifth, driving a car, is just below the midpoint. It is important to note that the exact coordinates in this multidimensional solution are not meaningful, but only *relative* coordinates between our items; researchers often rotate these solutions to obtain coordinates that fit more with intuition. In fact, researchers often, for convenience and visualization, label the axes. We have done this in Figure 9.1, for illustration.

Within Dimension 1 there is also evidence that people’s judgments reflect a sensitivity to conditions of physical proximity. Those joint actions that require physical touching (having sex and dancing the tango) rather than mere co-location (singing a duet, having a conversation)

³ The participants consisted of 14 females and five males (and one of unreported gender) living in the United States, and their average age was 36 years. They required an average of 6.5 minutes to complete the task.

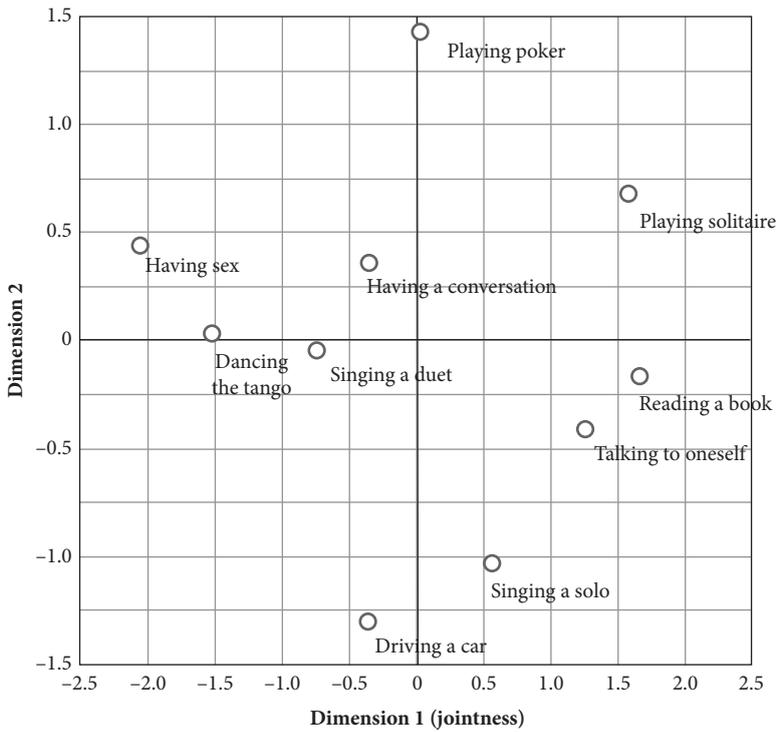


Figure 9.1 Multidimensional scaling and jointness: Study 1.

are further from the midpoint. Likewise, on the right-hand side the solo actions can be seen as falling according to those that are more or less physically isolating. Playing solitaire and reading a book are actions that are likely to be more physically isolating, whereas singing a solo is usually done in the company of others (an audience).

What about Dimension 2? If our participants were thinking about the *content* of the activities, then we would expect to see clustering based on the three pairs of activities that involved similar actions (i.e., singing, talking, and playing cards). However, an inspection of Figure 9.1 shows that the two card-playing activities are plotted far apart from each other in the upper right quadrant (although both are high on Dimension 2). Singing a duet and singing a solo are plotted in different quadrants (and also vary considerably on Dimension 2). Finally, having a conversation and talking to oneself are also in different quadrants, and appear toward

the middle of Dimension 2. Though there are hints of this sensitivity to content, the multidimensional scaling solution did not reveal it as clearly as one would expect.

Our results suggest that experimental philosophers can make use of techniques such as MDS and clustering to explore the dimensions that people spontaneously make use of when thinking about members of a category—in this case, different types of human activities. These results are supportive of the intuition that “jointness” plays a role in the conceptualization of activities. In the experiment described next, we wished to test the robustness of this conceptual dimension. Would participants continue to reveal a conceptual structure that had jointness as a dimension if we changed the instructions to be more indirect about this aspect of our stimuli?

Study 2: “Togetherness” Without Prompting?

Once again, we employed crowdsourcing to obtain intuitions about the similarity of various activities. As in Study 1, we recruited participants⁴ via Amazon’s Mechanical Turk service. The participants provided similarity ratings for the same 10 activities used in Study 1, but in this case, the web page bore a neutral title, “Evaluating Activities,” and they were asked to “rate the similarity of activities that people engage in.” The rest of the instructions and the procedure were identical to Study 1.

As before, the data were analyzed using the ALSCAL algorithm in SPSS, with ordinal measurement and Euclidian distance. As in the case of Study 1, a two-dimensional solution of participant judgments explained 85.9% of the variance in their responses.

As can be seen by comparing Figures 9.1 and 9.2, the results of both studies are highly similar. Once again, Dimension 1 can be explained in terms of “jointness.” The five joint activities (having sex, dancing the tango, singing a duet, having a conversation, and paying poker) all fall below the midpoint on the left side of the solution. All five of the solo activities (playing solitaire, reading a book, talking to oneself, singing a

⁴ The participants consisted of 18 females and 5 males (and two of unreported gender) living in the United States, and their average age was 40 years. They required an average of 4 minutes to complete the task.

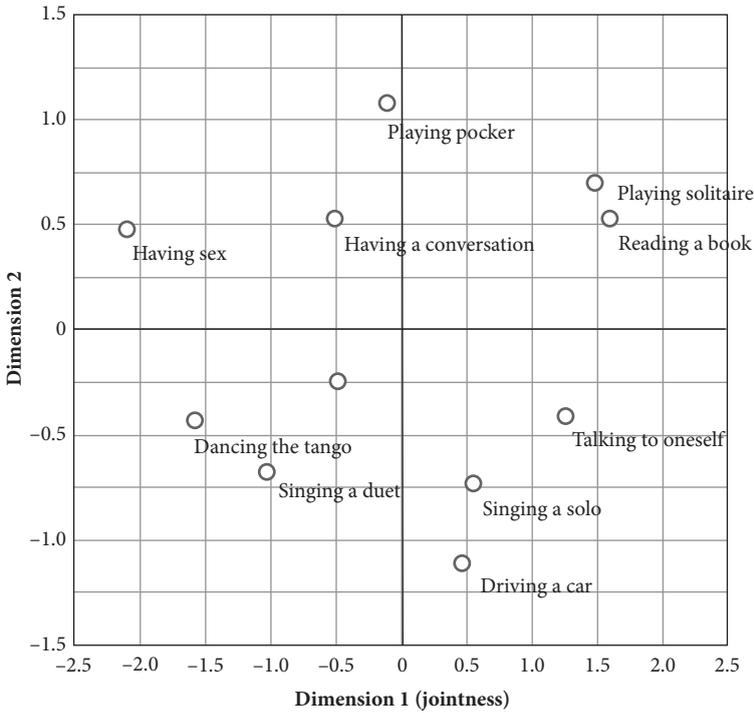


Figure 9.2 Multidimensional scaling and jointness: Study 2.

solo, and driving a car) fall above the midpoint on the right side of the solution.

Also as in Study 1, it is clear that participants are not making judgments based on the content of the activities. For all three pairs involving singing, talking, and playing cards, the joint activities (singing a duet, having a conversation, and playing poker) are in different quadrants than the solo activities (singing a solo, talking to oneself, and playing solitaire). This is a rather surprising result. The differences and similarities between types of actions would seem to be a salient feature but our results suggest that whether the action was done by a single individual or multiple individuals is more salient. The agent (who did it or does it) and not the action (what is done) seems to be a more salient feature for categorization (and this is so even when subjects are not prompted to categorize solo vs. joint actions).

Does the data from Study 2 provide any further insight into Dimension 2? One reviewer has suggested that subjects may be categorizing the activities in terms of the degree to which a successful performance depends upon an individual's performance. Having a successful conversation and successful sex (whatever that might mean) seems to depend on others in a significant way, whereas talking to oneself and singing a solo is under an individual's control and so its success is determined by the individual alone. Unfortunately, this doesn't seem to fit the general layout of Dimension 2. Successful sex does seem to require joint control/responsibility, but so too do dancing the tango and singing a duet and yet they are in the bottom quadrants of Dimension 2.

Again, our current data do not reveal a salient interpretation of Dimension 2, but reflection on this dimension has suggested a number of other questions for further research. For instance, actions might be categorized in terms of whether they are a means to an end or an end in themselves. Dancing, for instance, might be an end in itself, whereas driving a car is a means to some further end. Are intuitions sensitive to this distinction? There are some activities that are stereotypically solitary—not just done by an individual but done alone. They indicate a sort of anti-social behavior—solitaire and talking to oneself might be such cases. Are people making distinctions between doing things alone and doing things with other people present but individually?

Another way of analyzing these data is to employ hierarchical clustering, and this method suggests that the distinction between being alone and completing an action alone may be playing a role in people's judgments. In this method, each score begins as its own cluster, and then pairs and larger groupings emerge until all the data points are represented (Johnson, 1967). The clusters can be displayed in a dendrogram but can also be described verbally. For the set of six solo and joint action terms, and using a single linkage method, "talking to oneself" and "playing solitaire" emerge as the first cluster. We might characterize these actions as anti-social or solipsistic. A second cluster consists of these two terms along with "singing a solo." So the three solo activities can be described as clustering together and categorize actions which are done by an individual but are not necessarily anti-social. For the joint actions, "having a conversation" and "singing a duet" cluster together, so two of these also behave in the expected manner. The final action,

“playing poker,” only joins the other items when a single cluster is created. This is probably because playing poker is somewhat different from the other activities, which are cooperative in nature. Even though cooperation is required to play the game, the activity is ultimately competitive, and this may explain why the clustering algorithm doesn’t place this activity with the others. Varying the examples along cooperation vs. coordination, and cooperative vs. competitive contexts, may also reveal the nature of the second dimension.

The results of Study 2 make clear that the concept of “jointness” is one that people spontaneously use in thinking about the similarity of activities. Whereas the results from Study 1 could be explained by the fact that we explicitly invoked the idea of “doing things together,” no such explanation is available for the data from Study 2. In the absence of an explicit prompt, and with very neutral instructions (“evaluating activities”), our participants’ intuitions about these activities look very much like the intuitions of others for whom “jointness” was explicitly mentioned. It remains to be seen, however, whether the concept of jointness that the folk use is anything like that of the philosopher. Although clearly people are tracking whether the action is done by an individual alone or jointly (together with someone else) and not the content of the action, the conception of togetherness and aloneness might vary along a number of lines—including physical proximity, how anti-social the action is, and whether the action could be done by more than two people.

Study 3: “Togetherness” Without Proximity?

A reviewer of this chapter astutely noted that the items we used for Studies 1 and 2 involved a confound of physical proximity: joint activities correlated almost perfectly with activities that were done in close physical proximity. It is entirely plausible that this may be a criterion for joint action, at least of the canonical. Nevertheless, joint activities may be carried out without close proximity, such as conversation or Internet chatting; and solo activities can be done in close proximity, such as standing in a line or riding on a subway. We modified our items and pursued one final study to test whether jointness emerges out of item

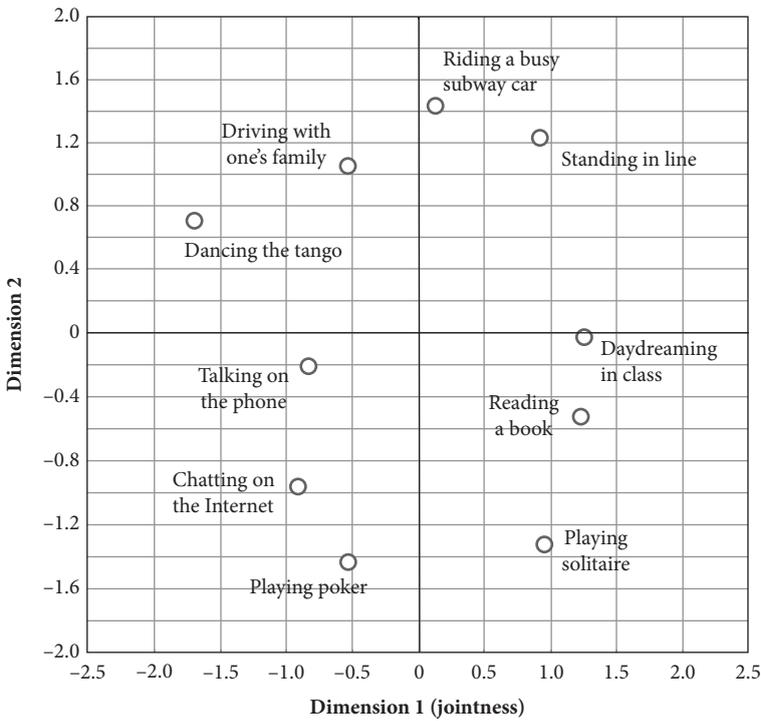


Figure 9.3 Multidimensional scaling and jointness: Study 3.

comparisons, and ran a similar multidimensional scaling solution.⁵ Items and results are shown in Figure 9.3.

Again, a clear pattern emerges across Dimension 1. Activities on the left-hand side of the graph are those done with others and solo activities fell to the right-hand side. Importantly, this pattern holds even for solo activities that are in physical interpersonal proximity (e.g. riding on a subway car) and joint activities done at a distance (chatting on the Internet or phone). Jointness emerges as the dominant feature by which subjects categorized actions. Subjects seem to be tracking the psychological connectivity of actors rather than simply physical

⁵ The participants consisted of 12 females and 16 males living in the United States, and their average age was 37.25 years. They required an average of 4.5 minutes to complete the task.

proximity. This is broadly supportive of philosophical theories of joint action that identify a special “flavor” of togetherness.

Philosophical Applications

What implications do these results have for philosophical analyses of joint action? One of the major debates concerning the nature of joint action is whether an analysis of joint action can be given that does not appeal to something irreducibly joint. Consider, for instance, Bratman’s analysis in terms of the intentions of individuals and their mutual interdependence. This analysis attempts to reduce joint intention to individual intentions which then can, it is assumed, be understood apart from anything joint. The benefit of such an account is that it avoids a charge of circularity, as one can analyze the notion of joint action and intention in terms that do not themselves appeal to notions such as “intending together.” But attempts at such reductions are often challenged. According to Searle (1990), for instance, collective intentionality is a primitive concept and cannot be reduced to a set of individual intentions. Individual intentions to perform one’s part in a joint action are formed on the basis of we-intentions and not the other way around. We-intentions are “primitive” and cannot be analyzed in terms of something more basic such as individual intentional states. Reference to a primitive notion of “doing things together” arises in the work of Margaret Gilbert as well.

I do, of course, posit a mechanism for the construction of social groups (plural subjects of belief or action). And this mechanism can only work if everyone involved has a grasp of a subtle conceptual scheme, the conceptual scheme of plural subjects. Given that all have this concept, then the basic means for bringing plural subject-hood into being is at their disposal. All that anyone has to do is to openly manifest his willingness to be part of a plural subject of some particular attribute.

(1989, p. 416)

The formation of plural subjects requires that the members have a basic or primitive concept of doing things together.

The data we introduce here suggests that ordinary folk do indeed have subtle concept of doing things together and that this concept plays a role in their ability to categorize activities. Far from being a figment of the philosopher’s intuitions, the concept of acting jointly seems to exhibit

itself in simple comparison tasks performed by the ordinary person. Though our results are not conclusive, they do suggest that philosophical analyses that appeal to primitive understandings of working together or intending together are not psychologically implausible and so reference to such primitive concepts in philosophical analyses are actually important in understanding the nature of joint action. It should be noted, however, that our studies fail to identify the nature of the psychological connectivity that subjects seem to track. That is, nothing about our studies reveals whether subjects identify the presence of we-intentions (Searle) or shared intentions (Bratman) or joint commitments (Gilbert) or whether such concepts are being used in categorization.

There are various ways in which the paradigm we introduce might be used to explore the boundaries of the concept of “jointness.” For instance, variation in the types of joint activities, those involving competitive contexts as opposed to cooperative contexts, might yield insights into whether and to what extent people see competitive contexts as genuinely joint. Philosophical analyses of competitive contexts seem to call for something less stringent than Bratman’s meshing of subplans, yet any competitive game requires some sort of collaboration. Our results provide hints of this, as the only competitive task (poker) seemed to behave differently in our cluster analysis. We also might develop experiments that tease out the second dimension along which people seem to be making comparisons. The extent to which the action is joint is a clear dimension that appears in our data both when subjects are prompted and when they are not, and in contexts of physical proximity and no physical proximity. But the second dimension remains unclear. If we vary the content of the action, the extent to which the action requires ongoing interaction, and the extent to which it requires an organizational context, we might be able to identify this dimension more clearly. Future investigations could uncover other flavors of joint activity.

Our results are inevitably preliminary. This application of MDS, in particular, requires getting participants to rate the similarity of every activity pair. This limits the range of stimuli that could be used in our design. For example, with 20 activities, we would require 190 comparisons; 30 activities would require 435. There are ways of getting around these and other issues (e.g., Ramsay, 1982; Ashby et al., 1994), which may allow future investigations to explore a wider range of activities. Our goal here, though, was to use the MDS approach “out of the box” and to

showcase the stability of the concept of joint action in some canonical cases. While these results are a first step in this direction, we hope to have shown that experimental philosophy can contribute in a positive way to philosophical theories of joint action. Not only do our results seem to suggest that naïve subjects are working with a basic notion of joint action, they also suggest that the jointness of an action is particularly salient such that comparisons are made with respect to this feature rather than with respect to the content of the action. Philosophical theories of joint action and intention that make reference to a certain distinct “flavor” of jointness or to a primitive concept of doing things together might find that common folk’s intuitions pick up on this distinct flavor as well. Because an activity’s participants are what sustain it, understanding how these participants conceptualize the activity will shed light on the social stability, or evanescence, of the activity itself.

References

- Ashby, F.G., Maddox, W.T., and Lee, W.W. (1994). On the Dangers of Averaging Across Subjects when using Multidimensional Scaling or the Similarity-choice model. *Psychological Science* 5(3): 144.
- Bratman, M. (1999). *Faces of Intention*. Cambridge, MA: Cambridge University Press.
- Gilbert, M. (1989). *On Social Facts*. New York: Routledge.
- Henley, N.M. (1969). A Psychological Study of the Semantics of Animal Terms. *Journal of Verbal Learning and Verbal Behavior* 8: 176–84.
- Jaworska, N., and Chupetlovska-Anastasova, A. (2009). A Review of Multidimensional Scaling (MDS) and its Utility in Various Psychological Domains. *Tutorials in Quantitative Methods for Psychology* 5: 1–10.
- Johnson, S.C. (1967). Hierarchical Clustering Schemes. *Psychometrika* 32: 241–54.
- Knobe, J. and Nichols, S. (2007). An Experimental Philosophy Manifesto. In Knobe, J., and Nichols, S. (eds), *Experimental Philosophy*. Oxford: Oxford University Press.
- Knobe, J. and Prinz, J. (2008). Intuitions about Consciousness: Experimental Studies. *Phenomenology and the Cognitive Sciences* 7: 67–85.
- Kruskal, J.B. (1964). Multidimensional Scaling by Optimizing Goodness of Fit to a Nonmetric Hypothesis. *Psychometrika* 29: 1–27.
- Munro, R., Bethard, S., Kuperman, V., Lai, V., Melnick, R., Potts, C., Schnoebelen, T., and Tily, H. (2010). Crowdsourcing and Language Studies: The New

- Generation of Linguistic Data. *Proceedings of the NAACL HLT Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*, 122–30.
- Phelan, M., Arico, A., and Nichols, S. (forthcoming). Thinking Things and Feeling Things. *Phenomenology and Cognitive Science*.
- Ramsay, J.O. (1982). Some Statistical Approaches to Multidimensional Scaling Data. *Journal of the Royal Statistical Society. Series A (General)*, 285–312.
- Rips, L.J., Shoben, E.J., and Smith, E.E. (1973). Semantic Distance and the Verification of Semantic Relations. *Journal of Verbal Learning and Verbal Behavior* 12(1): 1–20.
- Searle, J. (1990). Collective Intentions and Actions. In P. Cohen, J. Morgan, and M.E. Pollack (eds), *Intentions in Communication*, Cambridge, MA: Bradford Books, MIT Press.
- Shepard, R.N. (1962). The Analysis of Proximities: Multidimensional Scaling with an Unknown Distance Function. Part I. *Psychometrika* 27: 125–40.