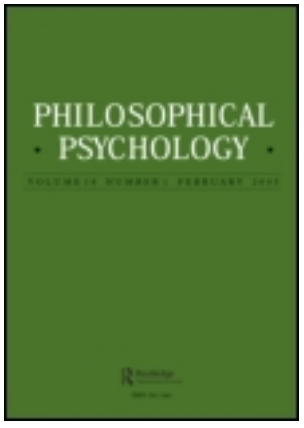

Philosophical Psychology
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/cphp20

Naturalizing joint action: A process-based approach

Deborah Tollefsen & Rick Dale

Available online: 24 Jun 2011

To cite this article: Deborah Tollefsen & Rick Dale (2012): Naturalizing joint action: A process-based approach, Philosophical Psychology, 25:3, 385-407

To link to this article: http://dx.doi.org/10.1080/09515089.2011.579418

Philosophical Psychology
Vol. 25, No. 3, June 2012, 385–407

Naturalizing joint action: A process-based approach

Deborah Tollefsen and Rick Dale

Numerous philosophical theories of joint agency and its intentional structure have been developed in the past few decades. These theories have offered accounts of joint agency that appeal to higher-level states (such as goals, commitments, and intentions) that are "shared" in some way. These accounts have enhanced our understanding of joint agency, yet there are a number of lower-level cognitive phenomena involved in joint action that philosophers rarely acknowledge. In particular, empirical research in cognitive science has revealed that when individuals engage in a joint activity such as conversation or joint problem solving, they become aligned at multiple levels (e.g., behaviors, or cognitive states). We argue that this phenomenon of alignment is crucial to understanding joint actions and should be integrated with philosophical approaches. In this paper, we sketch a possible integration, and draw out its implications for understanding of joint agency and collective intentionality. The result is a process-based, dynamic account of joint action that integrates both low-level and high-level states, and seeks to capture the separate processes of how a joint action is initiated and sustained.

Keywords: Alignment; Collective Intentionality; Conversation; Coordination; Joint Action

Over the past few decades, a number of philosophers have begun to investigate joint agency and its underlying intentional structure. Searle (1990, 1995), Bratman (1993, 1999, 2004, 2006), Gilbert (1989, 1994, 1996, 2003), and Tuomela (1992, 1993, 1995, 2007), among others, have offered accounts of joint agency that appeal to higher-level states (such as goals, commitments, and intentions) that are "shared" in some way. Though these accounts have greatly enriched our understanding of joint agency,

Deborah Tollefsen is an Associate Professor of Philosophy at the University of Memphis.
Rick Dale is an Assistant Professor of Psychology at the University of Memphis.
Correspondence to: Deborah Tollefsen, Department of Psychology, University of Memphis, Memphis, TN 38152-0001, USA. Email: dtollfsn@memphis.edu

ISSN 0951-5089 (print)/ISSN 1465-394X (online)/12/030385-23 © 2012 Taylor & Francis
http://dx.doi.org/10.1080/09515089.2011.579418

there are a number of lower-level cognitive phenomena that are involved in joint agency that philosophers have overlooked. In particular, it has now become evident that when individuals engage in a joint activity such as conversation or joint problem solving they become aligned at a variety of different levels. Their eye movements may become coordinated, their speech patterns more similar, and even their bodily movements can synchronize to some extent. The phenomenon of alignment is, we think, crucial to understanding joint actions and should, therefore, inform philosophical theories of joint action. The aim of this paper is to seek a more naturalized theory of joint action by integrating empirical findings concerning alignment with the philosophical literature on joint agency and collective intentionality.

In section 1, we review the philosophical literature on joint action and identify a number of ways in which current theories remain incomplete or inadequate. Neither goals nor plans can be considered without careful evaluation of the means of accomplishing the former, and of carrying out the latter. In short, any high-level, philosophical theory of joint action should be informed by mechanism-oriented considerations. In section 2, we turn to empirical research on alignment. This research suggests that alignment is found in countless cognitive corners, from the coordination of eye gaze patterns (Richardson & Dale, 2005; Richardson, Dale, & Kirkham, 2007) and low-level behavioral signals such as posture (Shockley, Santana, & Fowler, 2003), to higher levels of linguistic organizations such as syntax (Branigan, Pickering, & Cleland, 2000) and pragmatics (Roche, Dale, & Caucci, in press). In section 3, we introduce the idea of an ''alignment system'' involving lower-level coordinative structures that help to implement higher-level goals (Shockley, Richardson, & Dale, 2009). In section 4, we return to a particular philosophical theory of joint action—John Searle's—and show how it could be augmented by an alignment system to produce a more complete and empirically informed theory of joint action. We also lay out a number of avenues for future research based on this more empirically informed theory of joint action.

1. Philosophical Theories of Joint Action and their Inadequacies

Many of the philosophical theories of joint action on offer today are informed by the ''causal theory of action.'' The causal theory of action proposes that the distinguishing feature of voluntary action is that it is caused by appropriate antecedent mental events and episodes such as desires, beliefs, intentions, and so on. It is a necessary condition for a behavior to be an intentional action that it be caused by a mental event (and caused in the right way), and so to explain action is to specify the prior mental events that are the proximate cause of the action. The question, then, for those who want to explain joint action is: ''what are the prior mental events that are the proximate causes of joint action?''.

A number of philosophers have argued that an explanation of joint action requires appeal to "shared" intentions as the proximate cause of joint actions. Searle provides the following cases to motivate this need:

> 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 . . . are part of an outdoor ballet where the choreography calls for the entire corps de ballet to converge on a common point. We can imagine that the external bodily movements are indistinguishable in the two cases; people running for shelter make the same types of bodily movements as the ballet dancers. Externally observed, the two cases are indistinguishable, but they are clearly different internally. What exactly is the difference? (Searle, 1990)

Many philosophers have suggested that the difference lies in the presence of shared intentions. But there is very little agreement concerning the nature of shared intentions.

Raimo Tuomela (2005) argues that shared intentions involve individuals intending from the "we-mode," and mutual beliefs regarding the existence of such intentions. The we-mode, broadly the perspective one takes as a member of a group, is to be contrasted with the "I-mode," the perspective one takes as an individual. This contrast is designed to differentiate the personal (or "private," a term Tuomela tends to prefer) from the social, and hence to provide a distinguishing characteristic of the social. Though reasoning (or thinking, or intending, or believing) from the we-mode is something individuals do, the we-mode cannot be reduced to the I-mode. Tuomela provides the following analysis of what it is for an individual to we-intend (that is, to intend from the we-mode):

> (*WI*) A member A_i of a collective g *we-intends* to do X if and only if:
>
> (i) A_i intends to do his part of X (as his part of X);
> (ii) A_i has a belief to the effect that the joint action opportunities for an intentional performance of X will obtain (or at least probably will obtain), especially that a right number of the full-fledged and adequately informed members of g, as required for the performance of X, will (or at least probably will) do their parts of X, which will under normal conditions result in an intentional joint performance of X by the participants;
> (iii) A_i believes that there is (or will be) a mutual belief among the participating members of g (or at least among those participants who do their parts of X intentionally as their parts of X there is or will be a mutual belief) to the effect that the joint action opportunities for an intentional performance of X will obtain (or at least probably will obtain);
> (iv) (i) in part because of (ii) and (iii). (2005, p. 340)

Similarly, Searle posits an irreducible we-intention at the heart of joint action. According to Searle, individuals participating in a joint action all have intentions of

the form "we intend to J." How is it possible for an individual to have an intention of the form "we intend to J"? Searle contends that this capacity is biologically primitive and present in a variety of other species. This capacity presupposes other "Background" capacities (the Background is a technical term for Searle which refers to conditions necessary for certain cognitive activities and language). In particular, the capacity presupposes a Background sense of the other as a candidate for cooperative agency (Searle, 1990, p. 414).

Michael Bratman does not see the need to postulate an irreducible we-intention or mode of reasoning. Rather, Bratman argues that the shared intentions underlying joint action can be understood as a state of affairs consisting primarily of individual intentions of the form "I intend that we J" (where J is a joint action) and mutual beliefs regarding these intentions (1993, p. 99). Specifically, Bratman provides the following analysis:

We intend to J if and only if:

(1) 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.

The requirement that subplans mesh is meant to rule out cases in which participants both intend to do some joint action but they intend to go about fulfilling that intention in ways that would ultimately undermine the joint action. Meshing subplans requires that the means I use to complete my part of the joint action will not prevent you from completing your part of the joint action.

Finally, for Margaret Gilbert, shared intentions essentially involve a joint commitment. We intend to J if and only if we jointly commit to J-ing as a body (1989, 2003). Joint commitments are formed when each individual expresses his or her willingness to be so committed with others. Joint commitments are not individual commitments to do one's part. Rather, a joint commitment is the commitment of a group. It cannot be dissolved without the agreement of all parties and the existence of joint commitments brings about obligations and entitlements. Like Searle and Tuomela, who argue that we-intentions (or intentions formed from the we-mode) cannot be reduced to I-intentions, Gilbert's joint commitments are irreducible to individual commitments.

The details of these theories need not concern us here. What all these accounts share in common is the fact that they originate from the philosophical armchair. Though we agree that a theory of joint action needs to appeal to mental states like intentions, and we are persuaded by the arguments that these intentions should be "shared" in some manner, we think philosophical accounts of joint action should be informed by what is actually taking place within and between individual cognitive agents. Just as action theory, the theory of individual action, has been informed by the cognitive science of perception and action, so too a theory of joint action needs to be empirically informed. A "naturalized" theory of joint action is needed. The tendency of philosophical theories to provide highly idealized theories of joint action

and to ignore empirical research on human interaction results in at least five distinct problems: the execution problem, the implementation problem, the multiple initiation problem, the verifiability problem, and the over-intellectualization problem. Not every theory of shared intention and joint action suffers from all five problems, but we believe most theories suffer from at least one or more of them to some degree.

Execution problem. The focus in the philosophical theories of joint action has been, for the most part, on planning rather than execution of action. Bratman and Gilbert, for instance, are concerned with the intentions or commitments that are formed prior to the action. But given the connection between action and cognition, and action and perception, philosophical theories will remain incomplete without an understanding of how joint actions unfold over time. Many philosophical accounts remain empirically skeletal because they cannot address the dynamic cognitive updating that must occur in real-time during execution. Searle's we-intentions are intentions-in-action rather than prior intentions, but still we are given no account of how we-intentions interact with motor intentions, for instance, or how we-intentions are updated over time during execution of joint actions.[1]

Implementation problem. Philosophical theories of joint action focus almost entirely on higher-order cognitive states such as commitments, goals, and intentions. However, we know that in the case of individual intentional agency there are a number of lower-level cognitive phenomena that underlie such agency, such as perceptual processing, motor intentions, cognitive maps, categorization, and so on. Likewise, joint agency involves a number of lower-level phenomena, including joint attention and various alignment mechanisms (see below). Without an understanding of these lower-level phenomena, philosophical theories will remain incomplete. A similar point has been made by those who argue that high-level theories of cognition need to be informed by the science which identifies how cognition is implemented, specifically neuroscience (Bechtel & Mundale, 1999). Multiple-realizability has often been a premise in an argument that attempts to dismiss lower-level theories as irrelevant to understanding the nature of mental states and other cognitive phenomena. If pain, for instance, is multiply realized (in dogs, humans, and possibly aliens) then the way it is realized is irrelevant for understanding its nature. In response to this line of reasoning, Bechtel and Mundale (1999) argue for the centrality of mechanism in understanding how a particular process is carried out, making high-level theories crucially informed by finer-grained, low-level accounts.

We make a similar point here. Philosophical accounts of joint action seem to assume that implementation is a non-issue. The implicit reasoning seems to be something like this: provided high-level conditions are met (appropriate shared intentions, etc.), it is irrelevant what specific cognitive processing phenomena give way to them.[2] But more and more research in cognitive science suggests that the very goals and plans of eventual action depend upon specific motor intentions. That is, the means of carrying out a plan to accomplish a goal inherently require low-level phenomena of perceptual and bodily processes. In the case of joint action and

attention, this connection is especially stark. In short, the adequacy of philosophical theories of joint action depends on their being informed by empirical research. But the reverse holds as well: low-level accounts and high-level accounts are mutually constraining (McCauley & Bechtel, 2001). It is only by integrating both kinds of account that we will be able to adequately explain the capacity to engage in joint activities.

Multiple initiation problem. Joint action is initiated in diverse ways. Its initiation may come not only from goals and/or intentions but from lower-level phenomena that kick start joint action. Consider, for instance, conversation. Two people can form intentions and share intentions to have a conversation. They can plan to meet to discuss the candidates for a job position. But they can also engage in conversation spontaneously as a result of some lower-level emotional responses, bodily cues, or simply the vagaries of happenstance in time and physical location. The same conversation could be initiated by a shared emotional response to the candidate's job interview, or by even lower-level phenomena such as noticing shared attention. Philosophical accounts tend to focus on joint actions that come about by a conscious and planned manner, and many of them attempt to provide necessary and sufficient conditions for joint action (or shared intention) and hence rule out the possibility of joint action arising in different and less cognitively complex ways. Now one might argue that unless or until a joint intention (or shared intention) is in place, there is no joint action and so the mechanisms which initiate joint action and give rise to shared intention are not, themselves, important for a theory of joint action. But this seems to narrow the phenomenon of joint action prematurely. Action, individual or joint, is a *process* and the fact that the process is initiated in multiple ways is important for a more complete account of the phenomenon.

Verifiability problem. How do we arbitrate between existing philosophical theories? The philosophical literature on joint action has burgeoned over the past decade and a great deal of effort has gone towards developing counterexamples to proposed theories. This has led to ever more complex theories. What would help arbitrate between existing theories is if these theories could be operationalized in a way that would generate empirical results.[3] But in order to operationalize these theories we need to know more about the underlying mechanisms involved in acting together. Once we identify these mechanisms we can manipulate them, along with social context, in order to understand when and where shared intentions come on board.

Over-intellectualization problem. Joint actions can be performed by different sorts of beings. Animals, young children, adults, and artificial agents can engage in joint action. Because philosophical theories are modeled on joint actions among human adults they are often unable to account for joint action among different sorts of agents. If, for instance, joint action requires prior planning or joint commitments involving normative demands, it is unclear how animals could be able to engage in joint action, but surely they do. Philosophical theories of joint action tend to over-intellectualize joint agency. The psychological complexity of these analyses (consider, for instance, Tuomela's theory above) does not extend easily to non-adult

human and non-human joint action. It may be objected that to require a philosophical theory to explain all types of joint agency (across multiple types of agents) is itself to ''over-intellectualize'' these issues. We agree that this is a heavy theoretical burden. But we think that to focus exclusively on the adult human case has led to a rather myopic view of joint action and that the goal should be to understand joint action in all of its guises. Further, philosophical theories that focus on the adult human case do not lend themselves well to evolutionary perspectives. That is, even if these accounts explain how certain types of joint actions are performed (namely, those involving adult humans), we are left with no explanation of the way in which these joint actions are related to more basic forms of joint agency found in non-humans or how such complex forms evolved from joint actions found in more ''primitive'' versions in the history of our species.

These problems may be meliorated by recourse to empirical work on joint action. In the next section, we turn to a number of recent studies that suggest that when individuals engage in joint endeavors (including conversation[4]) there are a number of lower-level cognitive phenomena going on. In particular, there appears to be alignment across a broad range of levels of processing. The existence of diverse alignment processes suggests that joint agency rests on a more basic capacity to align with others.

2. Empirical Research on Alignment

Research on cognition and action has revealed that these two phenomena—often studied in isolation—are in fact intricately interconnected (e.g., Barsalou, 2008; Glenberg, 1997; Jeannerod, 2006; Varela, Thompson, & Rosch, 1992). For example, in recent work, the dynamic execution of action has been shown to still carry characteristics of the underlying cognitive processes generating that execution. McKinstry, Dale, and Spivey (2008) showed that when participants are evaluating fuzzy sentences such as ''is murder sometimes justifiable?'', their arm trajectories (which executes their decision) dynamically showcases the patterns of uncertainty that may underlie their decision itself. Arm trajectories, measured through computer-mouse usage, will be more or less ''wiggly'' or ''curved'' between affirmative or negative responses to fuzzy statements depending on the participants' level of cognitive certainty (see also Dale, Kehoe, & Spivey, 2007; Spivey, Grosjean, & Knoblich, 2005). This sort of action co-variation effect holds in low-level systems, such as the allocation of attention and visual processing (Song & Nakayama, 2009; Tipper, Lortie, & Baylis, 1992), and in even high-level processes, like social judgment and evaluation (e.g., Freeman, Ambady, Rule, & Johnson, 2008; Wojnowicz, Ferguson, Dale, & Spivey, 2009). In short, this shows that goals, decisions, and the execution of related actions, at root, could be more fundamentally part of the one overlapping, continuous process (Spivey, 2007). They can only be teased apart as an approximation.

While this work shows that cognition can flow smoothly into action, there is also evidence that the reverse holds: the conditions and constraints of execution can feedback onto cognition, and influence the decision prior to the execution itself. For example, action-compatibility studies have shown that the position of your body, such as moving the arm away from the body to make a response, can facilitate the interpretation of sentences, such as "the man threw the ball." The compatibility between potential action can facilitate cognition before and while it unfolds (Glenberg & Kaschak, 2002; as other examples see Tucker & Ellis, 2001, for perception; Ross, Wang, Kramer, Simons, & Crowell, 2007, for categorization; Grant & Spivey, 2003; Thomas & Lleras, 2007, for problem solving; and even modulation of emotion by facial muscle placement, see Larsen, Kasimatis, & Frey, 1992; see Rueschemeyer, Lindemann, Elk, & Bekkering, 2009, for a recent review). Thus cognitive states, such as goals and decisions, are not independent of their relevant action execution. This and other research shows that action feeds back to influence cognition, which (as reviewed above) smoothly flows back into action. During this cognition-action cycle the world may subtly (or less subtly) change the conditions for action as the cognitive system acts; thus one cannot purely separate these systems and expect a full account of how goals/actions together unfold.

If joint actions are based on two such cognitive systems coming together to achieve a goal, then the same insights hold. Cognition and action in the context of joint agency cannot be easily separated. Another person becomes the context in which action constraints and conditions take shape. Here we consider the joint activity of conversation (Clark, 1996), perhaps one of the most common joint actions. Several studies have demonstrated that the emerging cognitive processes supporting conversation are ones of "alignment" (Pickering & Garrod, 2004). 'Alignment' simply refers to the dynamic matching between behavioral or cognitive states of two people. Shockley et al. (2003) have shown that body posture aligns during naturalistic conversation, potentially generated by subtle matching of verbal cues (Shockley, Baker, Richardson, & Fowler, 2007), which also align during interaction (Giles, Coupland, & Coupland, 1991; see also Chartrand & Bargh, 1999). Through such low-level bodily alignment, Pickering and Garrod (2004) maintain that the joint activity of conversation succeeds most fluidly when representation and processes align across various levels of linguistic organization, from words to choice of sentence structure (e.g., passive versus active voice; Bock, 1986; Branigan, Pickering, & Cleland, 2000). Even at the highest levels of linguistic organization—such as figurative language usage—it appears that conversants exhibit alignment, such as strategic use of irony or sarcasm (Roche et al., in press). Recently there has been a significant growth of interest in identifying such fundamental mechanisms underlying joint action (Galantucci & Sebanz, 2009).[5]

As such alignment processes unfold during an interaction, it should be the case that the action output of the system should also align. As just noted, this may be the case with posture (Shockley et al., 2003) and other rhythmic behaviors (e.g., Richardson, Marsh, Isenhower, Goodman, & Schmidt, 2007; Schmidt, Carello, & Turvey, 1990), but other systems also exhibit this alignment. Richardson et al. (2007)

have demonstrated that during interaction, there is a tight coupling of visual attention. As people discuss a work of art, their eye movements become distinctly aligned in time. Indeed, it seems to be the case that the better the alignment, the better the participants are understanding each other or, in other words, the better they succeed in fulfilling the shared goal or intention of communicating with one another (Richardson & Dale, 2005; see also Ireland, Slatcher, Eastwick, Scissors, Finkel, et al., forthcoming; Tanenhaus & Brown-Schmidt, 2008). This alignment of attention, or joint attention, in conversation suggests that joint action, in general, involves aspects of an alignment "system."

What this brief survey of empirical research suggests is that when we glance at the details of the underlying cognitive system as joint actions are enacted, the details are deeper than high-level theories let on.[6] Action and cognition mutually constrain each other, and continue to do so as they cut across members of a potential joint activity, including mutual constraints from one person to another, on the very actions that may be selected to successfully carry out this activity. It seems to us that a more complete account of joint action requires an understanding of the fundamental regularities at these lower levels.

3. A Dynamic Theory of Joint Action

Consider an example of a joint activity that is seemingly complex but may be initiated in multiple ways: having a conversation. When standing at a bus stop near a stranger, a happenstance orientation of one's body towards another may prompt a curt statement about the weather, the time, or the bus schedule. Simply an inadvertent locked glance can be a low-level cue that may initiate a linguistic exchange. For those few seconds that this is occurring, the joint conversational activity has only a minimal "cognitive contract," a phrase that describes the high-level intentional structure often proposed by philosophers. Of course, the contractual nature of the exchange may grow more sophisticated as time goes on. What we are arguing is that the causal story is complicated by lower-level cognitive processes, and that a high-level account, such as those offered by philosophical accounts, by itself does not address such processes. This results in a lack of detail regarding how actions during joint activity unfold (execution problem), what cognitive or other processes are employed in the activity (implementation problem), and what starts the activity (multiple initiation problem). In a sense, these problems can be seen as consequence of the last problem, over-intellectualizing the necessary and sufficient conditions for joint action. By doing so, high-level theories abstract away from crucial causal components of joint activities. As we have noted, such abstractions may nevertheless be important characterizations of joint actions of particular kinds, under particular circumstances, and at particular times. The abstractions cannot, however, explain the origin and fate of joint actions, and their mechanistic make-up as they unfold—a causal story that is needed for a more complete understanding. In this section, we offer an account that aims to integrate both low-level and high-level aspects of joint

action, potentially bridging philosophical accounts to growing empirical findings in cognitive science.

The basis of our account is what we will call an ''alignment system'' present in humans. As reviewed in the previous section, there is considerable evidence that we can become rapidly behaviorally entrained along a variety of dimensions while engaging in joint tasks like conversation. We will remain agnostic regarding the exact nature of the alignment ''system.'' It may be a ''system'' only in the sense that it is a loosely interconnected set of cognitive processes that have evolved to facilitate the carrying out of group activities. For example, the presence of a mirror neuron system in humans has obtained much empirical support (Rizzolatti & Craighero, 2004), and may be the basis for shared understanding of actions across two people observing each other (and consequences of actions; Bekkering et al., 2009). This substrate could partly underlie our capacity to map our own actions onto the observation and understanding of others' actions by employing overlapping neural hardware. In addition, priming theories of alignment (Pickering & Garrod, 2004) have recently proposed that cognitive accessibility of certain behaviors (e.g., a chosen sentence structure) is induced by hearing another person use it, thus increasing the likelihood of one's producing a related behavior. This alignment by priming predicts a gradual unfolding of shared states between two people exchanging information. As one final example, dynamical systems theorists have used the concept of ''coordinative structure'' to explain how a large number of degrees of freedom (e.g., in muscle groups) are not centrally controlled but rather self-organized into coherent, functional units (e.g., dancing the jig versus throwing a baseball). The problem is one of reducing degrees of freedom through gradual, emergent mutual constraint among parts of the body (Turvey, 1990). In a similar fashion, it may be that joint tasks induce gradual mutual constraint across two or more people's bodily and cognitive states. This predicts that there is no central cognitive contract that is sufficient to produce joint activity. Instead, a joint activity is an emergent, self-organizing phenomenon produced through ''coordinative structures'' of two or more people (Shockley et al., 2009).[7] Indeed, there are other proposals and theoretical conditions that offer distinct accounts of shared lower-level states (e.g., among many, Gottman, Murray, Swanson, Tyson, & Swanson, 2005; Hatfield, Cacioppo, & Rapson, 1993; Marsh, Richardson, & Schmidt, 2009; Sebanz, Knoblich & Prinz, 2003; Warner, 1992).

If all these processes are working during joint action, then the alignment system is a heterogeneous mix of diverse components entraining two or more people to specific patterns of behavior. It is important to note that each example of an alignment component is relatively low level in nature. For example, the mirror neuron system may directly map perceived actions to one's own potential action execution, a rapid blend of self and other that needs no contractual inducement. Additionally, the priming account needs only the accessibility of particular mental states to become aligned during, for example, conversation. This can happen entirely unconsciously without the involvement of high-level contracts. So in joint actions where the alignment system is a crucial starter and sustainer of ongoing activities,

whatever ''cognitive'' contracts emerge, they may be better described as ''meta-cognitive'' contracts. Meta-cognition is the set of processes that track other aspects of cognition, such as self-regulation strategies that lead to adjustments in what one is doing cognitively during complex tasks (e.g., studying). Meta-cognitive processes are complex, self-monitoring, and seem to require a longer timescale to operate (Flavell, 1979). The detailed characterizations of cognitive states in the philosophical theories are deeply meta-cognitive: they posit mental states that themselves refer to the goals and plans of oneself and a partner. We would argue, from an evolutionary perspective, that any such meta-cognitive, contractual understanding of a group's behavior, as it relates to one's own, requires low-level processes to be present in the first place. Without any understanding of another's actions, or an ability to entrain to each other's mental or behavioral states, meta-cognitive contracts would have no mechanistic anchoring (no *cognitive* about which to be *meta*). Put simply, using Dennett's well-known terminology, they would be meta-cognitive ''skyhooks'' (Dennett, 1995).

Embracing the notion of an alignment system as a central underpinning draws out some theoretical implications for accounts of joint action. We express these implications as two distinctions that should be clarified in any philosophical or scientific account of joint action.

Surface synchrony versus deep commitments. Joint action is underlain by simple patterns of ''surface synchrony'' at its lowest level, while the deep commitments—goals, or intentions, those identified by philosophical theories—may emerge from these low-level patterns. These deep commitments, once present, may also ''trickle down'' to get surface synchrony kick-started for the purpose of joint action. That meta-cognitive commitments can initiate joint activity seems indisputable. Yet, there is also considerable evidence that surface synchrony can guide collective dynamics of a group's behavior. In non-human group behaviors, research on emergent crowd behaviors suggests that simple surface features constrain individuals to carry out large-scale self-organized group behaviors (e.g., in schools of fish; Huth & Wissel, 1992). Something similar may hold in human behavior. For example, many are familiar with the experience of ''joint following,'' where subtle directional cues during locomotion from one place to another may cause two people to wander off to an unintended location that was never contractually agreed upon simply by ''following'' each other through mutual, surface-synchronous constraints. In either case, it may be that the initiation of joint activities is supported by both, and in initial moments each may mutually constrain and shape the other as the higher-level deep commitments take form, and surface synchrony is established.

Starting versus sustaining. Joint actions are processes and not products. They must be initiated in some manner, and continually sustained by adaptive cognitive processes as members of a group and the environment change. We argue that these two aspects of joint actions are importantly separable in accounting for joint actions, as they (starting or sustaining) may be carried out by distinct mechanisms—either by automatic, rapid alignment components, or by continued meta-cognitive strategy or executive control (or both).

We propose that these distinctions lend insights into joint action by providing a *process-based* account of the initiation and maintenance of such an action. Consider figure 1. We portray these two distinctions as two axes of a plot reflecting the flow of current behavior between two dyads. Each reflects different manifestations of joint action that exemplify the separation between these two dimensions of activity. In the grayed regions, t_0 reflects the point of initiation. The top-left panel of the figure is a caricature of an unfolding joint action; it is initiated (by surface synchrony; SS) until it is sustained by both meta-cognitive "deep commitment" (DC) and low-level alignment processes together. The top-right panel represents a joint action the origin of which is a deep commitment, initiated by a high-level agreed-upon strategy. The bottom-left and bottom-right panels reflect potential joint action scenarios in which initiation and sustenance recruit differing systems. For example, in the bottom-left, a joint action may start with a joint commitment to engage in shared activity, but after it is initiated lower-level systems take over to guide the group behavior. The top-left region may characterize a conversation initiated by low-level glances, or happenstance body position, that quickly induces more high-level meta-cognitive "contracts" akin to philosophical theories. The top-right region reflects the canonical joint action scenarios often painted by theories of joint action. The bottom-left may reflect joint activity carried out by skilled dancers, initiated by commitment but

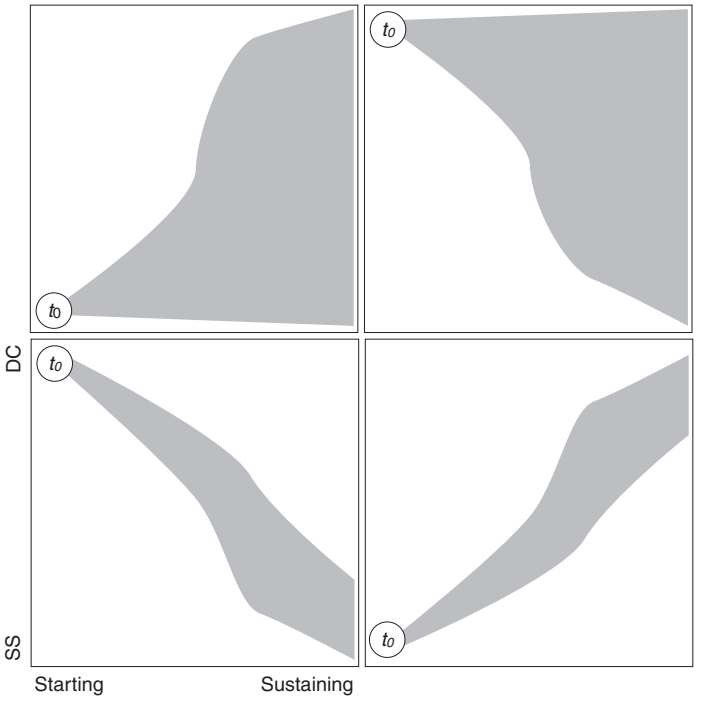

Figure 1 Process-based caricatures of hypothetical joint actions. A joint action may unfold from starting to sustaining by recruiting a variety of cognitive means, either low-level surface synchrony (SS), or higher-level cognitive contracts ("deep commitments"; DC).

sustained by skilled automatic processes that do not require constant cognitive contract. The bottom-right panel may reflect an argument between a couple that is sustained after a lower-level initiation, but the commitment conflicts with low-level alignment systems that do not help sustain the activity.

These two dimensions characterize the *process* of joint action. They integrate philosophical accounts (as they relate to deep commitments) with the assemblage of alignment subsystems that guide planning and execution of behaviors occurring in members of a joint action. This broader explanatory canvas does not seek sufficiency or necessity conditions for an intuition regarding a fixed concept called ''joint action,'' but instead seeks to characterize what it is that creatures do when they behave jointly. The account thus permits a more heterogeneous mix of behaviors to fit into the framework, permitting ''joint action'' to be a fuzzy concept—after all, joint action is a behavior of creatures in a contingent, probabilistic world, dependent upon cognitive systems that are partial and gradient in their nature (Spivey, 2007). In the next section, we discuss how it may fit with a prominent theory of joint action from philosophy, that of Searle.

4. Philosophical Applications

We have been arguing that philosophical theories need to be informed by empirical research. In particular, we have argued that joint action involves a lower-level alignment system and that this system plays a significant role in initiating and sustaining joint action. We have not argued, however, that joint or shared intentions do not play a role in joint agency. Rather, we have argued that these higher-level cognitive states are only part of the story. In this final section we discuss how this alignment system might augment Searle's particular philosophical account of joint action.

We begin by returning to Searle's cases involving groups of people, one a random collection and the other part of a ballet performance. As Searle notes, there seems to be no difference in their bodily movements, but internally they are very different. Searle's, and other philosophical accounts, suggest that the difference lies in the existence of we-intentions. But positing this as the definite mark of a joint action leads one to think of joint agency as something miraculous, as if the presence of a shared intention were a sort of switch that is turned that instantaneously transforms the same bodily movements into a joint action. This simply cannot be the only difference between these two groups. We can see this by considering a case in which joint action emerges from lower-level processes. The random collection of people may have no prior intention to run for shelter together, but in running together they may become perceptually coupled, and this perceptual coupling may give rise to other forms of alignment which might ''kick start'' we-intentions (large effects of such seemingly minimal alignment are well known in the behavioral sciences, such as the bandwagon and herding effects; Raafat, Chater, & Frith, 2009). Now one might reply that until we-intentions are present there is no joint action, but to do so is to

gloss over what may be the necessary conditions for joint action to occur and the multiple ways in which joint action might emerge from lower-level processes. Indeed, Searle himself would seem to agree. For Searle, we-intentions presuppose the Background. Searle appeals to the Background in his theories of intentionality and language understanding as well. He posits such a Background to avoid having to presuppose tacit rule following or unconscious rules or states that guide behavior. The Background, according to Searle, is non-intentional. He defines it as a set of non-intentional or preintentional capacities that enable intentional states of function. Searle suggests that these capacities are neurophysiological, but they can be described at a higher level. The Background that is presupposed in joint agency involves a sense of the other as a potential partner in coordinated activity: "collective Intentionality presupposes a Background sense of the other as a candidate for cooperative agency; that is, it presupposes a sense of others as more than mere conscious agents, indeed as actual or potential members of a cooperative activity" (Searle, 1990, p. 414).

Searle doesn't tell us very much about this "sense" of the other. One way to conceive of this set of capacities, however, is to understand them as structures or features of an alignment system. We-intentions may be necessary for joint action, but as we have argued, they are not sufficient (Bratman implicitly accepts this through his requirement of "meshing"; see Bratman, 1993). The alignment system provides the necessary structure in which we-intentions can be formed and sustained. Figure 2 sketches in a very preliminary manner the sort of multi-dimensional theory we are describing. It is modeled on the multi-level alignment theory of Pickering and Garrod (2004), but generalized across all systems. Admittedly, this is merely a caricature of the unfolding synchronization that may take place across a range of levels during joint actions, with the highest-level of organization the meta-cognitive we-intentions that sometimes initiate and contribute to sustaining (as indicated). We realize this is a simplification of the kind of interactivity that may take place within the system, along with idealized assumptions of modularity and so on which are justifiably questionable in their purer forms; but it is intended to illustrate that a joint activity and joint agency are underlain by a whole range of levels of analysis that research has shown may become coupled during group behavior.

The ballet troupe's higher-order we-intentions will inform their lower level processes and explains how their perceptual and motor systems can function together to achieve their goal. Similarly, the presence of an alignment system explains how we-intentions can be formed on the fly, so to speak, without prior planning or agreements. As a random collection of people run to find shelter they may become aligned on a variety of levels and this alignment can give rise to more sophisticated meta-cognitive states such as we-intentions. It is important to emphasize as well that an empirically informed theory of joint action will highlight the ways in which joint action is a dynamic process with various feedback loops happening across higher and lower levels.

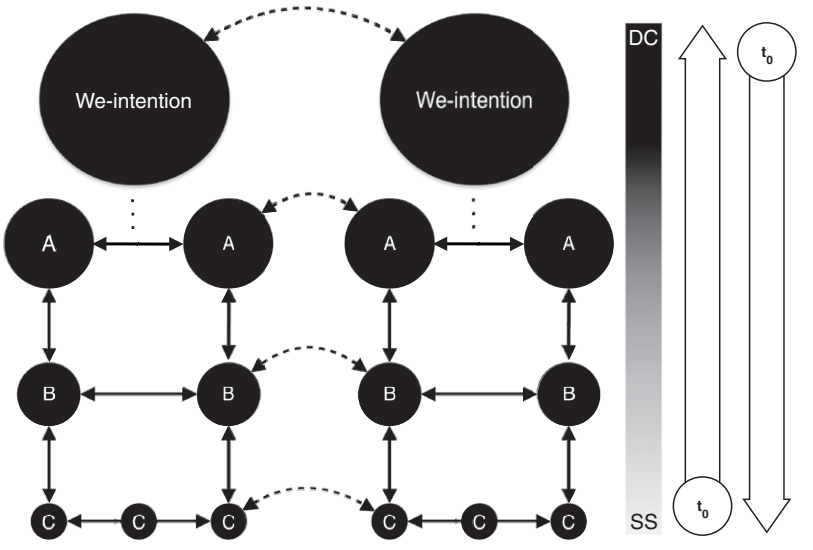

Figure 2 Two agents (left and right) become coupled at a variety of dimensions during interaction. These channels may be, for example, (A) psycholinguistic, (B) attentional, and (C) perceptuo-motor levels. The notion of deep commitments and surface synchrony can be approximated using this sort of idealization of the processes involved in joint action. The right-most arrows again illustrate that the starting and sustaining of joint actions may occur across the range of these levels rather than always at the canonical philosophical intentional end of the scale. Figure motivated by mechanistic theory of psycholinguistic alignment by Pickering and Garrod (2004).

Augmenting Searle's theory in this way allows us to overcome the inadequacies of standard philosophical accounts. Because Searle's account focuses on intentions-in-action rather than prior intentions, it provides the basis for a theory that can help us to understand how joint actions are executed. Thus it could help resolve the execution problem we raised for standard philosophical theories. Coupled with an alignment system that helps coordinate motor intentions and is responsive to we-intentions, we now have a better picture of how joint actions unfold over time and how various cognitive levels interact to sustain joint actions. The model allows us to conceive of joint action as processes rather than products and identifies a number of different levels that need to be explored in any given joint action. Specifically, our framework allows us to introduce the notion of joint action motor plans (or intentions). Just as there are motor intentions which control and monitor the movements of our body in fulfilling an intention, we propose that there are specific shared motor intentions which will determine the movement of bodies in a joint action. When we reach for the cup, our actions are limited to movements appropriate to that motor plan. Likewise, shared motor intentions, those that make reference to a joint action plan, will guide and help monitor the individual actions that make up the joint action. Deep commitments may trigger the implementation of

certain motor plans but likewise, they may be triggered by lower-level alignment processes.

The model we are proposing allows for the fact that joint agency can be initiated from bottom up and so handles the problem of multiple initiation. Deep commitments are clearly not sufficient for joint agency to be successful, and though we-intentions of some form seem to be the mark of joint *agency* (rather than just collective behavior), our theory suggests that these deep commitments can be the end result of alignment processes at the lower level. That is, joint agency is often the result of organization at the lower level rather than conscious prior plans at a higher level (Shockley et al., 2009).

Augmenting Searle's theory in this way also allows Searle to respond to one of the standard criticisms of his account and helps us to deal with the problem of over-intellectualization. Unlike Bratman and others, Searle does not require that we-intentions be mutually constraining. That is, he does not require that participants in a joint action have we-intentions that are interdependent in the way that Bratman or Gilbert require. It has seemed to many that the interdependency of the intentions involved in joint agency are exactly what makes them *shared* intentions. What is needed is something that coordinates minds and bodies over time, and Searle's potentially solipsistic we-intentions seem unable to do this. As Meijers puts it:

> The idea of sharing intentional states is underdeveloped in Searle's internalist conception of collective intentionality. What could this sharing be? The existence of several tokens of the same type of we-states in different individuals may be the only possibility, given Searle's framework. Such an account of sharing, however, has been criticized by Michael Bratman, Margaret Gilbert, Raimo Tuomela, and David Velleman, among others. For an intention to be shared, it is not enough that intentions are coincident.... Some form of unified agency is necessary. (2003, p. 175)

But if we augment Searle's theory with a lower-level alignment system it explains why the presence of we-intentions in individual minds can result in the coordination of minds and bodies over time without those we-intentions themselves having to be interdependent in the way that Bratman or Gilbert require. That is, the unified agency is a function of alignment processes at the lower level. This is important, as theories of joint agency that require that participants be aware of the existence of we-intentions (or joint commitments, or shared intentions of the form "I intend that we j") in other minds will have considerable difficulty in extending their account to young children and animals, who lack a robust theory of mind. Thus, augmenting Searle's theory in the way we suggest also helps us to address the over-intellectualization problem. Because Searlean we-intentions are biologically primitive, we-intentions are capable of being formed by animals and non-adult humans, and hence we don't run into the sort of multiple realizability problems standard philosophical accounts run into.

We leave open the possibility that joint agency in animals is different from joint agency in human beings (Bechtel & Mundale, 1999). It may very well be a less

sophisticated form of agency. Likewise, adult joint action may be a more sophisticated form of joint action than the non-adult human kind. What our theory shows is how various levels of cognition can work together to produce more complex forms of joint action. Animal joint agency may be a matter of various alignment systems and coupled motor intentions that coordinate animal bodies across space and time. This surface synchrony may be all there is to animal joint action in certain cases, whereas human joint action may involve deep commitments like we-intentions or prior commitments and plans. If we view joint action as sustained and produced by a multi-layered system then we can begin to see how such a multi-layered system developed over time to create more sophisticated forms of joint agency. Thus, an augmented Searlean theory of the type we develop here provides the framework for an account of human joint action that is more responsive to evolutionary considerations. We get a picture of how human adult joint action may have developed out of a rudimentary alignment system, and there is no risk of over-intellectualization.

We should note here that we do not wish to run the risk of underemphasizing the importance of higher-level commitments, and thus do not wish to ''under-intellectualize'' the problem of joint action.[8] Some joint actions may very well require the sorts of sophisticated apparatus that Tuomela and Bratman suggest. Perhaps such deep commitments are an additional layer in our dynamic system, permitting far more complex forms of joint activity and joint agency (such as in ''contracts'' more prosaically construed). But they cannot be the defining characteristic in a fuller causal story.

Augmenting Searle's theory with an alignment system also avoids the implementation problem we raised for standard philosophical theories. What happens at the lower levels will constrain what occurs at the higher levels, and if certain activities are being implemented in specific ways at the lower level then this will have an impact on the sort of higher-level intentional states that are guiding the joint action. There may very well be different alignment systems involved in different sorts of joint actions and this may have results for the sort of we-intentions we think are present at the higher level. Philosophical theories have tended to think of joint action as a merely a single coherent top-down phenomenon. What the Searlean theory we are proposing suggests is that joint action often originates from lower-level alignment and that the systems involved might ultimately determine the kind of deep commitments one finds at the higher level.

Finally, an empirically informed theory of the sort we describe opens up a number of different avenues for future research. Here we consider four possible directions that may further bridge philosophical theories with lower-level alignment research: mutual amplification, dimensional compensation, misalignment-needs-commitments, and an illusion of we-will. Each is based on the orthogonality of the two dimensions we described in the previous section (surface synchrony/deep commitments versus starting/sustaining), and on investigating the manner in which these dimensions interact during joint action.

Mutual amplification. If we are correct and an alignment system underwrites all joint action, then manipulation of aspects of the alignment system should have an impact on deep commitments and vice versa (similarly to proposals for psycholinguistic levels in Pickering & Garrod, 2004). Specifically, if one induces alignment at the lower level, one should see participants more likely to exhibit we-intention behavior. There should be more expressions of working together, feelings of solidarity, and so on (there is recent evidence that this is the case; Hove & Risen, 2009). Likewise, if deep commitments are initiated in terms of prior planning we should be able to design experiments that test whether the presence of deep commitments (say, an agreement to work together on a project) actually produces more alignment at the lower level compared to instances which have no such prior commitment. Interestingly, mutual amplification may shed light on the familiar social process of groupthink. It predicts that groups that are more aligned at the lower level will be more likely to exhibit the characteristics of groupthink.

Dimensional compensation. Since our theory predicts that alignment at different levels will make a difference to the sorts of deep commitments that arise, fixing certain aspects of the alignment system will allow us to identify when deep commitments come on board and specify more clearly the conditions under which they arise. For instance, it may be that when a group of individuals is prevented from aligning at the lower levels they feel the need to initiate commitments and plans at the higher level in order to figure out how to coordinate and guide their actions (as also described in Shockley et al., 2009). This suggests a second prediction from the account, in which low-level synchrony would not be available:

Misalignment-needs-commitments. Deep commitments may be necessary in some cases in order to identify ways in which individuals will become misaligned in order to complete parts of a joint activity. That is, deep commitments can often serve as a way of distributing the labor. Spontaneous alignment through surface synchrony does not support situations in which complementary or reciprocal actions are required, such as carrying an object together. In such cases, it may be necessary for the cognitive system to transform an aligned reference frame into a different one that is more conducive to a joint product or achievement. Because this is in direct ''violation'' of isomorphic alignment, it may be necessary for higher-level strategic commitments to modulate planning and execution. These scenarios may be readily developed in order to explore the relative explicitness of cognitive strategies (e.g., meta-cognitive strategies) in such complementary/reciprocal scenarios.

Illusion of we-will. If it is the case that low-level synchrony can induce higher-level deep commitments, then joint actions may also be subject to what Wegner has called the ''illusion of conscious will'' (Wegner, 2003). In a wide variety of studies, it has been demonstrated that a coincident stimulus (e.g., movement of a computer cursor), within certain relative temporal parameters, will induce a sense of conscious willing of having caused that stimulus to move, even if the temporal relations make it *impossible* that the cursor movement was so-willed. This suggests that lower-level perceptuo-motor patterns may be key to higher-level concepts such as willful action. Similarly, the inducement of (say) a conversation via lower-level cues or happenstance variables in

the laboratory may induce deep commitments that participants attribute to illusory intentions to have done so in the first place. In these cases, low-level alignment systems, anchored to surface synchronies in contingencies, could lead to illusory cognitive contracts of the kind described in philosophical theories. This ''illusion of we-will'' (since the we-will is an effect rather than a cause of the activity) would suggest that some joint actions, as our theory predicts, are indeed started by lower-level synchrony. Along with the previous prediction (misalignment-needs-commitments) the theory predicts two quite different aspects of higher-level deep commitments. In the former case, they are necessary components of performing complex non-aligned patterns of behavior or joint activities; in the illusion of we-will, they are epiphenomena, misattributed to original intentions.

5. Conclusion

Philosophical theories of joint action have gone to impressive lengths to articulate the high-level meta-cognitive commitments that underlie idealized examples of joint action and agency (e.g., washing the dishes *together*, preparing Hollandaise sauce *together*, performing in a ballet *together*, etc.). The complex constraints that have been identified are no doubt an important piece to understanding our capacity to engage in joint activities and agency. We have argued here that they cannot be the whole story, and that current philosophical accounts tend to suffer from five fundamental problems. By articulating a dynamic account of joint activity, through a heterogeneous alignment ''system'' that spans lower- and higher-level processes, these problems can all be addressed. By integrating this perspective of alignment with Searle's account, it appears one can achieve the best of both theoretical worlds, from high-level philosophical theories, to lower-level accounts of the cognitive processes involved in joint activities.

Acknowledgements

Preparation of this manuscript was supported by a grant from the National Science Foundation: BCS-0826825.

Notes

[1] Tuomela (2007) introduces a notion of we-willing into his account of joint action, and so appears to have a notion of intention-in-action. We thank a reviewer for pointing this out to us.

[2] Indeed, a reviewer of this article suggested that there ought to be a division of explanatory labor: philosophers should be concerned only with the normative and conceptual requirements of joint action, and psychology and cognitive science ought to worry about the underlying causal story. It is important to note that not all philosophers would agree to such a description of what philosophy is or should be doing. But even if we agree that there is or ought to be such a division of labor this does not entail that the two types of approaches cannot be brought together to provide a more complete account of joint action. It is this for

which we argue. The current paper is an attempt to integrate the causal story with the philosophical story, not replace the causal story with the conceptual story. We do believe, however, that a better understanding of the causal mechanisms involved in joint action will have an impact on the philosophical theories and, equally important, we think psychological research ought to be informed by philosophical theories.

[3] We are not suggesting that one cannot arbitrate between these theories on conceptual grounds. Rather, being able to generate empirical results via psychological research would provide additional means for arbitrating between existing theories.

[4] One reviewer of this article has suggested that to allow conversation to count as a joint action is to reject shared intention as a necessary condition for joint action, for there appears to be no joint intention or shared goal in the conversation one has with, for instance, a stranger at a bus stop. Instead, there may be numerous individual goals (e.g., ''to get the time of day''). If we allow that there are joint actions without shared goals or intentions then is it a consequence that some actions are ''joint'' that seem unintuitive, such as an interrogation, a boxing match, two pendulums on a wall? There are several things to note in response to this concern. Conversation is taken by many psychologists (Clark, 1996) and philosophers (Gilbert, 1989) to be a paradigm case of joint action. It has been suggested that the shared goal or intention in conversational contexts is ''that we communicate.'' What the conversation example brings up, however, is that joint agency, like individual agency, is a gradient phenomenon. There are actions that we do as individuals that are more or less intentional and deliberately executed. Compare riding a bicycle with wooing a lover. The former requires much less deliberative agency than the latter but both are things that we do (rather than things that happen to us). Likewise, there are some joint actions that are executed with less deliberation and conscious monitoring than others. Conversation may be more like riding a bike and less like wooing a lover. Finally, because we think agency (both individual and joint) is best thought of as a continuum with mere automatized action selection at one end and planned, deliberate, and consciously monitored action at the other end, we feel the demand for necessary and sufficient conditions on joint agency is unproductive.

[5] Diverse alignment patterns have been identified in literatures beyond psycholinguistics, such as in the study of adaptation and influence between interaction partners (Cappella, 1996; Warner, 1992) where other terms have been used, such as entrainment, synchrony, complementarity, reciprocity, etc. The literatures relevant to alignment could use more extensive integration (Sadler, Ethier, Gunn, Duong, & Woody, 2009). Here we focus on alignment because it has enjoyed some recent emphasis in the empirical literature on linguistic interaction.

[6] There is a deep literature on these issues, even just in the study of discourse and psycholinguistics. For example, how much alignment may be based on simple memory cuing processes, or higher-level coordinative mechanisms relating to so-called ''audience design,'' has generated considerable debate (Brennan & Hanna, 2009; Brown-Schmidt, 2009; Horton & Gerrig, 2005; Shintel & Keysar, 2009). These can make for very different explanatory tales of a joint action like conversation.

[7] Another related theory inspired by dynamical complex systems is to conceive of cognitive performance as emerging from an autocatalytic process of *mutual* constraint among perception and action. In autocatalysis, a system's parts feedback on each other to produce new behaviors of the overall system, such that standard linear tales of cause and effect are only approximations of the complex feedback dynamics displayed. In the same way, a radical proposal would be to conceive of joint activities, like conversation, as being driven by an autocatalytic process, an inter-personal perceptuomotor feedback cycle, which engenders higher-level commitments as an emergent phenomenon from this autocatalysis (see, for example, Chemero & Turvey, 2008).

[8] We thank a reviewer for pointing out this concern.

References

Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology*, *59*, 617–645.

Bechtel, W., & Mundale, J. (1999). Multiple realizability revisited: Linking cognitive and neural states. *Philosophy of Science*, *66*, 175–207.

Bekkering, H., Bruijn, E. R., Cuijpers, R. H., Newman-Norlund, R., Schie, H. T., & Meulenbroek, R. (2009). Joint action: Neurocognitive mechanisms supporting human interaction. *Topics in Cognitive Science*, *1*, 340–352.

Bock, J. K. (1986). Syntactic persistence in language production. *Cognitive Psychology*, *18*, 355–387.

Branigan, H. P., Pickering, M. J., & Cleland, A. A. (2000). Syntactic co-ordination in dialogue. *Cognition*, *75*, 13–25.

Bratman, M. E. (1993). Shared intention. *Ethics*, *104*, 97–113.

Bratman, M. E. (1999). *Faces of intention: Selected essays on intention and agency*. Cambridge: Cambridge University Press.

Bratman, M. E. (2004). Three theories of self-governance. *Philosophical Topics*, *32*, 21–46.

Bratman, M. E. (2006). *Structures of agency: Essays*. Oxford: Oxford University Press.

Brennan, S. E., & Hanna, J. E. (2009). Partner-specific adaptation in dialog. *Topics in Cognitive Science*, *1*, 274–291.

Brown-Schmidt, S. (2009). Partner-specific interpretation of maintained referential precedents during interactive dialog. *Journal of Memory and Language*, *61*, 171–190.

Cappella, J. N. (1996). Dynamic coordination of vocal and kinesic behavior in dyadic interaction: Methods, problems, and interpersonal outcomes. In J. H. Watt & C. A. Van Lear (Eds.), *Dynamic patterns in communication processes* (pp. 353–386). Thousand Oaks, CA: Sage.

Chartrand, T. L., & Bargh, J. A. (1999). The chameleon effect: The perception-behavior link and social interaction. *Journal of Personality and Social Psychology*, *76*, 893–910.

Chemero, A., & Turvey, M. T. (2008). Autonomy and hypersets. *Biosystems*, *91*, 320–330.

Clark, H. H. (1996). *Using language*. Cambridge: Cambridge University Press.

Dale, R., Kehoe, C., & Spivey, M. J. (2007). Graded motor responses in the time course of categorizing atypical exemplars. *Memory & Cognition*, *35*, 15–28.

Dennett, D. C. (1995). *Darwin's dangerous idea: Evolution and the meanings of life*. New York: Simon & Schuster.

Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, *34*, 906–911.

Freeman, J. B., Ambady, N., Rule, N. O., & Johnson, K. L. (2008). Will a category cue attract you? Motor output reveals dynamic competition across person construal. *Journal of Experimental Psychology: General*, *137*, 673–690.

Galantucci, B., & Sebanz, N. (2009). Joint action: Current perspectives. *Topics in Cognitive Science*, *1*, 255–259.

Gilbert, M. (1989). *On social facts*. Princeton, NJ: Princeton University Press.

Gilbert, M. (1994). Remarks on collective belief. In F. F. Schmitt (Ed.), *Socializing epistemology: The social dimension of knowledge* (pp. 235–236). Lanham, MD: Rowman & Littlefield.

Gilbert, M. (1996). *Living together: Rationality, sociality, and obligation*. Lanham, MD: Rowman & Littlefield.

Gilbert, M. (2003). The structure of the social atom: Joint commitment as the foundation of human social behavior. In F. Schmitt (Ed.), *Socializing metaphysics: The nature of social reality* (pp. 39–64). Lanham, MD: Rowman & Littlefield.

Giles, H., Coupland, J., & Coupland, N. (1991). *Contexts of accommodation: Developments in applied sociolinguistics*. Cambridge: Cambridge University Press.

Glenberg, A. M. (1997). What memory is for: Creating meaning in the service of action. *Behavioral and Brain Sciences*, *20*, 41–50.

Glenberg, A. M., & Kaschak, M. P. (2002). Grounding language in action. *Psychonomic Bulletin & Review*, *9*, 558–565.

Gottman, J. M., Murray, J. D., Swanson, C., Tyson, R., & Swanson, K. R. (2005). *The mathematics of marriage: Dynamic nonlinear models*. Cambridge, MA: MIT Press.

Grant, E. R., & Spivey, M. J. (2003). Eye movements and problem solving: Guiding attention guides thought. *Psychological Science*, *14*, 462–466.

Hatfield, E., Cacioppo, J. T., & Rapson, R. L. (1993). Emotional contagion. *Current Directions in Psychological Science*, *2*, 96–99.

Horton, W. S., & Gerrig, R. J. (2005). The impact of memory demands on audience design during language production. *Cognition*, *96*, 127–142.

Hove, M. J., & Risen, J. L. (2009). It's all in the timing: Interpersonal synchrony increases affiliation. *Social Cognition*, *27*, 949–960.

Huth, A., & Wissel, C. (1992). The simulation of the movement of fish schools. *Journal of Theoretical Biology*, *156*, 365–385.

Ireland, M. E., Slatcher, R. B., Eastwick, P. W., Scissors, L. E., Finkel, E. J., & Pennebaker, J. W. (forthcoming). Language style matching predicts relationship initiation and stability. *Psychological Science*.

Jeannerod, M. (2006). *Motor cognition: What actions tell the self*. Oxford: Oxford University Press.

Larsen, R. J., Kasimatis, M., & Frey, K. (1992). Facilitating the furrowed brow: An unobtrusive test of the facial feedback hypothesis applied to unpleasant affect. *Cognition & Emotion*, *6*, 321–338.

Marsh, K. L., Richardson, M. J., & Schmidt, R. C. (2009). Social connection through joint action and interpersonal coordination. *Topics in Cognitive Science*, *1*, 320–339.

McCauley, R. N., & Bechtel, W. (2001). Explanatory pluralism and heuristic identity theory. *Theory & Psychology*, *11*, 736–760.

McKinstry, C., Dale, R., & Spivey, M. J. (2008). Action dynamics reveal parallel competition in decision-making. *Psychological Science*, *19*, 22–24.

Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, *27*, 169–190.

Raafat, R. M., Chater, N., & Frith, C. (2009). Herding in humans. *Trends in Cognitive Sciences*, *13*, 420–428.

Richardson, D. C., & Dale, R. (2005). Looking to understand: The coupling between speakers' and listeners' eye movements and its relationship to discourse comprehension. *Cognitive Science*, *29*, 1045–1060.

Richardson, D. C., Dale, R., & Kirkham, N. Z. (2007). The art of conversation is coordination: Common ground and the coupling of eye movements during dialogue. *Psychological Science*, *18*, 407–413.

Richardson, M. J., Marsh, K. L., Isenhower, R. W., Goodman, J. R., & Schmidt, R. (2007). Rocking together: Dynamics of intentional and unintentional interpersonal coordination. *Human Movement Science*, *26*, 867–891.

Rizzolatti, G., & Craighero, L. (2004). The mirror-neuron system. *Annual Review of Neuroscience*, *27*, 169–192.

Roche, J. M., Dale, R., & Caucci, G. (in press). Doubling up on double meanings: Pragmatic alignment. *Language and Cognitive Processes*.

Ross, B. H., Wang, R. F., Kramer, A. F., Simons, D. J., & Crowell, J. A. (2007). Action information from classification learning. *Psychonomic Bulletin & Review*, *14*, 500–504.

Rueschemeyer, S., Lindemann, O., Elk, M. V., & Bekkering, H. (2009). Embodied cognition: The interplay between automatic resonance and selection-for-action mechanisms. *European Journal of Social Psychology*, *39*, 1180–1187.

Sadler, P., Ethier, N., Gunn, G. R., Duong, D., & Woody, E. (2009). Are we on the same wavelength? Interpersonal complementarity as shared cyclical patterns during interactions. *Journal of Personality and Social Psychology*, *97*, 1005–1020.

Schmidt, R. C., Carello, C., & Turvey, M. T. (1990). Phase transitions and critical fluctuations in the visual coordination of rhythmic movements between people. *Journal of Experimental Psychology*, *16*, 227–247.

Searle, J. (1990). Collective intentions and actions. In P. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in communication* (pp. 401–415). Cambridge, MA: MIT Press.

Searle, J. (1995). *The social construction of reality*. New York: Free Press.

Sebanz, N., Knoblich, G., & Prinz, W. (2003). Representing others' actions: Just like one's own? *Cognition*, *88*, 11–21.

Shintel, H., & Keysar, B. (2009). Less is more: A minimalist account of joint action in communication. *Topics in Cognitive Science*, *1*, 260–273.

Shockley, K., Baker, A. A., Richardson, M. J., & Fowler, C. A. (2007). Articulatory constraints on interpersonal postural coordination. *Journal of Experimental Psychology: Human Perception and Performance*, *33*, 201–208.

Shockley, K., Richardson, D. C., & Dale, R. (2009). Conversation and coordinative structures. *Topics in Cognitive Science*, *1*, 305–319.

Shockley, K., Santana, M. V., & Fowler, C. A. (2003). Mutual interpersonal postural constraints are involved in cooperative conversation. *Journal of Experimental Psychology: Human Perception and Performance*, *29*, 326–332.

Song, J. H., & Nakayama, K. (2009). Hidden cognitive states revealed in choice reaching tasks. *Trends in Cognitive Sciences*, *13*, 360–366.

Spivey, M. J. (2007). *The continuity of mind*. Oxford: Oxford University Press.

Spivey, M. J., Grosjean, M., & Knoblich, G. (2005). Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences, USA*, *102*, 10393–10398.

Tanenhaus, M. K., & Brown-Schmidt, S. (2008). Language processing in the natural world. *Philosophical Transactions of the Royal Society*, *B*, *363*, 1105–1122.

Thomas, L., & Lleras, A. (2007). Moving eyes and moving thought: On the spatial compatibility between eye movements. *Psychonomic Bulletin & Review*, *14*, 663–668.

Tipper, S. P., Lortie, C., & Baylis, G. C. (1992). Selective reaching: Evidence for action-centered attention. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 891–905.

Tucker, M., & Ellis, R. (2001). The potentiation of grasp types during visual object categorization. *Visual Cognition*, *8*, 769–800.

Tuomela, R. (1992). Group beliefs. *Synthese*, *91*, 285–318.

Tuomela, R. (1993). What is cooperation? *Erkenntnis*, *38*, 87–101.

Tuomela, R. (1995). *The importance of us: A philosophical study of basic social notions*. Palo Alto, CA: Stanford University Press.

Tuomela, R. (2005). We intentions revisited. *Philosophical Studies*, *125*(3), 327–369.

Tuomela, R. (2007). *The philosophy of sociality: The shared point of view*. Oxford: Oxford University Press.

Turvey, M. T. (1990). Coordination. *American Psychologist*, *45*, 938–953.

Varela, F. J., Thompson, E., & Rosch, E. (1992). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press.

Warner, R. (1992). Cyclicity of vocal activity increases during conversation. *Behavioral Science*, *37*, 128–138.

Wegner, D. M. (2003). The mind's best trick: How we experience conscious will. *Trends in Cognitive Sciences*, *7*, 65–69.

Wojnowicz, M., Ferguson, M., Dale, R., & Spivey, M. J. (2009). The self-organization of deliberate attitudes. *Psychological Science*, *20*, 1428–1435.