

Framing joint action¹

Abstract. Many philosophers have offered accounts of shared actions aimed at capturing what makes joint actions intentionally joint. I first discuss two leading accounts of shared intentions, proposed by Michael Bratman and Margaret Gilbert. I argue that Gilbert's account imposes more normativity on shared intentions than is strictly needed and that Bratman's account requires too much cognitive sophistication on the part of agents. I then turn to the team-agency theory developed by economists that I see as offering an alternative route to shared intention. I concentrate on Michael Bacharach's version of team-agency theory, according to which shared agency is a matter of team-reasoning, team-reasoning depends on group identification and group identification is the result of processes of self-framing. I argue that it can yield an account of shared intention that is less normatively loaded and less cognitively demanding.

Keywords: joint action; shared intention; social normativity; team-reasoning, frames; Michael Bratman; Margaret Gilbert; Michael Bacharach.

1. Introduction

In this paper I try to strike a middle ground between minimalist approaches to joint action, according to which nothing more is needed for there to be a joint action than that a common effect be brought about by several agents' actions and maximalist approaches that require joint actions to be intentionally joint.

Minimalist approaches want to catch a wide net over joint intentions but offer no account of the intriguing phenomenon of joint intentionality; maximalist approaches target this phenomenon, but they either fail to fully capture it or capture it at too high a price.

Being bent on thrift, I turn to economists to help me lower the costs. Specifically, I hope to get help from certain recent developments of game theory, known as theories of team agency. Theories of team agency (Bacharach 2005; Sugden 1993, 2003; Gold and Sugden 2007, 2008) challenge the individualistic rationality

¹ I thank the editors of this special issue, Stephen Butterfill and Natalie Sebanz, and an anonymous referee for their comments. I also thank Sacha Bourgeois-Gironde for introducing me to the literature on team-reasoning, in particular to Bacharach's work. This research was supported by ANR Grant 07-1-191653 from the French National Research Agency.

assumptions of standard game theory and offer a theory of team reasoning aimed at capturing the modes of reasoning that can be used by people as a group. There are two parts to a theory of team agency. One is concerned with characterizing the patterns of inference people use (or should use) when reasoning as a group. The other is concerned with explaining how groups are formed. While team agency theorists tend to agree on how team reasoning should be characterized, they differ in their hypotheses about how groups or teams are formed. My primary interest here will be with Bacharach's view that team formation and therefore team agency is the result of framing. My aim is not to argue that team agency can only be the result of framing as understood by Bacharach. I remain open to the possibility that there may be different routes to team-agency. Rather, I want to argue that insofar as framing is indeed one of the routes to team-agency, the toll demanded of us to take this route is less exacting than the toll we have to pay if we take other roads. In section 2, I characterize the minimalist and maximalist approaches to joint action and sketch some of the motivations underlying the maximalist approach. In section 3, I briefly review some of the leading accounts of joint actions within the maximalist tradition. In particular, I focus on Bratman's and Gilbert's distinctive accounts. While happy to admit that the shared intentions involved in joint action can sometimes take the forms proposed by Bratman or by Gilbert, I am skeptical that either account can apply to all joint actions (even as understood within the maximalist approach), Bratman's account because it requires too much cognitive sophistication and Gilbert's because it is overly normative. In section 4, I turn to theories of team agency. I first present some of the motivations for the idea of team reasoning, and then concentrate on Bacharach's view on the role of frames in group formation. Finally, in section 5, I consider some of the advantages this proposal has over other theories of group formation.

2. Minimalist vs. maximalist approaches to joint action

Through actions we bring about changes in the environment. Actions are typically described in terms of their effects. As a single action can have a wide range of effects, it is in principle describable in a host of ways. As a first pass, one may want to say that an event qualifies as a joint action if it is describable as the

common effect of what several agents did; that is, none of the agents involved did on their own bring about this effect. Note that what is at stake here is not whether a single agent could in principle have brought about such an effect. Some actions may be essentially joint inasmuch as their outcome is not achievable by a single agent (e.g., singing a duet or lifting a thousand pound object with bare hands), but this is not the case with all joint actions (e.g., I alone could have fixed the dinner we prepared together). Rather, what matters is that, in the case at hand, several agents did as a matter of fact contribute to the effect in terms of which the action is described. This we may call the *common effect requirement*.

Not all the changes we bring about in our environment count as actions. When we sneeze and thus release millions of microbes in the air, we certainly have an effect on our environment. Although there is a sense in which sneezing is something we do, few consider sneezing a genuine action. Central to the philosophy of action is the question of what distinguishes genuine actions from mere happenings and doings. Davidson (1980, essay 3) famously argued that for an event to qualify as an action it must be something the agent does that is intentional under some description. Many philosophers have agreed with him that there was an important tie between action and intention, although determining the exact nature of this tie has proven a difficult challenge and given rise to lively debates.

The challenge is not confined to individual action; indeed, many would think that in joint action the problem strikes with a vengeance. Minimally, we want to exclude cases where a common effect is caused by doings we do not wish to count as actions. Suppose, for instance, that five people are stuck between floors in an airtight elevator in a building empty for the weekend and die there from asphyxiation. Suppose, further, that had fewer people been stuck in the elevator, they would have had enough oxygen to survive until help finally came. Although it is right to say that their dying was the common effect of their breathing away the limited supply of oxygen, we would be very reluctant to consider their dying from asphyxiation a joint action. To exclude such cases, it seems it is enough to require that the agent's individual doings that together brought about the common effect be actions, that is, in Davidson's parlance, that they be intentional under some description. Let us call this second requirement the *individual intentional action requirement*.

Some may want to say, that for an event to qualify as a joint action it is both necessary and sufficient that it meets the common effect requirement together with the individual intentional action requirement. This seems to be the conception of joint action Butterfill (2010) has in mind when he argues that there can be joint actions without shared intentions. Specifically, Butterfill argues against certain further conditions being necessary conditions on joint action. In his view, agents participating in a joint action need not be aware of the joint-ness of the action, need not be aware of the other contributing agents as intentional agents, need not therefore act in part because of their awareness of joint-ness and of other's agency, and finally need not be aware of the other agents' attitudes toward the joint action. We may call this the minimalist view about joint action. In contrast to Butterfill's, most philosophical accounts of joint actions incorporate, in one form or another, some or all of these further conditions, where to have a shared intention is to meet these requirements. Their insistence that these conditions be met seems to stem from a common concern and a common intuition. The concern is that we need a principled way of distinguishing genuine joint actions from mere joint happenings or joint doings. The intuition can be seen as an intuition as to how Davidson's dictum that an action is something an agent does that is intentional under some description should be transposed when we move from individual to joint action. The intuition is that the relevant description in the case of joint action is a description of the action as joint. In other words, a joint action is something agents do that is intentional under some description of it as joint.

Minimalist approaches allow us to catch a wider net over joint actions.

Maximalists would protest that some of the actions caught in the net are not *bona fide* joint actions.

3. Maximalist approaches: defining shared intentions

Accounts of shared intentions² are attempts to cash out what it takes for agents to act in a jointly intentional manner. These accounts all agree that joint actions are more than mere summations of individual actions and that therefore the intentionality in joint action cannot reduce to the intentionality of the individual actions that together contribute to the joint action. They all agree therefore that something more is needed. But what exactly is the right formula? This is where disagreements start and where different recipes for shared intentions are offered, mixing different ingredients in different ways.

An early recipe was offered by Tuomela and Miller (1988). Essentially they proposed that shared intentions (we-intentions in their terminology) were analyzable as sets of individual intentions together with sets of mutual beliefs about the other agents' intentions and beliefs. This analysis was criticized by Searle (1990) as failing to account for the cooperative character of joint actions. Searle used counter-examples to show that the existence of mutual beliefs among members of a group is not sufficient to ensure that their individual actions together constitute a joint action. Thus, for example, business school graduates who have been exposed to Adam Smith's theory of the hidden hand may come to believe that the best way for somebody to help humanity is by pursuing his own selfish interests. Each may form a separate intention to thus help humanity by pursuing his own selfish interests and not cooperating with anybody and they may all have mutual beliefs to the effect that each has such an intention. In such a case, despite all the businessmen having the same goal as well as mutual beliefs about their respective intentions, there is no cooperation and no collective action. What they lack is an intention to cooperate mutually. Mutual beliefs among members of a group do not ensure the presence of such an intention.

This critique of Searle's may appear somewhat unfair as it neglects certain qualifications in Tuomela and Miller's analysis. Tuomela and Miller insist that an agent's individual intention is an intention to do his part "as his part" of the total

² Different authors use different terminologies, speaking of shared intentions, collective intentions, joint intentions or we-intentions. Here I use these labels interchangeably, unless otherwise stated. Similarly, 'shared agency', 'collective agency', 'joint agency' and 'team-agency' are used interchangeably, unless otherwise stated.

action. The problem, though, is that this qualification can be understood in different ways. If "doing his part as his part" is simply construed to mean that the agent believes that others are also acting and contributing to a total outcome, then Searle's counter-example is not deflected. If it is construed to mean doing his part as his part of the joint action, then circularity looms: the aim was to define joint actions in terms of we-intentions, but now joint actions are appealed to in defining we-intentions.³

According to Searle, the cooperative dimension of joint actions can be captured only if it is accepted that the intentions attributable to the individuals involved in joint actions are different in type from the intentions attributable to those same individuals when they engage in individual actions. The idea then is that to account for cooperation we have to introduce a specific type of mental states: we-intentions. What needs to be spelled out is the sense in which we-intentions are special and, relatedly, the sense in which they can be said to imply cooperation. *Prima facie*, there are three ways in which we-intentions can be special. The first possibility is that what makes we-intentions special has to do with the type of entities they can be attributed to. The second is that what makes them special are features of their contents and thus that the dimension of cooperation is linked to specific features of these contents. Finally the third possibility is that rather than the contents or the possible bearers of we-intentions, it is the psychological mode itself – i.e., the fact that the psychological mode is that of we-intending instead of I-intending – that implies the notion of cooperation.

Searle rejects the first option on the ground that it would force one to admit the existence of some forms of primitive collective entities, an ontological commitment he sees as unreasonable. For him, all intentional states are states of individuals. Thus, even we-intentions can only be had by individuals. Indeed, he claims that even a brain in a vat could have a we-intention. Searle also rejects the second option, claiming that the content of we-intentions is of a form already present in some complex cases of singular intentions, the content of which encompasses a by-means-of relation. The idea is that in the case of singular

³ But see Tuomela (2005) for a rebuttal of Searle's charge of vicious circularity. Importantly, Tuomela also insists that his and Miller's analysis was not meant as a reductive analysis of we-intentions. The *I-mode/we-mode* distinction made by Tuomela (2006; 2007) and his discussion of full-blown we-intentions as involving the we-mode make it clear that his aim is not reductive. Interestingly, there are important commonalities between Tuomela's notion of we-mode reasoning and Bacharach's team-reasoning (Hakli, Miller & Tuomela, 2010).

intention of, e.g., firing a gun by pulling the trigger, there is only one intention and one action, with the relation of the means-intention to the overall intention being only part-whole. Similarly, for Searle, in the case of collective actions, there is only one complex: the singular intentions of the participating agents are related to the collective intention as means to ends and this relation is simply part-whole. To borrow Searle's example of two cooks, say Paul and Gilbert, preparing a hollandaise sauce together, the content of Paul's we-intention would be something like 'that we make the sauce by means of me steering' and the content of Gilbert's intention could be rendered as 'that we make the sauce by means of me pouring'. It is important to note that there is nothing in the by-means-of relation per se that implies cooperation. For instance, I can intend that we go to the police station by means of me dragging you there, and clearly in such a case no cooperation need be involved. Thus, there is nothing in the analysis Searle offers of the form of the content of collective intentions that makes it necessary that the dimension of cooperation essential to collective intentions be reflected in their contents.

The option Searle favors is the third one, namely, that what makes we-intentions special is the psychological mode itself, not the possible subjects of we-intentions nor their contents. More specifically, what is special about we-intentions is that they are mental states that "make reference to collectives where the reference to the collective lies outside the bracket that specifies the propositional content of the intentional state" (1990: 408). But what is it then about this distinctive kind of mental state that accounts for the dimension of cooperation that Searle says is essential to joint action? Unfortunately, Searle has very little to say in answer to that crucial question. Since, for Searle, a we-intention is a mental state an individual can enjoy independently of whether or not other individuals enjoy the same or similar we-intentions, cooperation can't be construed as linked to the way the subject of a we-intention is formed. Second, as we have already noted, nothing in the structure of the content of we-intentions as laid out by Searle seems to capture the notion of cooperation, since there is nothing in the by-means-of relation per se that implies cooperation. According to Searle, in order to account for the cooperative character of we-intentions, we must appeal to Background capacities. What collective intentionality presupposes is "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" (1990: 414). According to Searle, such background capacities are not themselves representational; rather, they are sets of nonintentional or preintentional capacities that enable intentional states. In other words, they are biological phenomena rather than intentional phenomena. There Searle comes to rest, his motto seemingly being that what cannot be described in intentional terms, a philosopher must pass over in silence.

Though sharing his intuition that we-intentions cannot be reduced to summations of individual intentions (even supplemented with mutual beliefs), many philosophers regret that Searle does not characterize we-intentions more fully. Here are three of the main complaints. First, as pointed out by Gilbert (2007), even an account of we-intentions that takes them to be states of individuals (an assumption Gilbert doesn't share), needs to say something about how the we-intentions of several agents should fit together for them to successfully perform a joint action. Yet, Searle doesn't. Second, as remarked by Bardsley (2007) and Gold and Sugden (2008), Searle also fails to provide a satisfactory explanation of the sense in which individual intentions derive from we-intentions. Finally, Searle is unduly hasty in sweeping cooperation under the rug of Background presuppositions and biological phenomena. On the one hand, it is controversial to exclude at the outset the possibility that cooperation could be characterized at least in part in intentional terms. On the other hand, even if our sense of others as candidates for cooperative agency is strongly linked to Background capacities, it should be possible to say more about what these capacities are and under what conditions they are put into play. As we shall see in section 4, Bacharach's framing theory offers insights into these questions.

Before turning to the theory of team agency, I want to consider two further accounts of shared intentions, offered respectively by Bratman and Gilbert. Neither Bratman nor Gilbert agrees with Searle that shared intentions are a *sui generis* type of mental state attributable to a single individual agent. For Bratman, in the same way that it takes two to tango, it takes (at least) two to share an intention: a shared intention is a structure of interconnected intentions of individual agents. For Gilbert, shared intentions can only be attributed to a group as such, to what she calls a plural subject.

Bratman's approach to shared intentions is a constructivist approach that builds on his planning theory of individual agency. His aim is to provide sufficient

conditions for shared intentions using the conceptual and normative resources of this planning theory. In other words, he thinks we can account for the joint-ness of shared intentions if we construe them as structures of suitably inter-related intentions and other attitudes of participants, where the component intentions of the individuals belong, *pace* Searle, to the ordinary brand of intentions, only with special and distinctive contents and interrelations. He also thinks that the social norms that apply to and guide shared intentions can be shown to emerge from the norms that govern individual planning agency.

Bratman (2009a, 2009b) proposes that shared intention involves the following main building blocks:

- (1) Intentions on the part of each in favor of the joint activity.
- (2) Interlocking intentions: each intends that the joint activity go in part by way of the relevant intentions of each of the other participants.
- (3) Intentions in favor of meshing subplans: each intends that the joint activity proceeds by way of subplans of the participants that are co-realizable and can be consistently agglomerated.
- (4) Disposition to help if needed: given that the contribution of the other participants to the joint activity is part of what each intend, and given the demands of means-end coherence and of consistency that apply to intentions, each is under rational pressure to help others fulfill their role if needed.
- (5) Interdependence in the persistence of each participant's relevant intention: each continues to intend the joint activity if and only if (they believe) the other participants continue to so intend.
- (6) Joint-action-tracking mutual responsiveness: each is responsive to each in relevant subsidiary intentions and in relevant actions in a way that tracks the joint action.
- (7) Common knowledge among all participants of all these conditions.

I won't comment extensively on each of these conditions here, but I'll offer some remarks. First, although the first condition refers to the joint action, it is not the condition that is supposed to account for the joint-ness of the activity in the strong sense in which we are interested. If it were, the account would be circular. Indeed, Bratman takes care to point out that the concept a joint activity that figures in the contents of the intentions in (1) should be understood in a way that is neutral with

respect to shared intentionality. Instead, it is condition (2) that is the most central to Bratman's account of shared intentionality. It is the fact that for each participant, the content of their intention refers to the role of the intentions of other participants that, for him, captures the intentional joint-ness of their actions. It should also be noted that conditions (3), (4) and (6) can be derived from condition (2) taken together with the norms already associated with individual planning and acting. It should also be remarked that conditions (1) and (2) both violate the own action condition, i.e. the constraint that one can only intend one's own actions, with condition (1) ranging over others' actions and condition (2) over others' intentions. However, Bratman (1992) argues that conditions on intending *that* are weaker than conditions on intending *to*. Roughly put, for one to intend that *A*, it is not necessary that one suppose this intention will lead one to do *A*, it suffices that one suppose it will lead one to do something that has an influence on whether or not *A* obtains. Condition (6) clarifies the nature of this influence: each believes that his or her intention controls the intentions and actions of others by way of its support of the persistence of the other's relevant intentions.

Bratman's ambition in giving this account is avowedly modest. His interest is in cases of small-scale shared intentional agency in the absence of asymmetric authority relations. He purports to offer sufficient conditions for shared intentions rather than necessary and sufficient conditions, allowing therefore for possible alternative constructions and leaving open the possibility that shared intention is multiply realizable. There are indeed reasons to think that, even leaving aside cases where asymmetric authority relations are present, alternative construals of shared intentions will be needed. As we shall see in section 4, there are situations where agents cannot influence others' intentions and thus where some of the conditions specified by Bratman cannot be met and where joint action is nevertheless possible. Another limitation to the scope of Bratman's accounts comes from the fact that in his construction, the materials come cheap (it makes do with ordinary intentions and with the normativity already present in individual planning agency), but their assemblage is costly and demands cognitively sophisticated agents. As pointed out by Tollefsen (2005) and Pacherie and Dokic (2006), Bratman's analysis requires that the participating agents have concepts of mental states, since each participant should represent that the other participants have intentions and other attitudes relevant to the joint activity. It requires full-

fledged meta-representational abilities insofar as the contents of the intentions of each participant make reference to both their own intentions and the intentions of the other participants. The imposition of such strong cognitive requirements would imply that animals and small children who lack altogether or have not yet fully developed mentalizing and meta-representational abilities as well as adequate mastery of the norms of practical rationality associated with planning cannot share intentions and engage in joint action. Although it is currently debated whether non-human animals ever engage in true joint action, existing evidence indicates that children clearly do and do so before they acquire the kind of cognitive sophistication Bratman's analysis suggests is required (Rakoczy 2006, 2007; Tomasello and Carpenter 2007; Tomasello et al. 2005; Warneken et al. 2006, Warneken and Tomasello 2007).

As an account targeting small-scale shared intentional activity by suitably cognitively sophisticated agents in situations where they are in a position to mutually influence their intentions and related attitudes, Bratman's account has many assets. By conceiving of shared intentions as an interlocking web of intentions of individuals, it moves away from the classical reductive analyses of collective action, since it maintains that the crucial link among the attitudes of agents involved in joint activity is not just a matter of mutual belief or mutual knowledge is not sufficient to ensure that intention is shared or collective. At the same time, the account is metaphysically and normatively parsimonious. On Bratman's view, it is not necessary to introduce, as Searle does, a new *sui generis* kind of mental state, ordinary intentions do the trick. It is not necessary to introduce new substrata to whom shared intentions would be ascribed; instead, it is sufficient that we ascribe a set of intentions, *seriatim*, to individual human agents. It is not necessary to see the social normativity characteristic of shared intentions as a basic, non-reducible form of normativity; rather this social normativity emerges from the normativity already associated with individual planning agency.

In Gilbert's view, however, this parsimony of Bratman's account is inopportune. Participation in a shared intention involves mutual entitlements to others playing their parts and, correlatively, mutual obligations, each to the other, to act as appropriate to the shared intention. For Gilbert (1997, 2009), the entitlements and obligations inherent in shared intentions cannot be understood as moral in kind

and cannot be understood as emerging from the norms associated with individual planning agency. Rather, the social normativity associated with shared intention is of a *sui generis* kind. As Gilbert sees it, to account for this special normativity, shared intentions have to be construed in terms of joint commitments:

Persons *P1* and *P2* have a shared intention to do *A* if and only if they are jointly committed to intending as a body to do *A*. (Gilbert 1997: 73)

Gilbert describes the idea of a joint commitment as an analogue of the idea of personal commitment in individual agency. When an individual has formed an intention or made a decision, he has in virtue of this intention or decision sufficient reason to act in a certain way; that is, all else being equal, he is rationally required to act in that way. Thus, a personal intention or decision entails a personal commitment to act in a certain way. Analogously, a joint decision or intention to act involves a joint commitment. Importantly, however, Gilbert insists that joint commitments are not concatenations of personal commitments. Rather, in the basic case, a joint commitment is created when each of two or more people openly expresses his personal readiness jointly with the other to commit them all in a certain way and it is common knowledge between them that all have expressed their readiness. The author of a joint commitment comprises those who jointly committed themselves by their concordant expressions. Together they constitute the plural subject of the commitment. In Gilbert's view, plural subjects and joint commitments are indissociably linked: there can be no plural subjects without joint commitments and there can be no joint commitments that are not the commitments of a plural subject.

Gilbert (2009) glosses what she means by "intending as a body to do *A*" as emulating, as far as possible, by virtue of the actions of each, a single body (or agent) that intends to do the thing in question. This in itself is not very illuminating. One way to flesh out the idea of "acting as a body" is in terms of satisfying the type of rationality constraints that bear on individual agency. To intend as a body would then be a matter of acting in such a way that the actions of each together satisfy norms of consistency, agglomeration and means-end coherence. This would involve, in Bratman's terms, commitments to mutual compatibility of relevant sub-plans, commitments to mutual support, and joint-action tracking mutual responsiveness.

If we interpret her idea of "acting as a body" in this way, Gilbert may well be willing to accept that some of the normativity present in shared intentions emerges from the normativity already present in individual agency. But still, she would disagree with Bratman that all the normativity essential to shared intentions can be so derived. For her, mutual obligations and entitlements are part and parcel of what constitutes a shared intention and can only be accounted for in terms of joint commitments. For Bratman, although common, these obligations and entitlements are not essential to shared intentions and when present can be usually understood as belonging to the familiar moral kind. Here the disagreement threatens to turn into a battle of intuitions. Gilbert (1997, 2009) takes the claim that shared intentions essentially engage non-moral mutual obligations and entitlements as intuitively obvious and offers examples rather than arguments in support for this claim. Bratman (2009b) replies with counter-examples where mutual obligations of performance appear to be absent.

One may also question Gilbert's construal of the link between a plural subject and a joint commitment. Gilbert sees them as the two sides of the same coin: joint commitments are commitments of plural subjects and only when there is a joint commitment is there a plural subject. It seems one could accept the first claim while rejecting the second. Even if it is accepted that when we jointly commit to intending to *A*, each of us is committed as a member of the group and as such incurs obligations and entitlements, it may be that membership in a group is not a matter of joint commitment. I may for instance make a personal decision to become a member of an already existing group (as when one joins a political party or the neighborhood association) or I may simply conceive of myself as a member of a group. To the extent that membership in a group is a matter of personal decision, it is a commitment I can rescind on my own. To the extent that it is a matter of self-conception, it can cease if one's self-conception changes. If we accept that membership in a group or plural subject, on the one hand, and joint commitment to intending that *A*, on the other, can be separate processes, the former a precondition of the latter, then it should be possible to accommodate Gilbert's claim that we incur obligations as members of a plural subject committed to a certain course of action, while denying that anyone has an obligation to be and remain a member of the plural subject or an entitlement to the membership of others. On this view, those who fail to satisfy their obligations as members of a

group certainly forfeit their entitlement to be considered members of a group, but the group has no right to insist that they remain members of the group despite themselves. I strongly suspect that reluctance to accept Gilbert's stance on plural subject and joint commitments is at bottom motivated by the conviction that sociality should not be had at the expense of personal autonomy.

Gilbert purports to provide individually necessary and jointly sufficient conditions on shared intention: to wit, two or more people share an intention to do *A* if and only if they are jointly committed to intending as a body to do *A*. The price to pay for this ambition is thus the introduction of the notion of joint commitment which Gilbert conceives as a basic, non-reducible kind of commitment. For her, a joint commitment creates for those jointly committed a set of obligations (and corresponding entitlements), that can neither be construed as forms of moral obligations, nor be derived from the constraints inherent in individual rational planning. Their normativity is therefore neither that of moral norms neither that of norms of individual rationality. It is rather a *sui generis* form of social normativity. The notion of joint commitment is thus basic or non-reducible insofar as joint commitments are the source of this *sui generis* form of social normativity. One may fear with Bratman that, with this notion of joint commitment, her account takes on board more normativity than is strictly needed. One may also question the strength of the link between the notions of joint commitment and of plural subject. Loosening this link may allow us to pare down the social normativity in shared intention to more acceptable levels, while preserving the gist of Gilbert's insight that the social normativity in shared intention is *sui generis*. But if plural subjects are not (or not always) explicable in terms of joint commitments, we will need alternative accounts of their formation. In contrast to Gilbert's, Bratman's account aims at metaphysical and normative parsimony, but the downside of this parsimony is cognitive prodigality. In addition, both accounts require that some form of communication, verbal or otherwise, be possible between agents, allowing them to influence each other's intentions or to form explicit or tacit agreements to jointly commit themselves to doing something together. Thus, they cannot in principle account for the emergence of shared agency in situations where agents cannot communicate or influence each other in other ways. While there is no doubt that communication can support and facilitate cooperation and shared agency (Ostrom 1990, Ostrom *et al.* 1992), it is

questionable whether communication is a necessary precondition of shared agency.

In contrast, theories of team agency, to which I now turn, are motivated in part by the need to explain how coordination and shared agency can emerge, absent any possibility of communication between the agents. As I hope to show, they can also help us build an account of shared intentions that doesn't assume extreme cognitive sophistication or wholesale normativity.

4. Team Reasoning

Theories of team agency were developed by economists Sugden (1993, 2003) and Bacharach (2006) in response to problems that arise in classical game theory. Classical game theory assumes that the players of a game are ideally rational agents and have perfect information: they maximize expected utility, given the expected behavior of others, and they have common knowledge of the game itself and of the rationality of other players. One central motivation for theories of team reasoning is that there are games that create problems for conventional game theory. In these games, there exists some strategy that is arguably rational and that many people adopt in real life, but which can't be explained or predicted as rational by classical game theory.

One such puzzle is the Prisoner's Dilemma, a typical version of which is given in Figure 1. For each player, *defect* is the dominant strategy (i.e., regardless of what the opponent does, *defect* earns a higher pay-off than *cooperate*). Conventional game theory, in its explanatory form, therefore predicts that both players will choose *defect* and, in its normative form, prescribes that they do. Yet, the players would be better off, if they had both chosen *cooperate*. A substantial number of people see that, since in experiments in which people play the Prisoner's Dilemma for money, anonymously and without repetition, between 40 and 50 % of the participants choose *cooperate* (Sally 1995).⁴

Figure 1 about here

A second puzzle is the game of Hi-Lo, a version of which is given in Figure 2. In this game players must choose between two actions, *a* and *b*. They receive

⁴ See also the empirical results of Colman et al. (2008) showing that in social dilemma games most players prefer team-reasoning strategies to strategies supporting unique Nash equilibria, although individually rational players should choose equilibrium strategies.

something only if they both choose the same, but they get more if they both choose a than if they both choose b . This puzzle is, if anything, even more puzzling than the first one. It is intuitively obvious that the rational choice for both players is a . Yet, conventional game theory has no explanation of what makes the choice of a rational. All the theory says is that if either expects the other to play a , then a is the rational thing to do, but also that if either expect the other to play b , then b is the rational thing to do.

Figure 2 about here

Many game theorists have interpreted these results as showing that the rationality assumptions of classical game theory were too strong. If the assumption of perfect rationality is relaxed it is not difficult to construct theories that purport to explain the choice of a . Yet, it may seem counterintuitive to have to invoke failures of rationality in order to yield a solution to an apparently trivial problem of coordination. In contrast, theories of team agency retain the classical rationality assumptions; what they think is wrong with the classical game theory is its focus on individual choice. The key move in these theories consists in replacing the question "What should I do?", asked separately by each individual, with the question "What should we do?". If instead of reasoning as separate individuals, the players reason as members of a team, then it will be collectively rational for them to choose (a,a) over (b,b) , (a,b) and (b,a) and, similarly to choose $(cooperate, cooperate)$ over $(defect, defect)$, $(cooperate, defect)$ or $(defect, cooperate)$.

To articulate this insight, a theory of team agency should do two things: (1) provide a theory of team-reasoning, i.e. an analysis of the reasoning an individual uses, if, thinking of herself as a member of a group, he or she asks the question, "What should we do?" instead of the question "What should I do?"; and (2) provide a theory of team formation, i.e., explain how the individual comes to ask one of these questions rather than the other?. A complete theory of team agency should thus have two complementary parts.

Roughly, as Bacharach puts it, "somebody 'team-reasons' if she *works out the best possible feasible combination of actions for all the members of her team, then does her part in it*" (2006: 121). What a theory of team reasoning does is refine and spell this out in a number of ways, characterizing the inference schemas that

capture the modes of reasoning involved in team-reasoning.⁵ I won't dwell on this aspect of the theory here. Rather, I concentrate on the issue of group formation, or of how individual agents come to identify with a group, and more specifically on Bacharach's treatment of this issue.

For Bacharach, thinking of oneself as a member of a group is a matter of framing. A frame is a set of concepts or descriptors used when thinking about a situation. To take a trivial example, thinking of a glass as half-empty or as half-full are alternative ways in which one can frame a given situation, and whether one uses one frame or the other may have important consequences for how one behaves with respect to this situation. According to Bacharach, whether an agent identifies with a group or not is a matter of what frame she uses to represent herself and the agents with whom she is interacting, and the frame she uses will determine the logic by which she will reason about what to do. Bacharach takes group identification to be a basic human propensity. His concern is with identifying conditions for the production of group identification.

Bacharach relies on theories of group identification developed in social psychology. According to Brewer and Garner (1996), people don't have a single self-concept, but rather a range of self-representations falling within three main categories: personal, relational and collective. The personal self is a self-conception as having a unique, differentiated identity. The relational self is a self-conception derived from connections and role relationships with significant others. The collective self is a self-conception defined in terms of membership in social categories or groups. Self-conception is self-framing, and like other forms of framing has characteristic instability and context-dependence. Psychologists have identified a number of conditions that tend to produce group identification, including belonging to the same social category (e.g. being a woman, a philosopher, a Parisian), to the same *ad hoc* category (being born on the 1st of June), face-to-face contact, "we" language, shared experience (e.g., being an air crash survivor), having common interests, being subject to a common fate, interdependence, and a competing outside group (e.g. analytic vs. continental philosophers). Whether a situation promotes group identification and to which

⁵ See, for instance, Gold & Sugden (2007, 2008) for a presentation and discussion of those reasoning schemas.

group depends on whether the situation presents some of these properties and whether they are salient enough to prime the corresponding group frame.

The two features most relevant to inducing a "we" frame in the kinds of games Bacharach is interested in are common interests and interdependence. In Bacharach's definition (2006: 82), two players have common interests if, given their options for action, there are at least two possible outcomes, such that the interests of both are better served in one than in the other. On this definition, having common interests is not synonymous with having identical interests and is indeed compatible with the players also having very different interests.

Importantly, this is the case in the Prisoner's Dilemma, where with regard to the two remaining possible outcomes, the interests of the players are in conflict.

The second important feature that can induce group identification in games is strong interdependence. Informally, there is interdependence in a game if there is an outcome of common interest that can only be achieved together, and there is strong independence, if this outcome is not preferred by both agents to all other feasible outcomes. According to Bacharach if the two features of common interest and strong interdependence are present in a situation and salient enough to be perceived, they can induce a "we" frame (i.e. self-identification as a member of the team of players). In the Hi-Lo game these two features are present and are highly salient, a "we" frame, hence team-reasoning, will normally be induced. The situation in the Prisoner's Dilemma is more ambiguous; like the duck-rabbit illusion, it can be seen in two different ways. On the one hand, common interest and strong interdependence are present and become salient when one concentrates on the main diagonal ((*cooperate, cooperate*) and (*defect, defect*)). On the other hand, if one looks at the outcomes on the other diagonal ((*cooperate, defect*) and (*defect, cooperate*)), one sees a conflict of interest and the possibility of being double-crossed by the other player if one chooses *cooperate*. It is therefore important to stress that for Bacharach the presence of common interests and strong interdependence is no guarantee that a "we" frame will be used. What frame is induced turns on the relative salience of various features of a situation and on the strength of their tendency to stimulate or inhibit group identification.

Bacharach also insists that the use of a frame is not a matter of choice. In the same way that you do not choose to see the duck-rabbit figure as a duck rather than as a rabbit, you don't choose to think of yourself as a member of team rather than as a

separate individual. This claim may be too strong as it stands. While it may well be true that we do not normally choose to see the duck-rabbit figure as a duck rather than as a rabbit, once we are aware of the ambiguous nature of the figure, we are in a position to gain some control over whether we see it as a duck or a rabbit. We have some voluntary control over our attention and can use it to manipulate salience. We can choose, for instance, to focus our attention on those features of the figure that make the duck interpretation more probable. It seems in principle possible that we could manipulate our self-conception in a similar way. I suspect that at bottom what Bacharach really opposed was the idea that the adoption of a frame can be a matter of *rational* choice. His reluctance appears founded, if we consider what a meta-level version of the Prisoner's Dilemma. By that I mean a version of the dilemma where the choice is not directly between cooperating and defecting but between asking the question "What should we do?" (*We-question*) or the question "What should I do?" (*I-question*) in order to solve the problem. If a player chooses the *We-question*, she will select the (*cooperate, cooperate*) option and do her part in it, namely *cooperate*. If a player chooses the *I-question*, she will as a result choose *defect*. As shown in Figure 3, the payoff matrix of the meta-level version of the Prisoner's Dilemma is exactly the same of the payoff matrix of the original Prisoner's Dilemma. We are thus back to square one. If the players ask the question "What question should I ask?", the answer is that the *I-question* should be chosen, but if they ask the question "What question should we ask?", the answer is the *We-question*.

Thus, on pain of infinite regress, the cost to be paid for preserving the rationality assumptions of the classical game theory while resolving its puzzles, is in accepting that the adoption of one mode of reasoning over another cannot be a matter of rational choice. This, of course, is not to say that adopting one frame over another in given situations cannot be externally or evolutionary rational (be to the agent's benefit and enhance fitness). Indeed, it should be noted that in his book Bacharach surveys a body of evolutionary theory that provides evidence that humans have evolved to be cooperators and that the psychological mechanisms that support group identification are the proximate mechanisms that make cooperation possible.

Going over these evolutionary considerations goes beyond the scope of the present paper. What I want to consider in the remainder of this paper are the new

perspectives on shared agency offered by Bacharach's view of group identification as a matter of framing.

5. New perspectives on shared agency

Among the views surveyed in section 3, the one Bacharach is closest to in spirit is probably Searle's. Bacharach's theory can be interpreted as fleshing out Searle's insight that collective intentionality presupposes a sense of the other as a candidate for cooperative agency. Bacharach's discussion of group identification aims in effect at spelling out at least certain of the production conditions for this phenomenon and at characterizing its effects, most importantly changes in modes of reasoning. Another important commonality between Searle and Bacharach is that for both shared agency remains vested in individuals. For Bacharach, shared agency does not involve transferring agency from individuals to plural subjects; rather it is a matter of individuals conceiving of themselves as members of a group and engaging in team-reasoning. It is, so to speak, a within-subject transformation of agency. Yet, for there to be an actual shared intention, several agents must engage in team-reasoning. Thus, in Bacharach's framework, a shared intention would be characterized as follows:

Two persons P1 and P2 share an intention to A, if:

- (i) each has a self-conception as a member of the team consisting of P1 and P2 (*collective self-framing*);
- (ii) each reasons that A is the best choice of action for all members of the team (*team-reasoning*); and
- (iii) each therefore intends to do his part of A (*participatory intention*).

One important advantage of the theory of team agency is that it allows cooperation to emerge even in situations where agents cannot communicate or influence each other in other ways. In particular, in the one-shot version of the Prisoner's Dilemma and the Hi-Lo game the theory of team-reasoning aims at solving, the players are not able to influence each other. This is in contrast to both Bratman's and Gilbert's theories of shared intention. As team-reasoning theorists have pointed out (Bacharach 2006: 138 sq; Bardsley 2007; Gold and Sugden 2007, 2008), in Bratman's account, the decisions and actions of the agents are governed by classical (i.e. individualist) canons of rationality. Unless they can influence each others' intentions through their own intentions or actions, they

won't be able to rationally generate determinate expectations about others' actions. Thus, in situations where agents are unable to influence each other, there is no way they can rationally decide to cooperate. Although this is less frequently noted (but see Bardsley 2007: 154), Gilbert's theory faces a similar problem. Once they have formed a joint commitment, agents can, as a plural subject, rationally decide in favor of the cooperative option. Yet, it is through explicit or tacit agreements that they form joint commitments, and these agreements require some form of prior interaction.

The fact that on Bacharach's theory group identification and team-reasoning are the result of self-framing and need not involve joint commitments also deflates the normative import of shared agency. An agent whose current self-conception is as a member of a team is subject to the norms of team-reasoning. This entails obligations to herself as a member of the team (she will be rationally required to do what team-reasoning tells her is her part of the best combination of actions for members of the team), but not necessarily obligations to others or entitlements to their performance. Her engaging in team-reasoning may well lead her to form expectations as to what others will do, but expectations are not entitlements. This is not to say that forming joint commitments may not be a very efficient way to induce group identification, hence team-reasoning and joint action, but to say that this is not the only way.

Finally, thinking of joint actions and shared intentions in terms of frames allows us to considerably reduce the cognitive demands we impose on agents to be participants to joint actions. Recall that on Bratman's analysis, agents must have full-fledged mentalizing capacities in order to form shared intentions: they must have concepts of mental states, since each participant should represent that the other participants have intentions and other attitudes relevant to the joint activity, and they must have well-developed meta-representational abilities insofar as the contents of the intentions of each participant make reference to both their own intentions and the intentions of the other participants.

Bacharach's version of team-agency theory is also less demanding than Sugden's version of team-agency theory. Although both agree that an intention is a shared intention when it is arrived at through a process of team-reasoning and that group identification is a prerequisite of team-reasoning, they hold different views of what leads to group identification. For Bacharach, group-identification is a

framing phenomenon automatically induced by psychological mechanisms. Frames are not chosen and an agent therefore does not choose to engage in team-reasoning rather than in individual reasoning. In contrast, for Sugden (2003), team-reasoning is merely a 'logic' among others, that is, an internally consistent system of axioms and inference rules. Whether one endorses a particular logic, thereby accepting as true any conclusions that can be derived within it, is a matter of choice. On Sugden's assurance view of team-reasoning, it is rational for an actor to endorse a principle of team reasoning which prescribes acting as a team member only conditional on assurance that others have endorsed the same principle. For the necessary assurance to be provided, Sugden requires that for all members of the group there be common reason to believe that each member of the group endorses and acts on team reason. With respect to the cognitive demands imposed on agents, Sugden's position is thus intermediate between Bratman's and Bacharach's. In contrast to Bratman, Sugden does not require that the contents of the intentions of each actor make reference to both their own intentions and the intentions of the other actors. In contrast to Bacharach, Sugden requires assurance, that is, he requires that actors have reasons to believe that each other endorses team-reasoning, that each other has reason to believe that each other endorses team-reasoning, and so on.

On Bacharach's version of team-agency theory, assurance plays no essential role. Psychological mechanisms of group identification can lead people to spontaneously team-reason, without first considering whether others will be so inclined. Having team-reasoning depend on such psychological mechanisms makes engagement in joint action much less cognitively costly. Group identification involves sensitivity to certain features of situations, including the presence of common interests and interdependence. The detection of common interests and interdependence presupposes some capacity to represent others as animate, goal-directed, and intentional agents. Many developmental psychologists would agree that agency-detection and goal-attribution are precursors of mindreading and that intention and desire-ascription are early components of mindreading.⁶ Yet, demanding that agents have these skills is a far cry from

⁶ A range of researchers have argued that infants are sensitive to some aspects of goal-directed activity and discriminate between intentional and accidental actions (Gergely, Nadasky, *et al.* 1995; Csibra 2008; Tomasello and Rakoczy 2003; Woodward 1998; Woodward and Sommerville 2000). Developmental psychologists also widely agree that children's understanding of desire and

demanding that they have full-fledged mindreading abilities, involving mastery of a fuller range of mental concepts and sophisticated reasoning about intentions and other attitudes.

The modesty of the demands made by on mindreading abilities is not offset by extreme demands on reasoning skills. The basic inferential principles used in team-reasoning are quite analogous to the inferential principles used in individual reasoning. This is not to say that team-agency theory has no use for sophisticated mindreading abilities or reasoning skills; only that those skills are not a prerequisite of shared agency. Indeed, more robust mentalizing capacities may be needed when situations are less immediately transparent than they are typically assumed to be in game-theory. Normal form game matrices present combinations of actions with their associated utilities, thus presupposing that the problems associated with working out what the different action alternatives are and what their respective utilities are have already been solved. In such games, it is therefore typically obvious for the agents what action profile maximizes the utility of the team, and what component of this profile each should perform. In many real-life situations, however, working out one's part isn't trivial. Even as simple a situation as Searle's case of the two cooks preparing a sauce together raises problems: both may have reached, through team reasoning, the decision to prepare the sauce together, but this in itself doesn't make it obvious who should do the pouring and who the stirring. In other cases, it may not be obvious to the agents not just who should do what, but also how they should proceed to achieve their shared goal. Experimental work on joint action in young children suggests that what limits what they can do jointly is not their inability to share goals, but rather their very rudimentary skills at coordinating their actions towards a shared goal (Warneken & Tomasello 2007). It is probable that when we move beyond very simple or routine joint actions, the appropriate meshing of subplans often requires that agents take into consideration the beliefs and intentions of their partners. The need to engage more robust mentalizing abilities also arises when there are discrepancies between our expectations about our partners' actions and what they actually do. Although, according to Bacharach's version of team-agency theory, agents need not have expectations about others' actions, intentions and beliefs in

intention develops earlier than their understanding of belief (Baron-Cohen 1993; Wellman & Bartsch and Wellman 1995). There is also evidence that initially children have difficulty clearly

order to team-reason, team reasoning may yield such expectations. When they are not met, it can be important to understand why, and mindreading can help us make sense of these discrepancies. Finally, as pointed out by Hakli, Miller and Tuomela (2010), Bacharach's theory takes group preferences for granted. Yet, it may not be obvious what utility values individuals associate to various outcomes or how individual preference orderings relate to collective preference orderings.

Bacharach's theory doesn't make other theories of shared agency superfluous or redundant. It should certainly be allowed that joint actions can involve shared intentions with the characteristics described by Bratman or with the properties insisted on by Gilbert. At the same time, it casts doubt on the idea that either account delivers a set of necessary and sufficient conditions for shared intentions. If as the team-agency theory suggests, shared agency can be the result of framing rather than be produced by joint commitment, then shared intentions need not be as normatively loaded as Gilbert argues. If shared agency can be the result of framing, it needs not require the type of cognitive sophistication Bratman's theory demands. Bacharach's approach thus allows us to reduce both the normative and the cognitive cost of shared intentions, while still allowing us to capture a form of joint action more substantial than what the minimalist approach is content with. Thus also, while the primary motivation for team agency theory was quite independent of developmental issues, it offers new perspectives on the development of shared agency.

Yet, Bacharach's theory is not a panacea and has limitations of its own. First, it is not always out in the open what the action alternatives are and what the group preferences are. Bacharach's theory should therefore be supplemented with accounts of how action alternatives and group preferences can be worked out by agents. Second, although Bacharach's idea that group identification is a matter of framing, hence that team reasoning can be spontaneous, is an important insight, I am not sure we should follow him in thinking that team-reasoning can never be a matter of choice. It seems in principle possible to exert some voluntary control, direct or indirect, over the psychological processes involved in self-framing once we become aware of their role. Finally, an account of shared intention should tell us not just what it takes to engage in joint action, but also what it takes to ensure it

distinguishing intentions from desires (Astington 1991, 1994; Perner 1991).

is successfully performed. Team agency theory allows us to understand how it is possible for people to rationally decide to pursue a shared goal. Yet, it has very little to say about how in many real-life situations people succeed in coordinating their actions towards a shared goal. To answer this question, we must appeal to resources that go beyond what the theory has to offer. These further resources may include further mindreading skills, and also, but this is another story, other kinds of mechanisms allowing individuals to share representations, to predict actions, and to integrate the actions of self and other.⁷

⁷ See Sebanz, Bekkering, and Knoblich (2006), Sebanz and Knoblich (2009) for discussions of these further mechanisms.

References

- Astington, J. (1991). Intention in the child's theory of mind. In D. Frye & C. Moore (Eds.), *Children's Theories of Mind: Mental States and Social Understanding*, . Hillsdale: Lawrence Erlbaum Associates.
- Astington, J. (1994). *The Child's Discovery of the Mind*. Cambridge, Mass: Harvard University Press.
- Bacharach, M. (2006). *Beyond Individual Choice*. N. Gold and R. Sugden (Eds.). Princeton: Princeton University Press.
- Bardsley, N. (2007) On collective intentions: collective action in economics and philosophy. *Synthese*, 157, 141-159.
- Baron-Cohen, S. (1993). From attention-goal psychology to belief-desire psychology. In S Baron-Cohen, H. Tager-Flusberg and D. J. Cohen (Eds.), *Understanding other minds: Perspectives from autism* (pp. 59-82). Oxford: Oxford University Press.
- Bartsch, K., and Wellman, H. M. (1995). *Children talk about the mind*. New York ; Oxford: Oxford University Press.
- Bratman, M. E. (1992). Shared cooperative activity. *The Philosophical Review*, 101(2), 327-41.
- Bratman, M. E. (2009a). Shared Agency. In C. Mantzavinos (Ed.), *Philosophy of the Social Sciences: Philosophical Theory and Scientific Practice* (pp. 41-59). Cambridge: Cambridge University Press.
- Bratman, M. E. (2009b). Modest Sociality and the Distinctiveness of Intention. *Philosophical Studies*, 144, 149-165.
- Brewer, M. & Gardner, R. M. (1996). Who is this "we"? Levels of collective identity and self-representations. *Journal of Personality and Social Psychology*, 50(3), 543-549.
- Butterfill, S. (2010). Joint action without shared intention. Manuscript.
- Colman, A., Pulford, B. & Rose, J. (2008). Collective rationality in interactive decisions: Evidence for team reasoning. *Acta Psychologica* 128: 387–397.
- Csibra, G. (2008). Goal attribution to inanimate agents by 6.5-month-old infants. *Cognition*, 107(2), 705-717.
- Davidson, D. (1980). *Essays on Actions and Events*. Oxford: Oxford University Press.
- Gergely, G., Nadasky, Z. Csibra, G & Biro, S. (1995). Taking the Intentional Stance at 12 Months of Age. *Cognition*, 56, 165-193.
- Gilbert, M. (1997). What is it for Us to Intend? In G. Holmstrom-Hintikka & R. Tuomela (Eds.), *Contemporary action theory*, (Vol. 2, pp. 65–85). Dordrecht: Springer.
- Gilbert, M. (2007). Searle and collective intentions. In S. Tsohatzidis (Ed.), *Intentional acts and institutional facts* (pp. 31-48). Dordrecht: Springer.
- Gilbert, M. (2009). Shared intention and personal intentions. *Philosophical Studies*, 144, 167–187.
- Gold, N. and R. Sugden (2007) Collective Intentions and Team Agency. *Journal of Philosophy*, 104(3), 109-37.

- Gold, N. and R. Sugden (2008) Theories of Team Agency. In F. Peter and S. Schmidt (Eds.) *Rationality and Commitment*. Oxford: Oxford University Press.
- Hakli, R., Miller, K. & Tuomela, R. (2010). Two kinds of we-reasoning. *Economics and Philosophy*, 26, 291–320.
- Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.
- Ostrom, E., Walker, J. & Gardner, R. (1992). Covenants with and without a sword: self-governance is possible. *American Political Science Review* 86: 404-417.
- Pacherie, E. & Dokic, J. (2006). From mirror neurons to joint actions. *Journal of Cognitive Systems Research*, 7, 101-112.
- Perner, J. (1991). On representing that: The asymmetry between belief and desire in children's theory of mind. In D. Frye & C. Moore (Eds.), *Children's Theories of Mind: Mental States and Social Understanding*. Hillsdale: Lawrence Erlbaum Associates.
- Rakoczy, H. (2006). Pretend play and the development of collective intentionality. *Cognitive Systems Research*, 7, 113-127.
- Rakoczy, H. (2007). Play, games, and the development of collective intentionality. In C. Kalish & M. Sabbagh (Eds.), *Conventionality in cognitive development: How children acquire representations in language, thought and action. New Directions in Child and Adolescent Development*. No. 115 (pp. 53-67). San Francisco: Jossey-Bass.
- Sally, D. (1995). Conversation and Cooperation in Social Dilemmas: A Meta-Analysis of Experiments from 1958 to 1992. *Rationality and Society*, 7, 58–92.
- Searle, J. (1990) Collective Intentions and Actions. In P. Cohen, J. Morgan, and M.E. Pollack (Eds.), *Intentions in Communication* (pp. 401-416). Cambridge, MA: Bradford Books, MIT Press.
- Sebanz, N., Bekkering, H. & Knoblich, G. (2006). Joint action: bodies and minds moving together. *Trends in Cognitive Science*, 10, 70–76.
- Sebanz, N., & Knoblich, G. (2009). Prediction in Joint Action: What, When, and Where. *Topics in Cognitive Science*, 1, 353–367.
- Sugden, R. (1993). Thinking as a Team: Toward an Explanation of Nonselfish Behavior. *Social Philosophy and Policy*, 10, 69–89.
- Sugden, R. (2003). The Logic of Team Reasoning. *Philosophical Explorations*, 6, 165–81.
- Tomasello, M. & Carpenter, M. (2007). Shared intentionality. *Developmental Science*, 10, 121-125.
- Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, 28, 675 - 691.
- Tomasello, M., & Rakoczy, H. (2003). What makes human cognition unique? From individual to shared to collective intentionality. *Mind and Language*, 18(2), 121-147.
- Tollefsen, D. (2005). Let's pretend! Children and joint action. *Philosophy of the Social Sciences*, 35(1), 75–97.
- Tuomela, R. & Miller, K. (1988). We-intentions. *Philosophical Studies*, 53, 367-389.
- Tuomela, R. (2005). We-intentions revisited. *Philosophical Studies*, 125, 327–369.

- Tuomela, R. (2006). Joint Intention, the We-mode and the I-mode, *Midwest Studies in Philosophy* XXX, 35-58.
- Tuomela, R. (2007). *The Philosophy of Sociality: The Shared Point of View*. New York: Oxford University Press.
- Warneken, F., Chen, F., & Tomasello, M. (2006). Cooperative Activities in Young Children and Chimpanzees. *Child Development*, 77, 3, 640 – 663.
- Warneken, F. & Tomasello, M. (2007). Helping and cooperation at 14 months of age. *Infancy*, 11, 271-94.
- Woodward, A. L. (1998). Infants Selectively Encode the Goal Object of an Actor's Reach. *Cognition*, 69, 1-34.
- Woodward, A. L., and Sommerville, J. A. (2000). Twelve-Month-Old Infants Interpret Action in Context. *Psychological Science*, 11(1), 73-77.

		Player 2	
		<i>Cooperate</i>	<i>Defect</i>
Player 1	<i>Cooperate</i>	2,2	0,3
	<i>Defect</i>	3,0	1,1

Figure 1. The Prisoner's Dilemma

		Player 2	
		a	b
Player 1	a	2,2	0,0
	b	0,0	1,1

Figure 2. The game of Hi-Lo

		Player 2	
		<i>We-question</i>	<i>I-question</i>
Player 1	<i>We-question</i>	(C,C), (C,C) 2,2	(C,C), D 0,3
	<i>I-question</i>	D, (C,C) 3,0	D, D 1,1

Figure 3. The meta-version of the prisoner's dilemma