Rev.Phil.Psych. (2011) 2:335–353
DOI 10.1007/s13164-011-0054-3

Socially Extended Intentions-in-Action

Olle Blomberg

Published online: 12 May 2011

Abstract According to a widely accepted constraint on the content of intentions, here called *the exclusivity constraint*, one cannot intend to perform another agent's action, even if one might be able to intend *that* she performs it. For example, while one can intend *that* one's guest leaves before midnight, one cannot intend *to* perform her act of leaving. However, Deborah Tollefsen's (2005) account of joint activity requires participants to have intentions-in-action (in John Searle's (1983) sense) that violate this constraint. I argue that the exclusivity constraint should not be accepted as an unconditional constraint on the contents of intentions-in-action: one may intend *to* perform a basic action that belongs both to oneself and to another agent. Based on the phenomenology of tool use, I first argue that intentions-in-action of one's basic actions may be *technologically extended*, meaning that their contents are not restricted to concern the agent's bodily movements. In analogy with this, I then argue that the phenomenology of some skillful joint activities supports the idea that one's basic intentions-in-action may be *socially extended*, in violation of the widely accepted exclusivity constraint. Tollefsen's account is specifically constructed to account for the joint activities of infants and toddlers who lack the capacity to think of others as planning agents and grasp their plan-like intentions (a capacity required by Michael Bratman's (1992, 1993, 2009a, b) influential account of joint activity). At the end of the paper, I raise some doubts regarding the extent to which infants and toddlers have socially extended intentions-in-action.

1 Introduction

According to Michael Bratman's influential philosophical account of joint activity, if you and I are engaged in a small-scale joint activity *J*, this activity must be an outcome of our shared intention to *J* (Bratman 1992). For us to have a shared intention to *J*, each of us

O. Blomberg (✉)
Department of Philosophy, School of Philosophy, Psychology and Language Sciences, University of Edinburgh, Dugald Stewart Building, 3 Charles Street, George Square, Edinburgh EH8 9 AD, UK
e-mail: olle.blomberg@gmail.com

must intend that we (continue to) *J* (Bratman 1992, 1993, 2009a, b). For example, if you and I are tangoing together, each of us must (among other things) have had an intention of the form "I intend that we tango." Following Christopher Kutz, I will call intentions of this form *group-intentions* (2000, p. 21). A group-intention is an intention that belongs to an individual participant, but whose content refers to the joint activity of all participants. Sometimes the presence of overlapping group-intentions is claimed to be a necessary condition for an activity to count as a truly joint one (Bratman 1992, 1993; Pettit and Schweikard 2006; Alonso 2009). At other times, it is presented as at least present in a core class of cases (Bratman 1997, 2009a, b; Kutz 2000).[1] However, appeals to group-intentions are problematic since the content of such intentions—"that we *J*"—appear to violate a widely accepted constraint on what can be intended by an agent. The constraint, which I call *the exclusivity constraint*, says that the performance of another agent's action cannot be part of the content of one's intention to do something (Searle 1983; Kutz 2000; Bardsley 2007; Bratman 2009a, b). This constraint appears to be logically implied by what Bratman calls the *own-action condition* on something being an intention, which says that one can only intend to perform one's own actions (1997, p. 53, 2009b, p. 156). Assuming that an action can only belong to one agent, the exclusivity constraint follows from the own-action condition.

Bratman argues that his planning conception of intention allows him to be more liberal about what it is intelligible to intend (see also Pettit and Schweikard 2006, p. 21). In Bratman's words, "the planning conception of intention supports the legitimacy of the appeal to my intention that *we J*." (1992, p. 331) On this conception, an intention is a plan-like mental state that commits the agent to act. The intentions of others can function appropriately in our plans since we are often able to reliably predict that another agent will form a certain intention. This allows you, for example, to intend that I tell you the time by means of intending to ask me what time it is (Bratman 1997, p. 58). Simple examples like this show, convincingly I think, that the exclusivity constraint is not a constraint on the contents of plan-intentions.

However, reliance on this planning conception of intention risks making participation in joint activity too conceptually or cognitively demanding (Tollefsen 2005; Pacherie 2007, p. 166; Pacherie and Dokic 2006, p. 110). Participants must be capable of both *having* group-intentions and of *perceiving* or *inferring* the group-intentions of other participants. Given that group-intentions must be plan-intentions, participants must at least be capable of having and recognising plan-intentions. But infants and toddlers engage in what appear to be joint activities (social pretend play, for example), despite the fact that they are not sophisticated "mindreaders". That is, they lack an explicit understanding of others as agents whose behaviour is driven by mental states such as beliefs, desires, and plan-like commitments to act. Before 3 to 5 years of age, children fail at so-called elicited-response false belief tasks, in which they have to take into account an agent's false belief about (for example) an object's

[1] Bratman (1997) suggests that shared intention may be multiply realisable, so perhaps not all forms of shared intention have group-intentions among their building blocks. Kutz (2000) writes: "Group-intentions are ordinary, instrumental individual intentions whose subject is the individual agent and whose object is a collective act or outcome: I intend that we will dance the tango. Clearly some paradigmatic forms collective action incorporate our action as the direct aim." (p. 21) Kutz' main point, however, is that there is a wide range of activities where not all participants have such intentions but which nevertheless qualify as joint activities. A discussion of this point is beyond the scope of this paper.

location in order to correctly answer a question about where the agent will look for the object (this is a very robust finding, see Wellman et al. 2001). While it is clear that children's spontaneous behaviour (such as their looking behaviour) is sensitive to the mental states of others long before 3 years of age, including their false beliefs (see Baillargeon et al. 2010), such implicit understanding does not allow children to understand others as agents guided and constrained by plan-like intentions. Furthermore, some joint activities that adult participants spontaneously engage in do not involve the kind of deliberation and planning that is arguably required for a shared intention (in Bratman's sense) to be formed. For example, you and I may be strangers who show up at a milonga where we spontaneously start dancing tango with each other upon merely making eye contact. Or we might be acquaintances who bump into each other on the street and fall into an informal conversation that moves from topic to topic in an improvised manner.[2] While it is no doubt possible to assimilate such cases into Bratman's planning framework (see e.g. 1999, p. 139), I suspect that a more plausible account of such cases is possible outside it.

Here, I focus on the challenge that Bratman's account faces in light of joint activities with infant and toddler participants, who lack capacities for planning and recognising the plans of other agents. To accommodate such activities, joint action theorists must either provide an account of joint activity that does not appeal to group-intentions at all, or argue that group-intentions and recognition of group-intentions are possible outside Bratman's planning framework. By relying on John Searle's account of intention and action (1983, ch. 3, see also 1980), I will show that the latter option is plausible, and argue that we can have what I call *socially extended intentions-in-action*. In some circumstances, agents cannot merely intend *that* we *J*, but can also perform an action which includes the performance of another agent's action in its Intentional content.[3] That is, one can intend (us) *to J* in such a way that the own-action condition is met even if the content of the intention violates the exclusivity constraint.

In the next section, I outline Deborah Tollefsen's (2005) analysis of joint activity, which is tailored to deal specifically with the joint activities of infants and toddlers. Her analysis follows the structure of Bratman's account closely, but the participants' intentions are Searlean intentions-in-action (Searle 1983) rather than Bratmanian plan-intentions.[4] She thus takes something very controversial for granted, namely that an agent can have an intention-in-action with the content "something like 'that we *J*'" (Tollefsen 2005, p. 93). (What she here takes for granted is thus what I argue is possible in this paper.) In section 3, I present the exclusivity constraint and the own-action condition in more detail. Sections 4 and 5 contain the meat of the paper. In section 4, I argue that we can have basic intentions-in-action that are

[2] Bratman (2009a, p.150; also 1999, p.130) takes "our having a conversation together" as an example of the kind of "small scale shared intentional agency" that he is interested in. Kutz (2000) also mentions "conversing" as an example of a joint activity (p. 2).

[3] I will follow John Searle (1983) in writing 'Intentional' with a capital 'I' when I refer to Intentional content in the technical philosophical sense of the satisfaction conditions of mental states. In other words, not only intentions or actions have Intentional content, but so do beliefs, desires, hopes, etc. The Intentional content of a belief is the conditions under which the belief is true. The Intentional content of an intention is the conditions under which the action counts as having been successfully performed.

[4] Searle (1990) has an account of what he calls "collective intentions-in-action" that is very different from Tollefsen's account. I will not discuss his account in this paper (but see footnote 17).

"technologically extended", and whose conditions of satisfaction thus range beyond movements of the agent's body. In section 5, I argue that intentions-in-action in an analogous way can be "socially extended", and that, as a consequence, the exclusivity constraint should not be accepted as an unconditional constraint on our intentions-in-action. While I raise some doubts concerning whether Tollefsen's account is a good starting point for a *general* account of young children's participation in joint activities, her account is nevertheless fruitful for understanding other types of joint activities where there is little room for deliberation and planning. Such activities will be ones where the participants are highly skilled at performing the joint activity and where they can trust not only their own skill but the expertise of the other participants on which they rely. They will also typically be ones where the actions of the participants are tightly coupled to each other. Examples that come to mind include that of two professional dancers or ice skaters performing a dance together, or that of two partners making love. Finally, in section 6, I consider and respond to various objections to the idea that intentions-in-action may be socially extended.

2 Tollefsen's Account

According to Bratman's account, a joint activity is the result of a shared intention, and a shared intention is simply a pattern of "interlocking" plan-intentions of the participants about which they have common knowledge. For you and I to have a shared intention, the following conditions must be fulfilled (Bratman 1993, p. 106):

We intend to *J* if and only if[5]

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

For example, for me and you to be going on vacation together, this activity must be the outcome of our having (had) a shared intention to go on vacation together.[6] For this in turn to be the case, each of us must have a group-intention of the form "I intend that we go on vacation together". Each of us must also intend that our subplans for carrying out our group-intentions should be made consistent (mesh). If I intend that we spend our time while on vacation partying, but you intend that we spend that time on early morning walks in the countryside, and this clash of subplans is known to us, then we do not yet have a shared intention to go on vacation together. Finally, the fact that we have these group-intentions and that each intend

[5] Bratman elsewhere suggests that shared intention may be multiply realisable (see 1997). In effect, the fulfilment of some other package of conditions may also result in two or more people having a shared intention. Each of the conditions in Bratman's account is an insufficient but necessary part of an unnecessary but sufficient package of conditions that realises a shared intention.

[6] When Bratman uses locutions such as "We intend *to J*" or "They have a shared intention *to J*", these should not be taken as attributions of ordinary intentions to perform an action to a collective. Hence, attributing a shared intention "to go on vacation together" to two individuals does not violate the exclusivity constraint.

that our subplans mesh must be common knowledge between us. This is a paradigmatic example of a joint activity that may be illuminated by Bratman's account. Both the initial formation of the shared intention and its execution involves deliberation and planning, as well as interdependent commitments to act and common knowledge among participants. Bratman's account is meant to be quite general in scope though. It is supposed to account for the jointness of activities such as two people singing a duet, playing basketball together, engaging in conversation (Bratman 1999, p. 130), or dancing tango (Bratman 1997, pp. 50–51).

As Tollefsen (2005) points out, condition (2) can only be fulfilled if you and I are proficient mindreaders, since it is not merely necessary that we have meshing intentions and plans, but that each of us recognise that this is the case. However, children do not appear to develop a robust representational "theory of mind" until between 3 to 5 years of age (Wellman et al. 2001).[7] Given that what 1-to-3-year-olds are doing together is appropriately described as joint activity, then this seems to be of kind of joint activity in which condition (2) is not satisfied. Tollefsen (2005) calls this *the mutual responsiveness problem*.[8]

I assume here that the mutual responsiveness problem is a real problem for Bratman's account when it is applied to the joint activities of young children.[9] But it should be noted that this is not beyond controversy. The idea that children become sophisticated mindreaders between 3 and 5 years of age is mainly based on their performance on so-called elicited-response false-belief tasks, but what is crucial here is really their understanding of other's plan-like commitments to act rather than their understanding of other's false beliefs. In addition, some findings strongly indicate that children have some implicit grasp of the false beliefs of others already at 15 months of age (Onishi and Baillargeon 2005; see also Baillargeon et al. 2010). However, in light of Bratman's planning account of intention, which identifies intentions with mental states that fulfil certain roles in practical reasoning and deliberation, I take it that the understanding that is required to avoid the mutual responsiveness problem is of an explicit rather than an implicit sort. Furthermore, there is as far as I know no positive evidence for the idea that infants and toddlers have an understanding of others as planning agents guided by plan-like commitments to act.[10]

[7] Following Tollefsen (2005, p. 81), I assume that a robust theory of mind includes the following: (i) an understanding of other persons in terms of their thoughts, intentions, and beliefs; (ii) an understanding that other persons' thoughts, beliefs, and intentions may differ from one's own; and (iii) an understanding that others have thoughts and beliefs that may not match with the current state of affairs (false beliefs).

[8] Tollefsen points out that Bratman's common knowledge condition also gives rise to a problem, since having common knowledge arguably also requires having a robust theory of mind. However, a discussion of this problem and Tollefsen's proposed solution is beyond the scope of this paper.

[9] This may not be a problem for Bratman himself since he is primarily interested in the shared agency of planning agents, such as "adult humans in a broadly modern world" (2009b, p. 153). However, others have adopted his account to understand what is involved in children's participation in joint activity (e.g. Tomasello et al. 2005, p. 680; Carpenter 2009, p. 381).

[10] Malinda Carpenter (2009) states that 12-month-olds "arguably show evidence of understanding something about others' intentions or plans for action—the means others have chosen to use to achieve their goals." (p. 382) She interprets this as showing that they understand something about Bratmanian plan-intentions and that, thus, the mutual responsiveness problem does not really arise. But Carpenter seems to use the notion of a 'plan for action' in a very broad sense that encompasses low-level control of movement sequences during action. Bratman's notion of planning is much more narrow.

At any rate, children engage in what appear to be joint activities with each other or with care takers before the age of 3. Between the age of 9 and 15 months, they start to engage in coordinated joint visual attention with both care takers and peers (Bakeman and Adamson 1984; Butterworth 2001). Around the same time they also start to participate in playful activities together with others, such as building a block tower together (Hay 1979; Verba 1994). In their second year, they even start to engage in "shared pretence scenarios" (Rakoczy 2008, p. 507; Harris and Kavanaugh 1993). For example, a shared pretence scenario might start with an adult pretending to pour tea into two cups and pretending to spill some, to which the child might respond appropriately by pretending to wipe the tea away from the spot where the adult pretended to spill it. Such social play appears to be a clear case of a joint activity.

According to Tollefsen, the mutual responsiveness problem can be avoided if conditions (1) and (2) are reinterpreted so that the intentions they refer to are *intentions-in-action* rather than plan-intentions. 'Intention-in-action' is a technical concept introduced by Searle (1983).[11] It refers to the Intentional component of an action, which specifies the (typically) detailed and fine-grained conditions that the action's movement component must meet in order for the action to be successful.[12] One of these 'conditions of satisfaction' is always that the intention-in-action itself is the cause of the action's movement component. In other words, intentions are causally self-referential. Crucially for the argument in this paper, the content that one accesses in one's experience of acting is the content of the action's intention-in-action. But according to Searle, one can have an intention-in-action without having a *conscious* experience of acting (1983, pp. 91–92). What Searle has in mind here, I think, is that one may automatically perform an intentional action while one is focused on some other primary task or is daydreaming. For example, while thinking about how to formulate this sentence, I may scratch my head or take a sip of coffee without paying any attention to performing these actions. Typically, intentions-in-action are expressed verbally in the form "I am *A*-ing" or "I am doing *A*" (1983, p. 84). Deliberate actions also have what Searle calls a *prior intention*, but for actions that are not premeditated, this is the whole story. Prior intentions are expressed in the form "I intend to *A*" or "I will *A*" and have more coarse-grained contents than intentions-in-action (1983, p. 93).

Tollefsen (2005) argues that intentions-in-action are perceptually overt by means of various behavioural cues such as "facial expression, extended hands, [or] expressions of cooperativeness." (p. 93) The conditions of satisfaction of an intention-in-action are manifest in such cues, and according to Tollefsen children can "literally see these conditions." (p. 91) If this is true, then an intention-in-action version of condition (2) can be met, even when the participants are children who lack sophisticated mindreading capacities (p. 93). There is at least ample evidence which shows that young children can distinguish animate agency from other

[11] Searle simply writes "intention in action", without the dashes. I have added the dashes to emphasise that this is a technical concept of Searle's. Note that I have also inserted dashes wherever Searle is quoted using this concept, so that it always reads "intention-in-action" in this paper.

[12] Searle writes that "in any real-life situation the intention-in-action will be much more determinate than the prior intention, it will include not only that the arm goes up but that it goes up in a certain way and at a certain speed, etc." (1983, p. 93).

movements, as well as understand that an agent's behaviour is directed towards a specific goal (Tomasello et al. 2005; Gergely and Csibra 2003). This does not necessarily imply that they can perceive Searlean intentions-in-action, but the evidence is at least consistent with this hypothesis. Furthermore, Tollefsen is not alone in interpreting the evidence in this way (see e.g. Pacherie 2000, 2007; Tomasello et al. 2003). I tentatively accept this interpretation of the evidence.

While Tollefsen's modification of condition (2) is a promising solution to the mutual responsiveness problem, I have pointed out that her modification puts the account in a real (rather than merely apparent) conflict with the exclusivity constraint.[13] Children may be able to *perceive* intentions-in-action with a content "something like 'that we *J*'" (p. 93), but how can they or their care takers *have* intentions-in-action with such contents?[14] Such contents are ruled out by the exclusivity constraint. While Tollefsen loosely refers to the content as "something like '*that* we *J*'", intentions-in-action are intentions *to* perform an action, not intentions *that* some state of affairs comes about.[15] From now on, I will therefore refer to the content of these intentions-in-action as "to *J*" rather than as "that we *J*".

3 The Own-Action Condition and Two Constraints on the Content of Intentions

Nicholas Bardsley refers to what I call the exclusivity constraint as "an uncontroversial constraint […] that an individual's intentions cannot be said to range over others' actions",[16] and he uses it to rule out analyses of shared (or collective) intention as inadequate (2007, p. 144). This constraint is indeed widely accepted and uncontroversial when it comes to what one can intend *to* perform. While one can intend *that* one's guest leaves before midnight, one cannot intend to perform her act of leaving (Bratman 1993, pp. 101–102, 2009b, p. 157). Bratman accepts that "what one attempts are [only] one's own actions." (1992, p. 330). In agreement with Bratman, Philip Pettit and David Schweikard state that "[i]t is true that I cannot intend *to* X, where X-ing is a joint performance. But I may still be able to intend *that* we X together." (2006, p. 21) According to Bratman, the contents of what one can intend *to do* are constrained by the own-action condition on something being an intention, according to which "the *subject* of an intending is always the intended *agent* of the intended activity." (2009b, p. 156) Group-intentions, then, should be glossed as "I intend that [we bring about some state of affairs]", rather than as "I intend us to [perform some action]". Accordingly, I take it that Bratman accepts the exclusivity constraint as an unconditional constraint on what one may intend to do.

[13] Insofar as one closely links Searlean intentions-in-action with mechanisms of "motor imagery", then Elizabeth Pacherie's proposed account of joint action in (2007) seems to be very similar to Tollefsen's. Pacherie (2000) herself argues that empirical work on motor imagery throws light on Searle's notion of intention-in-action.

[14] There is one brief reference on to the possibility that such contents may be problematic: "Aside from the difficulties with the notion of an individual intending that we *J* […]." (Tollefsen 2005, p. 93) She does not spell out what these difficulties are though, let alone how they could be overcome.

[15] Perhaps this is why Tollefsen qualifies the content specification with "something like".

[16] It is clear from the context that Bardsley is primarily referring to "intention-in-action" here.

Other joint action theorists also seem to accept the exclusivity constraint as a constraint on what one can intend to do, among them Kutz (2000, pp. 21–22) and Searle (1983, p. 110). Kutz calls group-intentions "non-standard intentions" because they are not tightly linked to actions and "cannot be directly transformed from intentions that P to intentions to P." (p. 21) He even suggests that group-intentions might best be thought of as "either expressing a hope that we will do something, or standing in for an individual intention to promote our doing something." (p. 22) In effect, Kutz excludes the possibility that intentions to P may have a content similar to group-intentions, presumably because he accepts the exclusivity constraint.

Put in terms of Searle's notion of intention-in-action, the own-action condition says that the content of one's intention-in-action can only concern the event that, together with the intention-in-action itself, constitutes one's action. This event of one's own action is typically assumed to be one's bodily movement. It is therefore helpful to introduce another possible constraint on the content of intentions-in-action, which I will call *the own-body constraint*. According to this constraint, the content of an agent's intentions-in-action only concerns the agent's own bodily movements. I will argue later that if the own-action condition is taken to imply the own-body constraint, then the condition will often be breached in cases where humans perform actions with tools.

Searle clearly embraces the exclusivity constraint in some form.[17] He claims that if there is "intervening Intentionality" in the causal chain between the tokening of an intention-in-action and the event that is the intention's outcome, then that action cannot be successful. Searle illustrates this with the following example:

> Thus, suppose that unknown to me my arm is rigged up so that whenever I try to raise it, somebody else causes it to go up, then the action is his not mine, even though I had the intention-in-action of raising my arm and in some sense that intention caused my arm to go up. […] And that this is the right way to construe intentions-in-action is at least indicated by the fact that, when my intentions-in-action make explicit reference to the intentions of other agents, then in general the actions become the actions of those agents. Thus, suppose I know how my arm is rigged up and I want it to go up. My intention-in-action then is *getting the other agent to raise it*, not *raising it*. *My* action is getting him to raise it, *his* is raising it. (1983, p. 110)

The idea, then, is the following: In general, if I have an intention-in-action that includes in its conditions of satisfaction that it causes you to have an intention-in-action that in turn causes an action *A*, then *A* is performed by you. This excludes me from also performing *A*. The action *I* perform is instead getting you to perform *A*. Searle thinks that this line of reasoning remains valid even when I am unaware of the mediating role of your intention-in-action. However, it is clearly not valid in all such cases. We can see this by considering one of Searle's own examples (see 1983, p. 98): Gavrilo Princip's blow against Austria and vengeance of Serbia (performed by

[17] Searle thinks that individuals can have a special kind of mental attitude, which he calls "collective intention", which has the form "We intend to *A*" or "We are *A*-ing" (1990). But he embraces the exclusivity constraint as a constraint on the contents of personal intentions-in-action of the form "I am *A*-ing".

means of his killing of the Archduke Franz Ferdinand in Sarajevo on June 28, 1914). Insofar as "striking a blow against Austria" is an action at all, then it is Gavrilo's intentional action in spite of the fact that most of the links in the causal chain from intention-in-action to massive societal effects must have consisted of "intervening Intentionality". Perhaps Gavrilo was in some sense aware of this intervening Intentionality, at least upon reflection. In that case, Searle's line of reasoning is not even valid for some cases where an agent is aware of the mediating role of others' intentions in the causal chain between intention-in-action and outcome. Given Searle's choice of example in the quote above (raising one's arm), it is reasonable to take the exclusivity constraint to hold only, or only hold unconditionally at least, in cases where *A* is a *basic action*. A basic action is an action *A* that an agent can intend to perform "without intending to do another action by means of which he intends to do *A*." (Searle 1983, p. 100) In most circumstances, squeezing the trigger of a gun is a basic action, since one does not do this by means of intending to contract one's muscles in the arm and hand in such a way that the trigger gets pulled. This is unlike the case when Gavrilo intends to strike a blow against Austria. He intends to do this by means of firing a gun against the Archduke. In the case of basic actions, it is hard not to accept the exclusivity constraint as a fundamental constraint. But I will argue in the next section that the constraint does not always hold, even if its scope is restricted to intentions-in-action of basic actions.

To sum up, Tollefsen's modification of Bratman's account leads to a problem. By relying on the notion of intentions-in-action rather than plan-intentions, Tollefsen saws off the branch that allows the original account to avoid violating the exclusivity constraint. However, Tollefsen does not explain how individuals can have intentions-in-action of the form "I am intending-in-action to *J*'". I will now explain how this is possible.

4 Technologically Extended Intentions-in-Action

My strategy will be to first argue that the conditions of satisfaction of the intention-in-action of a basic action are not always restricted to only concern an agent's own bodily movements. Human agents can have what I call basic but *technologically extended* intentions-in-action. (An intention-in-action is "basic" if it is the intention-in-action of a basic action.) After establishing this, I argue in the next section that basic intentions-in-action can also, contrary to the exclusivity constraint, be *socially extended*. In effect, I am arguing that what may make up "one's own actions" is a lot less restricted than action theorists tend to assume. The own-action condition implies neither the own-body constraint, nor the exclusivity constraint. The reach of an agent's basic actions is not necessarily limited by the agent's body surface, nor by another agent's co-ownership of the action's physical movement or event.

As far as I can see, there is nothing in Searle's characterisation of intention-in-action as such that restricts the conditions of satisfaction to range only over the agent's bodily movements. I here seem to be in disagreement with Joëlle Proust (2003) who interprets Searle as embracing the own-body constraint. She claims that "an intention-in-action fails to represent any sort of further goal, such as 'switching on the light,' 'breaking a vase,' and so on. What it determines is, rather, a bodily

movement." (p. 106) It is true that Searle frequently writes that an intention-in-action causes (and presents itself as causing) "a bodily movement", but I think this is a reflection of the choice of examples that he illustrates his theory with, such as raising one's arm, rather than the expression of a general constraint that is supposed to be part of his theory. Indeed, he says explicitly that we can make "intentional bodily movements where the conditions of satisfaction of our intentions go beyond the bodily movements." (1983, p. 99) However, Searle calls the intentions of such actions "complex intentions" (1983, p. 98). He characterises the self-reflexive content of the intention-in-action of a man who fires a gun in the following way: "This intention-in-action causes it to be the case that the trigger pulls, which causes it to be the case that the gun fires." (1990, p. 409) While the extended "reach" of the action is to be found in the content of the intention, this reach is represented *as* extended, as pertaining to events beyond the agent's immediate control. But there need be nothing complex (in terms of content) about intentions-in-action that reach beyond movements available to the agent considered as an unaided body.

Consider the well-known example of the blind man and his cane who explores his surroundings by moving and tapping the white cane in front of him (Merleau-Ponty [1945] 2002, pp. 165-166). What, we can ask, is the blind man's 'experience of acting'? His experience is not that of manipulating his arm, wrist and fingers in order to cause the cane to move in certain ways, as if he was holding a cane for the first time in his life. Rather, the blind man simply taps the ground directly with the tip of the cane. The cane itself has faded out of his awareness and attention, just like our arms and hands are outside the focus of awareness and attention when we "use them" in action. We do not intend to move our joints, wrist and fingers in order to cause our hand to reach out and grasp an object, instead we simply reach out and grasp it directly. Similar things can be said about the perception involved in the skilled use of tools. The blind man's perceptual experience is that of touching the ground at the tip of the cane: he touches the ground *through* the cane. The experience is not that of touching the cane in his hand in order to *infer* what the ground is like. Such extension of our capacities for perception and action is typical of fluent use of tools, where the tools become "transparent" to the user. Following Andy Clark (2008, p. 31), I think that the best picture of the agency involved in such cases is that "of an extended or enhanced agent confronting the (wider) world" rather than that of a bare biological agent facing a tool.

Interestingly, neuroscientific findings suggest a subpersonal basis for the shift from phenomenological opaqueness to transparency. Maravita and Iriki (2004) argue that an agent's so-called body schema—a neural representation of the agent's body shape and posture—changes as the agent becomes fluent in using a tool to perform a task. This results in an extension of the agent's peri-personal space (its proximal action space). Bimodal neurons in the body schema of Japanese macaques normally only fire when a macaque touches something with its hand or sees something near its hand (hence *bi*modal). But after a macaque has learned to use a rake to reach for and collect food, these bimodal neurons also fire when the macaque touches something with the rake or sees something in the space around the rake (Maravita and Iriki 2004). Behavioural experiments suggest that such an extension of peri-personal space is also an effect of tool use in adult human subjects. In some neuropsychological patients, so-called "visual-tactile extinction" (see Di Pellegrino et al. 1997)

has been taken to indicate the boundary of peri-personal (or peri-hand) space. If a single tactile stimulus is presented to such a patient on their contralesional hand,[18] which is hidden from view on a table in front of them, then they have no problem detecting the stimulus. The same thing happens if they are simultaneously presented with a distracting visual stimulus on the ipsilesional side of their body, but not close to the ipsilesional hand. However, if the visual stimulus is presented immediately adjacent to the ipsilesional hand, then their performance in detecting the tactile stimulus drops significantly: the tactile stimulus is extinguished. Farné et al. (2007) studied one such patient's susceptibility to visual-tactile extinction before and after having used a rake to collect objects placed on a table. After tool use the space in which distracting visual stimuli resulted in extinction of the tactile stimuli was no longer restricted to the space immediately adjacent to the patient's ipsilesional hand, but extended into space where the rake's axis and tip had been.

While not conclusive, these parallel results suggest that extension of visual-tactile extinction is due to an extension of the patient's peri-personal space. This extension is induced by active tool use and encoded in the body schema (see Holmes et al. 2007 for some doubts). Thanks to the plasticity of the body schema, tools become incorporated into the suit of resources at hand that the cognitive system simply takes for granted (see Clark 2008, sect. 2.5). On this interpretation, these taken-for-granted resources define the agent's proximal action space. It should be noted though, that some behavioural experiments on adult human subjects are in tension with this extended body schema hypothesis (see Holmes et al. 2004; Holmes et al. 2007).[19]

My argument is not hostage to any particular interpretation of these experimental results. But if the extended body schema hypothesis was confirmed, then it would demonstrate an interesting isomorphism between the personal-level phenomenon of "transparency in use" and its subpersonal underpinning. Personal-level phenomenology suggests that tool-using actions are often basic actions, just like bodily actions such as raising one's arm, and if the extended body schema hypothesis is correct, then this may be correlated with the tool's incorporation into the subpersonal representation of the agent's bodily action capacities. However, the phenomenological observations about tool use would not themselves be undermined if this particular hypothesis was rejected. Within the framework of Searle's theory of

[18] The contralesional hand is the left (right) hand if the lesion is in the right (left) hemisphere. The other hand is the ipsilesional hand.

[19] For example, Holmes et al. (2007) studied the effect of tool use on subjects' perception and action capabilities by comparing under what conditions and to what degree a visual distractor stimulus induced so-called "visual-tactile interaction". Visual-tactile interaction was measured by the effect that the visual distractor (one or two flashes) had on a subject's ability to discriminate between two tactile target stimuli (one or two vibrations felt through a stick). They found that visual-tactile interaction was more marked when (i) subjects held a single stick in their right hand to discriminate target stimuli presented on the same (right) side, compared to when (ii) the subjects held the stick in their right hand to discriminate target stimuli presented on the other (left) side (so that the axis of the stick crossed the body's midline), and compared to when (iii) subjects either held a single stick in their right hand or held one stick in each hand to alternate between discriminating target stimuli presented on their right side and on their left side. Holmes et al. (2007) suggest that these results can be explained by appeal to a general effect that maintained spatial attention toward one side of the body has on visual-tactile interaction on that side (it becomes more likely). If the body schema encodes an extended peri-personal or peri-had space after tool use, then it is not clear why the degree of visual-tactile interaction should vary between these experimental conditions.

intention and action at least, the phenomenology of action should be taken seriously. After all, the content of the intention-in-action is the content that is accessed when one has the experience of acting.

Searle makes some observations that fit this take on transparent tool use. First, he notes that what counts as a basic action is relative to an agent's skills (1983, p. 100). Hence, a skilled tool user can act on the world through the tool without having to do this by means of intending to manipulate the tool. Secondly, Searle does not want to restrict the kinds of causal chains that can be involved in the execution of an intention-in-action to those occurring within the body (see 1983, p. 110). If I have an intention-in-action to raise my arm that causes my arm to raise, then this counts as my intention-in-action being carried out successfully, even if someone, unbeknownst to me, has rewired my brain so that the motor commands from my brain go half way around the world via the internet before reaching the effectors. In effect, there is nothing in Searle's account that precludes that the causal chain between an intention-in-action and the movement it represents extends beyond the agent's body. Besides Searle's habit of calling the intentions-in-action of tool-using actions "complex", there is nothing in his account that precludes the possibility of basic but technologically extended intentions-in-action.

I conclude that the contents of an agent's basic intentions-in-action are not constrained by the own-body constraint. I do not claim that the own-body constraint on the content of basic actions is explicitly accepted by many action theorists, but the idea that basic actions are bodily movements seems to be a tacit and unexamined assumption in much philosophy of action. For example, Davidson claims "that all primitive actions are bodily movements", where 'bodily movements' is "openhanded enough to encompass such 'movements' as standing fast, and mental acts like deciding and computing." (1980, p. 49) This assumption may be one source of resistance among action theorists to the possibility of intentions-in-action "to $\mathcal{J}$". However, at least within the framework of Searle's theory of intention and action, the assumption should be rejected.[20]

Since the phenomenon of 'transparency in use' has been used to argue for the view that the vehicles of the contents of perception and action extend beyond the biological boundary of the human (or macaque) organism (Clark 2008, ch. 2; see Rupert 2009, ch. 8, for a critique of this kind of argument), it is worth pointing out that widening what counts as permissible contents of an agent's basic intentions-in-action does not depend on accepting such "vehicle externalism" (Hurley 1998). The question of whether or not to accept the own-body constraint is a matter of what content it makes sense for an intention-in-action to have, given the kind of mental state that an intention-in-action is. It is not a matter of whether the vehicle of the intention-in-action may extend beyond the body surface of the agent. Neither is it a matter of whether the environment of the agent plays a role in fixing the content of its intentions-in-action. In other words, rejecting the own-body constraint does not depend on accepting either vehicle externalism or content externalism. The same can be said with regard to rejecting the exclusivity constraint.

[20] For a much more thorough discussion about basic actions and the use of tools, see Pols (forthcoming, especially ch. 3).

5 Socially Extended Intentions-in-Action

In order to reject the exclusivity constraint as an unconditional constraint on what one can intend to do, I will appeal to considerations similar to those that make rejection of the own-body constraint plausible. In other words, I will argue that intentions-in-action can be socially extended, not only technologically extended.

What motivated my notion of basic but technologically extended intentions-in-action was the phenomenon of transparency in skilled tool use. Do people experience a similar shift from opaqueness to transparency with increased participation in a joint activity? I think this sometimes happens. Consider the example of two skilled ice skaters performing a figure dance. From an adopted first-person perspective, Axel Seemann (2009) characterises their sense of joint control in the following way:

> [I]t isn't that *I* experience myself as being in a position to determine *your* doings by my actions. It is, rather that the experience really is one of *us* controlling our doings. And this experience is an embodied one: in the most obvious kind of case, such as the example of the figure dancers, our bodies really seem to form an experiential unit. The awareness I enjoy of your body in a dance is quite unlike the sensation of your flesh pressing against mine that I might be exposed to in a crowded subway carriage. Your body seems to form part of our joint interface with the world. (2009, p. 508)

To my ears, Seemann's observation here rings true. Like in the case of skilled tool use, agents sometimes experience each other as transparent extensions of their own activity. The coordination of actions recedes to the background of attention and the agents experience themselves acting as one attending to their joint activity. The intentions-in-action here are basic, they are not executed via the execution of other intentions, whether one's own or those of one's co-agent.

My claim here is not that the whole activity—the entire dance on the ice for example—is part of the content of the participants' basic intentions-in-action, but rather a move or turn within that activity. The following joint move may be part of each intention's content: As one of the dancing ice skaters traces a curve on the ice, he lifts his partner up into a position where she has one leg held above her head and the other leg stretched out parallel to the ice (a so-called full Biellmann position). Here, the intention-in-action of each of the skaters will be dependent on the intention-in-action of the other, and the move as whole will be a basic action for each of them. The bodily movements of their partner will be part of the conditions of satisfaction of the intentions-in-action of each of them. Furthermore, the self-reflexive condition that an intention-in-action causes the movement or event that is specified in its content is satisfied for both of them. The intention-in-action of each skater causes not just their own bodily movements, but also the bodily movements of the other. After all, in the absence of their intention-in-action, the movements of the other would not be performed.

I have argued that the notion of 'socially extended intentions-in-action' is intelligible. It is also a useful notion, since it captures the kind of intentions-in-action that participants engaged in some skillful joint activities have, such as figure skaters or dancers. One mark of such joint activity would be that the participants have

socially extended intentions-in-actions that overlap. This is one way in which we might understand the idea that intentions-in-action could have the content "to *J*" that is part of Tollefsen's account. It is also an alternative way of elaborating Searle's analysis of intention-in-action in *Intentionality* (1983) to deal with joint activity, which does not appeal to a special "we"-mode of intention-in-action (as in Searle 1990).

6 Objections and Rejoinders

One initial objection to socially extended intentions-in-action might be the following. The sources of knowledge about our own and others' actions are different in important ways. We have no proprioceptive information about the bodily posture or movements of others. To the extent that such information is required for carrying out intentions-in-action and for monitoring the dynamic unfolding of action, there is reason to doubt that intentions-in-action can be socially extended. But no such proprioceptive information is available in tool use either, so such information does not seem to set a limit to what can part of the content of basic intentions-in-action.[21]

However, evolutionary considerations suggest that the boundary of the biological organism has an epistemic importance which at least makes it unlikely that extended intentions-in-action are prevalent. Even if vision is important for awareness and knowledge of our own actions, it has been argued that proprioceptive information has a kind of privileged status in our subpersonal cognitive machinery that is not enjoyed by perceptual information about the extra-bodily environment. In arguing against the hypothesis of extended cognition (or vehicle externalism), Kim Sterelny (2004) draws attention to the fact that an organism is a conglomeration of parts that have evolved together and thus become co-adapted to each other over time. Evolution has ensured that the internal "environment" of the organism is a friendly one: Information exchange is reliable, trustworthy and efficient. "[O]ver evolutionary time", as Sterelny (2004) puts it, "the internal informational environment of an agent will become more transparent." In contrast, the external environment is a "shared and sometimes contested space" inhabited by other agents who are potentially out to deceive and manipulate the organism. Hence, while there is selection pressure on perceptual systems to become more reliable and efficient, it is unlikely that external resources, such as tools, public information, or (I would add) other agents become as transparent to us as our internal resources. This is because perceptual systems pick up information from an epistemically hostile environment. But as Clark (2008, p. 103) points out in reply to Sterelny, these evolutionary considerations do not show that technological extensions of cognition are impossible, just that we should be sceptical about their prevalence.

Given that Sterelny is correct in characterising internal information flow as trustworthy and reliable, and external information flow as always potentially noisy

[21] Indeed, in circumstances where proprioceptive and visual information about our own actions are in conflict, our conscious reports about our own bodily movements seem to rely more on vision than proprioception, as for example the 'rubber hand illusion' demonstrates (Botvinick and Cohen 1998).

and deceptive, technologically extended intentions-in-action ought to be quite rare, and arguably, socially extended intentions-in-action ought to be even more rare. The possibility of defection and deception is much greater in the case of other agents than in the case of tools. The objection, then, is that a socially extended intention-in-action will depend on the absence of a kind of vigilance and double-checking that we should expect to be the norm in interaction with other agents. However, like in the case of technological extension, these considerations do not show that socially extended intentions-in-action are impossible, just that that we should be sceptical about their prevalence. In addition, the epistemic difference between the internal and the external environment should not be exaggerated.[22] If, as Sterelny himself believes (see 2003), group selection has played an important role in the evolution of our pervasive tendency to cooperate, then human groups are in a sense made up of co-adapted parts that have evolved together. Less controversially, infants and toddlers are profoundly dependent on their caretakers and the group into which they are born. Even if there are cognitive mechanisms in older children and adults with the function of double-checking and vetting during joint activity, it does not seem plausible that such mechanisms are operating in young children, at least not during interaction with close kin or other familiar individuals. Furthermore, my argument crucially depends on the phenomenology of participation in joint activity. In cases where an agent experiences his or her participation in a joint activity as transparent, then that should be reflected in the content of his intention-in-action. I am assuming here that Searle is right to identify an agent's experience of acting with the content of the agent's intention-in-action: The content of one's experience of acting just *is* the content of one's intention-in-action.

However, perhaps the experience associated with a socially extended intention-in-action is always mistaken or illusory. One might grant that a basic intention-in-action may have conditions of satisfaction that makes it socially extended, but argue that, in fact, the self-reflexive condition that the intention-in-action causes a jointly performed movement or event will never be satisfied. This would account for the phenomenological observations that I have been swayed by, but would arguably make socially extended intentions-in-action less interesting since they would not be appropriately causally connected to the (joint) activity. But it is not clear why one should think that the self-reflexive condition cannot sometimes be satisfied. After all, an intention-in-action is just supposed to be one in a host of necessary factors that bring a movement or event about. The agent's body and the world must cooperate for the right things to happen even in the case of ordinary solo action.

One could object that socially extended intentions-in-action cannot have a *joint* or *shared* activity as its content. I have argued that the bounds of what counts as one's own basic actions are wider than action theorists have thought, but this does not make those actions joint or shared. Hence, the objector continues, I have not here gone to the heart of what is philosophically puzzling about shared intention, namely how the question of whether to do this or that can be a matter for us to settle

[22] In a recent paper, Sterelny (2010) comments on his earlier argument against the extended cognition hypothesis. He submits that his "initial contrast [between the internal and external environment] was overdrawn" and that his "earlier ideas on the importance of contested space were overinfluenced by Machiavellian models of social interaction." (p. 474).

together. As David Velleman (1997) points out, I can only intend to do something that is up to me to settle. But if it is up to me to settle whether or not to do something, then it cannot simultaneously be up to you. The puzzle is how an agent can exercise control over something while at the same time delegating this control to someone else.

What I am proposing here is not a solution to this puzzle. I am not making a proposal about how agents can share a future-directed intention or decision to engage in a joint activity, I am rather trying to suggest what characterises some joint activities, irrespectively of whether the participants jointly chose to enter the activity. I take it that many joint activities that young children engage in are characterised by what Bratman calls "mutual responsiveness in action", but not necessarily "mutual responsiveness of intentions" (see 1992, p. 339). What is joint is control and guidance of the unfolding of the joint activity, not a choice about whether or not to initiate it. Furthermore, I am not saying that the presence of overlapping socially extended intentions-in-action is the only condition required for an activity to be jointly intentional. I have merely focused on Tollefsen's first condition, but for two participants to engage in joint action, conditions (2) and (3) must also be met. The jointness of joint action must be captured by the analysis as a whole, not by the each participant's intention-in-action "to *J*".

That it is possible for participants' intentions-in-action to be socially extended does not, of course, establish that the intentions-in-action of infants and toddlers can be socially extended. Indeed, infants and toddlers do not appear to be in a good position to establish the kind of fluent expertise that makes coordination during joint activity transparent. For this reason, I doubt that an account based on socially extended intentions-in-action, which I suggest that Tollefsen's account is, is applicable to children's joint activities *in general*, including cases such as two toddlers building a block tower together or engaging each other in social pretend play. But perhaps the conditions needed for two children, or a child and a care-taker, to have overlapping socially extended intentions-in-action do obtain in some simple forms of social play that involve repetition and predictable turn-taking.[23] This seems especially plausible in asymmetrical joint activities where a child is playing with an adult, since an adult co-participant is typically more disposed to be controlled and "manipulated" by the child than a peer will be (the adult can thus make up for the child's lack of fluent expertise). In a footnote, Tollefsen even raises the intriguing possibility that, in light of young children's profound dependency on their care takers, their intentions-in-action are actually socially extended as a rule rather than as an exception (2005, p. 93 n. 24). Given the role of phenomenology in my argument for the possibility of socially extended intentions-in-action, and the difficulty of accessing what infants and toddlers experience, it is hard to draw any firm conclusion about what the role, if any, socially extended intentions-in-action might play in young children's participation

[23] One of the reviewers suggested that breastfeeding might be an example of a joint activity where a child and a care taker have meshing socially extended intentions-in-action. This, I agree, looks like a plausible case. Breastfeeding will often occur before the child meets the requirements that are implied by conditions (2) and (3) in Tollefsen's account, but in many cultures, breastfeeding continues well beyond the child's first year, in which case all three conditions may indeed be met.

in joint activities.[24] The conclusion that such intentions-in-action have a role to play in some joint activities with adult participants, however, is much more secure.

Even with adult participants, the kind of sustained experience of fluency described by Seemann is the exception rather than the rule. As Kutz (2000) points out, "the complexity arising from problems of coordination renders collective activity salient, making it stand out against a background of unreflective, self-regarding activity." (p. 12) In line with my reflections on Sterelny's earlier argument against extended cognition, he also notes that the possibility of free-riding makes joint activity unlikely to become automatised and unreflective (p. 13). However, Kutz grants, in line with what I have been saying, that "well-rehearsed joint action may require no conscious deliberation or reflection." (p. 12) Indeed, in the case of some joint actions, such as when the ice dancers perform their well-rehearsed joint move, conscious deliberation and reflection will not even be possible given the speed and timing requirements that must be met to avoid failure (and most likely, injury).

7 Conclusions

I have pointed out that couching Bratman's account of shared intention in terms of participants' intentions-in-action in order to avoid the mutual responsiveness problem, as Tollefsen (2005) does, is problematic. Such an account requires participants to have intentions-in-action with the content "to *J*", but unlike plan-like intentions "*that* we *J*", such intentions really do violate the exclusivity constraint. According to this constraint, one's intentions-in-action cannot range over the actions of another agent. I have argued that this constraint should not be accepted as an unconditional constraint on the possible contents of intentions-in-action. There is nothing in Searle's analysis of intention-in-action per se that precludes the contents of one's basic intentions-in-action to range beyond one's bodily movements, and the phenomenon of transparency during skillful tool use indicates that intentions-in-action often actually do have such wide-ranging contents. Thus, the own-action condition does not imply what I have called the own-body constraint. By analogy with such technologically extended intentions-in-action, phenomenological considerations also suggest that intentions-in-action can be socially extended. This means that the bounds of what counts as one's own actions are wider than action theorists typically assume. Neither technologically extended, nor socially extended intentions-in-action, fail to fulfil the own-action condition. As a result, the exclusivity constraint should be rejected. Two agents can have socially extended intentions-in-action with overlapping contents, so that they both guide and monitor the unfolding of their joint activity.

While I have expressed doubt regarding Tollefsen's account as a plausible *general* account of young children's joint activities, I believe the account fills a lacuna in the philosophy of joint action. Bratman (1992) mentions that "mutual responsiveness in action" is a defining feature of joint activity, but he does not have much to say about this feature. Tollefsen's account—supported by my notion of socially extended

[24] Thanks to an anonymous reviewer for pointing this out.

intentions-in-action—specifies an interpersonal pattern of mental states which I suggest is characteristic of some skillful joint activities that involve tightly coupled mutual responsiveness in action.

Acknowledgements For critical comments, helpful suggestions and encouragement in response to earlier paper drafts and talks, I want to thank Natalie Gold, Steve Butterfill, Georg Theiner, Auke Pols, Andy Clark, Matt Nudds, Richard Harper, Ashley Taylor, and my audiences at the Collective Intentionality VII and European Society of Philosophy of Psychology 2010 conferences. I also want to thank two anonymous reviewers and the Edinburgh reading group on group agency (Orestis Palermos, Jonas Christensen, Eric Kerr, and Andy McKinlay) for helpful input, and Liz Irvine for proof-reading this paper. Finally, I am grateful to Microsoft Research for funding my PhD research.

References

Alonso, F. 2009. Shared Intention, reliance, and interpersonal obligations. *Ethics* 119(3): 444–475.

Baillargeon, R., R.M. Scott, and Z. He. 2010. False belief understanding in infants. *Trends in Cognitive Sciences* 14(3): 110–118.

Bakeman, R., and L.B. Adamson. 1984. Coordinating attention to people and objects in mother-infant and peer-infant interaction. *Child Development* 55(4): 1278–1289.

Bardsley, N. 2007. On collective intentions: collective action in economics and philosophy. *Synthese* 157 (2): 141–159.

Botvinick, M., and J. Cohen. 1998. Rubber hands 'feel' touch that eyes see. *Nature* 391(6669): 756.

Bratman, M.E. 1992. Shared cooperative activity. *The Philosophical Review* 101(2): 327–341.

Bratman, M.E. 1993. Shared intention. *Ethics* 104(1): 97–113.

Bratman, M.E. 1997. I intend that we *J*. In *Contemporary action theory: social action*, eds. G. Holmström-Hintikka, R. Tuomela, 49-63. Dordrecht: Kluwer. Also published in Bratman, M.E. 1999. *Faces of Intention*. Cambridge University Press, pp. 142-161.

Bratman, M.E. 1999. Shared intention and mutual obligation. In *Faces of Intention*, Cambridge University Press, pp. 130-141.

Bratman, M.E. 2009a. Shared Agency. In *Philosophy of the social sciences: philosophical theory and scientific practice*, ed. C. Mantzavinos. Cambridge: Cambridge Scholars Press.

Bratman, M.E. 2009b. Modest sociality and the distinctiveness of intention. *Philosophical Studies* 144(1): 149–165.

Butterworth, G. 2001. Joint visual attention in infancy. In *Blackwell Handbook of Infant Development*, ed. J.G. Bremner and A. Fogel, 213–240. Oxford: Blackwell.

Carpenter, M. 2009. Just how joint is joint action in infancy? *Topics in Cognitive Science* 1(2): 380–392.

Clark, A. 2008. *Supersizing The Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.

Davidson, D. 1980. *Essays on Actions and Events*. Oxford: Clarendon Press.

Di Pellegrino, G., E. Làdavas, and A. Farnè. 1997. Seeing where your hands are. *Nature* 388(6644): 730.

Farnè, A., A. Serino, and E. Làdavas. 2007. Dynamic size-change of peri-hand space following tool-use: determinants and spatial characteristics revealed through cross-modal extinction. *Cortex* 43(3): 436–443.

Gergely, G., and G. Csibra. 2003. Teleological reasoning in infancy: the naïve theory of rational action. *Trends in Cognitive Sciences* 7(7): 287–292.

Harris, P.L., and Kavanaugh, R.D. 1993. Young children's understanding of pretense. *Monographs of the Society for Research in Child Development,* 58(1) [Serial No. 231], v-92.

Hay, D.F. 1979. Cooperative interactions and sharing between very young children and their parents. *Developmental Psychology* 15(6): 647–653.

Holmes, N.P., G.A. Calvert, and C. Spence. 2004. Extending or projecting peripersonal space with tools? Multisensory interactions highlight only the distal and proximal ends of tools. *Neuroscience Letters* 372(1–2): 62–67.

Holmes, N.P., D. Sanabria, G.A. Calvert, and C. Spence. 2007. Tool-use: capturing multisensory spatial attention or extending multisensory peripersonal space? *Cortex* 43(3): 469–489.

Hurley, S.L. 1998. *Consciousness in Action*. Harvard University Press.

Kutz, C. 2000. Acting together. *Philosophy and Phenomenological Research* 61(1): 1–31.
Maravita, A., and A. Iriki. 2004. Tools for the body (schema). *Trends in Cognitive Sciences* 8(2): 79–86.
Merleau-Ponty, M. 1945. *Phénomènologie de la perception*. Paris: Gallimard. Translated. by C. Smith. (2002). *Phenomenology of Perception.* London: Routledge.
Onishi, K.H., and R. Baillargeon. 2005. Do 15-month-old infants understand false beliefs? *Science* 308 (5719): 255–258.
Pacherie, E. 2000. The content of intentions. *Mind & Language* 15: 400–432.
Pacherie, E. 2007. Is Collective Intentionality Really Primitive? In *Mental processes: representing and inferring*, ed. M. Beaney, C. Penco, and M. Vignolo, 153–175. Cambridge: Cambridge Scholars Press.
Pacherie, E., and J. Dokic. 2006. From mirror neurons to joint actions. *Cognitive Systems Research* 7(2–3): 101–112.
Pettit, P., and D. Schweikard. 2006. Joint actions and group agents. *Philosophy of the Social Sciences* 36 (1): 18–39.
Pols, A.J.K. forthcoming. Acting with Artefacts. Doctoral dissertation, Technische Universiteit Eindhoven.
Proust, J. 2003. Action. In *John Searle*, ed. B. Smith, 102–127. Cambridge University Press.
Rakoczy, H. 2008. Pretence as individual and collective intentionality. *Mind and Language* 23(5): 499–517.
Rupert, R.D. 2009. *Cognitive Systems and The Extended Mind*. Oxford University Press.
Searle, J.R. 1980. The intentionality of intention and action. *Cognitive Science* 4(1): 47–70.
Searle, J.R. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press.
Searle, J.R. 1990. Collective Intentions and Actions. In *Intentions in Communication*, ed. Philip R. Cohen, Jerry Morgan, and Martha E. Pollack, 401-415.
Seemann, A. 2009. Joint agency: intersubjectivity, sense of control, and the feeling of trust. *Inquiry* 52(5): 500–515.
Sterelny, K. 2003. *Thought in a Hostile World: The Evolution of Human Cognition*. Wiley-Blackwell.
Sterelny, K. 2004. Externalism, Epistemic Artefacts and the Extended Mind. In *The Externalist Challenge*, ed. R. Schantz, 239–254. Walter de Gruyter.
Sterelny, K. 2010. Minds: extended or scaffolded? *Phenomenology and the Cognitive Sciences* 9(4): 468–481.
Tollefsen, D. 2005. Let's pretend! Children and joint action. *Philosophy of the Social Sciences* 35(1): 75–97.
Tomasello, M., J. Call, and B. Hare. 2003. Chimpanzees understand psychological states-the question is which ones and to what extent. *Trends in Cognitive Sciences* 7(4): 153–156.
Tomasello, M., M. Carpenter, J. Call, T. Behne, and H. Moll. 2005. Understanding and sharing intentions: the origins of cultural cognition. *Behavioral and Brain Sciences* 28(5): 675–691.
Velleman, J. 1997. How to share an intention. *Philosophy and Phenomenological Research* 57(1): 29–50.
Verba, M. 1994. The beginnings of collaboration in peer interaction. *Human Development* 37(3): 125–139.
Wellman, H.M., D. Cross, and J. Watson. 2001. Meta-analysis of theory-of-mind development: the truth about false belief. *Child Development* 72(3): 655–684.