Sharing and Ascribing Goals

PIERRE JACOB

Abstract: This paper assesses the scope and limits of a widely influential model of goal-ascription by human infants: the *shared-intentionality* model. It derives much of its appeal from its ability to integrate behavioral evidence from developmental psychology with cognitive neuroscientific evidence about the role of mirror neuron activity in non-human primates. The central question raised by this model is whether sharing a goal with an agent is necessary and sufficient for ascribing it to that agent. I argue that advocates of the shared-intentionality model underestimate both the distinction between the target and the goal of a goal-directed action and the gap between sharing and ascribing a goal.

Introduction

Much experimental work in developmental psychology over the past fifteen years strongly suggests that before the end of their first year, preverbal human infants are able to represent the actions of a variety of agents as goal-directed. Although the nature of the basic mechanisms at work in human infants' goal-ascription remains controversial, one influential model of goal-ascription by human infants has emerged: I shall call it the *shared-intentionality* model. This model derives much of its appeal from its ability to integrate behavioral evidence from developmental psychology together with cognitive neuroscientific evidence about mirror neuron (MN) activity based on single-cell recording in the brain of non-human primates. It seems to be corroborated by recent developmental evidence showing the role of early motor experience in shaping the ability of human infants to represent the actions of others as goal-directed, which has been interpreted as evidence for the role of MN activity in early goal-ascription by human infants.

Address for correspondence: Institut Jean Nicod, Ecole Normale Supérieure, Pavillon Jardin, 29, rue d'Ulm, 75005 Paris, France.
Email: jacob@ehess.fr

In philosophy and cognitive science, there is a stronger and a weaker sense of the expression 'shared intentionality'. In the stronger sense, individuals are said to share collective intentionality (or we-intentions) when they have the capacity to engage in joint (or collective) actions (e.g. tango dancing, playing chamber music or playing football).[1] Collective intentionality has recently been highlighted as a distinctive human cognitive trait.[2] In the cognitive neuroscience of action, there is also a weaker sense in which the agent of a transitive action such as grasping a peanut and an observer of that action are said to share a motor representation of the agent's action, on the grounds that many of their brain structures overlap.[3] In this weaker sense, being able to engage in joint action is not a necessary condition for shared intentionality. Nor is it distinctive of human cognition: mirror neuron (MN) activity, which was discovered in the ventral premotor cortex of non-human primates, has been widely interpreted as evidence for the existence of shared motor representations in the brains of an agent and an observer.[4] Only the weak sense of 'shared intentionality' will be relevant to this paper, the main goal of which is to assess the developmental evidence for and against the shared-intentionality model of goal-ascription by preverbal human infants.

While some of the recent behavioral investigation of goal-ascription by human infants focuses on their early ability to imitate a model's observed action, most of it has been conducted within the violation-of-expectation framework. This framework rests on the assumption that individuals look longer at an unexpected than an expected event, so that by measuring infants' looking times, psychologists can get evidence about the nature and content of their expectations. These experiments fall into two broad categories: one choice-based, and the other efficiency-based. In choice-based experiments, infants are induced to expect that an agent has a preference for one of two toys. In efficiency-based experiments, infants are induced to expect that an agent will select an action as an efficient means to achieve a goal in the presence of situational constraints (e.g. obstacles in the environment) (cf. Figure 1).

The paper falls into five sections. In the first section, I spell out the main assumptions of the shared-intentionality model and summarize some of the basic cognitive neuroscientific evidence in its favor from single-cell recording in the monkey brain. In the second section, I review some of the developmental evidence for this model from choice-based experiments conducted within both the violation-of-expectation framework and the so-called 'imitation' framework. In the third section, I review some of the choice-based evidence against the shared-intentionality

---

[1] Cf. Bratman, 1992; Gilbert, 1989; Pacherie and Dokic, 2006; Searle, 1995; Tuomela, 1995.
[2] Cf. Rakoczy and Tomasello, 2007; Tomasello, 2008; Tomasello *et al.*, 2005; Tomasello and Carpenter, 2007.
[3] Cf. Jeannerod, 1999, 2006; Decety and Grèzes, 2006; Decety and Sommerville, 2003; Jackson *et al.*, 2006; Singer and Lamm, 2009.
[4] Cf. Rizzolatti *et al.*, 1996; Gallese, 2001; Gallese *et al.*, 2004; Rizzolatti and Sinigaglia, 2008.
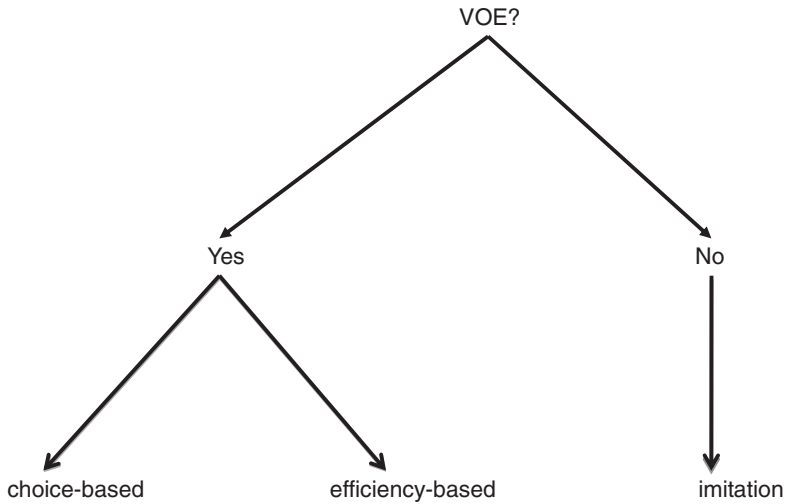
Figure 1 *The investigation of goal-ascription by human infants*

model. In the fourth section, I review some further developmental evidence against the model, based on the single-target-efficiency paradigm. In the last section, I criticize the claim that some of the recent evidence showing that infants' motor experience enhances their understanding of goal-directed actions performed by others is also evidence for the role of mirroring in human infants, and offer an alternative account of the role of infants' motor experience in goal attribution.

1. Mirroring Another's Goal

At the core of the shared-intentionality model are the two related ideas that a fundamental part of ascribing a goal to another is to share the goal with him or her and that one cannot share goals with every conceivable agent. In accordance with Meltzoff's (2005, 2007) *Like-me* framework, advocates of the shared-intentionality model stress the role of *familiarity* in shaping human infants' early ability to reason about and ascribe goals: the fundamental intuition underlying the model is that the early ability to reason about another's goal is shaped by familiarity with the agent's action. But on the face of it, many different kinds of agents and actions will be perceptually familiar to an infant, depending on her social and cultural environment. The shared-intentionality model tries to spell out the relevant notion of familiarity in terms of a two-step appeal to the notion of *similarity* relative to the observer (i.e. the self). The *perceptual* familiarity of an agent is evaluated in terms of the human-like aspect of her bodily appearance. The *motor* familiarity of an agent's action depends on whether the observer could perform the agent's act,

i.e. on the similarity between the agent's performed act and the observer's motor repertoire. The shared-intentionality model can be captured in the five following assumptions.

(1) The first assumption is that, in accordance with Meltzoff's (2005, 2007) *supramodal* hypothesis, preverbal human infants are equipped with a supramodal code that enables them to represent in a single format both the observable acts performed by others and their own felt acts.[5]

(2) The second assumption is that in virtue of the supramodal code, preverbal human infants have the ability to single out observable agents that are *perceptually* like them (or human-like) from other agents.

(3) The third assumption is that preverbal human infants have the ability to single out observable actions that fall within their own motor repertoire, i.e. that are executable, from observable actions that are not executable.

(4) The fourth assumption is that when a preverbal human infant observes a *Like me* (i.e. human-like) agent perform an executable action that falls within the infant's own motor repertoire, the infant mentally rehearses (i.e. covertly imitates) the agent's action (or maps it onto her own motor repertoire).[6]

(5) Finally, the model assumes, along with Rizzolatti *et al*. (2001), that by mentally rehearsing (or covertly imitating) the agent's action, the infant comes to share, and thereby to understand, the psychological state (e.g. the agent's goal) that caused the agent to perform her action in the first place.

In fact, there are two versions of assumption (5): a stronger and a weaker one. The question is whether sharing a goal and ascribing a goal are two different psychological events or one and the same. On the stronger version, they are one and the same: sharing another's goal is constitutive of ascribing a goal to an agent; there is simply no difference between instances of the former and instances of the latter. On the weaker version, there is a difference between the two events: sharing is not ascribing, but sharing another's goal is a crucial causal step in the process leading to goal-ascription.[7]

The shared-intentionality model of goal ascription seems able to unify results from developmental psychology and cognitive neuroscience, in particular the activity of mirror neurons (MNs), first discovered in the early 1990s by single-cell recording in area F5 of the ventral premotor cortex (and subsequently in the inferior parietal lobule) of macaque monkeys. On the face of it, MN activity seems to strongly

---

[5] As Meltzoff (2007) puts it, infants 'can use the self to understand the actions, goals, and psychological states of others and conversely can learn about their own powers and the possibilities and consequences of their acts by observing the behavior of others'.

[6] For some relevant discussion, see Gergely, 2007.

[7] This distinction between the strong (constitutive) and the weaker (causal) interpretation of the contribution of sharing a goal to ascribing it to another is emphasized by Goldman (2009).

support assumptions (4) and (5) of the model. MNs are sensorimotor neurons that fire during both the execution and the perception of transitive actions directed towards a physical target (e.g. grasping). Two negative experimental results based on single-cell recording in area F5 of the monkey are worth mentioning at the outset because they seem to have a counterpart in developmental psychology. On the one hand, Rizzolatti *et al.* (1996) report that when a recorded monkey watched a human grasp, not with his own hand, but with pliers, MNs did not fire.[8] On the other hand, Umiltà *et al.* (2001) report that when a monkey watched a human agent's hand perform a pantomime of a grasping action and the animal knew that there was no target, MNs did not fire either.

MNs were so-called because MN activity in an observer's brain was taken to *mirror* (resonate with or simulate) MN activity in the agent's brain (without giving rise to the observer's execution of the observed act). Action-mirroring, which generates a neural similarity between an agent and an observer who comes to *share* a motor representation of the action performed by the agent, was first interpreted in terms of a *direct-matching* model of *action-understanding* (cf. Rizzolatti *et al.*, 2001; Rizzolatti and Craighero, 2004). On this two-step model, the perception of the agent's act first causes the observer to automatically map the agent's act onto her own motor repertoire, i.e. to mentally rehearse the agent's act on her own premotor cortex, in accordance with assumption (4) of the shared-intentionality model of goal-ascription. Secondly, the mental rehearsal of the agent's act enables the observer to understand 'the meaning' or *goal* of the observed act. The direct-matching model of action-understanding has given rise to two kinds of critical discussion, one of which was launched by a dilemma posed by Csibra (2007), while the other concerns the nature of the link between mirroring and understanding an agent's goal.

Csibra (2007) raises the following dilemma for advocates of the direct-matching model: what is being mirrored by MN activity must be either the agent's movements or the agent's goal, but it cannot be both at the same time. As recognized by Rizzolatti and Sinigaglia (2010), at least among humans, an agent's movements do not stand in a one-to-one relation to her goal: not only can an agent recruit different movements in the service of a single goal, but she can also recruit one and the same movement in the service of different goals. So, if what is being mapped onto the observer's motor repertoire by MN activity is the agent's overt motor act (or bodily movements), then it is unlikely that mirroring could deliver a representation (or understanding) of the agent's goal. Conversely, if the output of mirroring is a representation (or understanding) of the agent's goal, then it is unlikely to be generated by the mapping of the agent's observed movements onto the observer's motor repertoire.

---

[8] Later results showed that after active and/or observational training of the monkey, MNs in area F5 do fire when the animal either performs a grasping action with pliers or observes a human perform such an action (cf. Umiltà *et al.*, 2008 and Ferrari *et al*, 2005).

In response, Rizzolatti and Sinigaglia (2010) choose the second horn of Csibra's dilemma: on the basis of experimental results reported by Umiltà *et al.* (2008), they argue that in the monkey brain, MN activity encodes *goals*. Umiltà *et al.* (2008) trained monkeys to grasp objects using both normal pliers and so-called 'reverse' pliers: grasping an object with normal pliers involves closing the fingers, but grasping with reverse pliers involves opening the fingers. Single cell recordings in area F5 in executive tasks by Umiltà *et al.* (2008) show that F5 neurons discharged during the same phase of grasping 'regardless of whether this involved opening or closing of the hand' (Rizzolatti and Sinigaglia, 2010, p. 266).[9] In other words, what seems to matter to the firing of F5 neurons is the agent's goal (grasping) irrespective of the difference between the agent's closing and opening his fingers.

Furthermore, Rizzolatti and Sinigaglia (2010) go on to endorse a dual view of MN activity in human and non-human primates: whereas in non-human primates, only an agent's goal can be mirrored, they argue that in humans, both an agent's goal and her movements can be mirrored by MN activity in an observer's brain. As they put it, 'in the earlier studies on the mirror mechanism, it was indeed not clearly specified that the parieto-frontal mirror mechanism in humans is involved in two kinds of sensory-motor transformation—one mapping the observed movements onto the observer's own motor representation of those movements (movement mirroring), the other mapping the goal of the observed motor act onto the observer's own motor representation of that motor act (goal mirroring).'[10]

The second set of issues raised by the direct-matching model of action-understanding concerns the link between mirroring and understanding an agent's goal, which is crucial to assumptions (4) and (5) of the shared-intentionality model. In a classic paper, Rizzolatti *et al.* (2001, p. 661) drew a contrast between the direct-matching model and a purely visual analysis of an agent's action. They argued that whereas 'the results of the visual analysis of an observed action [. . .] are devoid of meaning', an observer's understanding of an agent's action generated by mirroring derives its meaning from the observer's own 'motor knowledge' (about how to perform the observed action) 'because the motor representation of that action is activated in our brain'. Advocates of the direct-matching model have stressed the immediate (or non-mediated), experiential and non-conceptual (i.e. presumably non-inferential) character of the understanding of another's action provided by MN activity (cf. Gallese *et al.*, 2004). According to the latest advocacy of the direct-matching model, ''although there are several mechanisms through

---

[9] Umiltà *et al.* (2008) report the results of single-cell recordings in area F5 mostly for *executive* tasks of grasping with normal and reverse pliers, but on p. 2112 they also refer to 'an example of F5 mirror neuron recorded during grasping with tools and during the observation of the same motor act', thereby offering evidence for *mirroring* of goals.

[10] One crucial issue that I will not address here is whether understanding an agent's goal is generated by MN activity (i.e. mirroring of goal) or whether MN activity is itself generated and enhanced by a prior and independent process of goal understanding based on contextual cues, as argued by Csibra (2007) and Jacob (2008, 2009).

which one can understand the behavior of other individuals, the parieto-frontal mechanism is the only one that allows an individual to understand the action of others 'from the inside' and gives the observer a first-person grasp of the motor goals and intentions of other individuals'' (Rizzolatti and Sinigaglia, 2010, p. 264).

Thus, on the latest version of the direct-matching model, MN activity in the monkey brain *codes* the agent's goal. If so, then MN activity in an observer's brain enables the observer to mirror the agent's goal (as a result of the former's perception of the latter's action) and the observer forms a goal with the same content as the agent's. Both the agent and the observer entertain the same goal, but only the agent executes an action that will achieve the goal. In accordance with the strong version of assumption (5) of the shared-intentionality model of goal-ascription, sharing the agent's goal is sufficient for the observer to understand and thereby to ascribe it to the agent.[11]

What advocates of the mirroring view of goal-understanding (i.e. the strong version of assumption (5) of the shared-intentionality model) seem to overlook, however, is the gap between *having* a goal and *ascribing* a goal. Mirroring an agent's goal could at best causally enhance an observer's ascription of a goal to the agent, but it could not constitute such an ascription.[12] To mirror another's goal is *to have* a goal. As stressed by Jeannerod (1999, 2006), not all individuals who form goals become agents: many of an individual's goals may remain unfulfilled. But whether or not an individual's goal is fulfilled, sharing a goal is having a goal and having a goal is not ascribing a goal to another. Ascribing a goal to an agent consists in forming a *belief* (or *judgment*) about an agent that he or she has a goal or is performing some goal-directed action.[13] There is a crucial difference between the psychological *attitude* of forming a goal and that of ascribing a goal: to form a goal is to mentally represent an act that will bring about a change from one state of affairs to another. Having a goal is to be directly motivated to act: goals have a *world-to-mind* direction of fit (they can be satisfied or frustrated, but neither true nor false).[14] By contrast, by ascribing a goal, one is not directly motivated to act: to ascribe a goal to an agent is to form a *belief* (or *judgment*) whose content is that the agent has a determinate goal. Like other beliefs and judgments, goal-ascriptions have a *mind-to-world* direction of fit (they can be true or false).[15] Furthermore, a creature can arguably form a goal, perform goal-directed actions and be an agent, whether or not she possesses the *concept* GOAL.[16] However, I assume that one could not

---

[11] If, as recognized by advocates of the direct-matching model, there are other ways of understanding an agent's goal, then sharing another's goal might not be necessary.

[12] This is the weak version of assumption (5) of the shared-intentionality model.

[13] This distinction is clearly drawn by Goldman (2009).

[14] See Section 2 for a further discussion of what it takes to form a goal.

[15] On the distinction between the mind-to-world and the world-to-mind directions of fit, cf. Anscombe, 1957; Searle, 1983; and Jacob and Jeannerod, 2003.

[16] Words in capital letters refer to the concepts of the psychological states whose meanings they express.

	necessary to possess concept GOAL?	motivating	truth-conditional	world-to-mind direction of fit	mind-to-world direction of fit
having/sharing a goal	−	+	−	+	−
ascribing a goal	+	−	+	−	+

Table 1 *Difference between the psychological attitude of forming a goal and that of ascribing a goal.*

ascribe a goal to some agent, i.e. form a *belief* or *judgment* that an agent has some goal or other, unless one possessed and deployed (or used) the concept GOAL. (These distinctions are summarized in Table 1.) Having outlined the main assumptions of the shared-intentionality model and some of the issues it raises, I now turn to developmental evidence for assumptions (1)-(3) of the shared-intentionality model of goal-ascription.

2. Choice-Based and Imitation-Based Evidence for the Shared-Intentionality Model

The shared-intentionality model derives much empirical support from experiments conducted within both the violation-of-expectation framework and the so-called 'imitation' framework. I start with the former. What I call the choice-based paradigm derives from a series of experiments reported by Woodward (1998, 1999). In an initial phase of the experiments, after being presented with two toys (e.g. a teddy bear on the left and a ball on the right), 9-, 6- and 5-month-old infants were habituated to seeing a human hand in a grasping posture repeatedly reach for, and rest on, one of the two toys (e.g. the teddy bear). During the test trials, the toys' spatial positions were switched and the infants saw the human hand reach for and rest on either the same toy as before (i.e. the teddy bear) or the other toy (i.e. the ball). The aim was to test whether preverbal human infants were more surprised by the agent's changing her target (new toy event) or her path (new side event). Woodward (1998) reports that 9-, 6- and even 5-month-olds looked reliably longer at the new target (old side) event than at the old target (new side) event, indicating that they expected the target to remain the same.

However, when the human hand was replaced by either an inanimate rod topped with a sponge or a mechanical claw, infants did not look longer at the new rather than the old target (Woodward, 1998).[17] Similarly, the asymmetry between looking

---

[17] On the face of it, these negative results are reminiscent of early reports of single cell recording in the monkey showing that MNs in area F5 do not fire in response to a grasping action performed by a human with pliers (cf. Rizzolatti *et al.*, 1996).

longer at the new than the old target also broke down when instead of a human hand in a grasping posture, infants saw the back of a human hand in a non–grasping posture repeatedly drop on the same toy (Woodward, 1999).[18] In accordance with assumptions (1)-(3) of the shared–intentionality model, these results strongly suggest that preverbal human infants' early goal–ascriptions are biased in favor of human–like agents who perform executable actions (that the observer can map onto her own motor repertoire).

Choice–based experiments conducted within the so–called 'imitation' paradigm yield similar results.[19] *Covert* imitation is at the heart of the shared–intentionality model of goal–ascription: according to assumptions (4) and (5) of the model, seeing an agent perform a goal–directed action causes the observer to mentally rehearse (or *covertly imitate*) the agent's action, which in turn causes the observer to *share* the agent's goal and thereby to ascribe the goal to the agent. Now from *covert* imitation to *overt* imitation, there is only a small step.[20] In a pair of so–called 'imitation' studies reported by Hamlin *et al.* (2008), after being presented with twelve toys, 7–month–olds saw a human agent sitting in front of a pair of the toys. In the first study, the infants saw the agent either contralaterally reach-and-grasp one of the two toys or ipsilaterally place the back of her hand onto one of them. Then the infants were invited to choose one of the toys. Infants systematically selected the toy that had been the target of a grasping action, not a back-of-the-hand action. In a second study, infants saw the agent either perform an unfulfilled reach-and-grasp action directed towards one of the two toys or point to one of them. When invited to choose one of the two toys, the infants were reliably more likely to select the toy that had been the target of the failed reach-and-grasp action rather than of the pointing act. These 'imitation' studies seem to corroborate both the positive and the negative results reported by Woodward (1998, 1999) in the violation-of-expectation framework.

In the first of another pair of so–called 'imitation' studies reported by Mahajan and Woodward (2009), after being familiarized with sixteen different toys, 7–month–olds saw the experimenter's face above a barrier made of black foam with a concealed opening in it through which she could either move her own hand or move a box towards one of a pair of toys. (Before seeing either the hand or the box action, infants were given evidence that the box was self-propelled.) When invited to choose one of the two toys, infants in the hand condition chose the target significantly more often than infants in the box condition. In order to test for the possible effect of the presence of the experimenter's head above either

---

[18] Also on the face of it, this negative result seems similar to the result of a single cell recording experiment in the monkey showing that MNs in area F5 do not fire in response to a pantomime of grasping when the animal knows that there is no target (cf. Umiltà *et al.*, 2001).

[19] As I shall argue shortly, I don't think that what the relevant experiments test are infants' imitative capacities.

[20] For relevant further discussion, see Gergely, 2007.

her hand or the box, Mahajan and Woodward (2009) designed a second study in which infants saw the same hand and box actions as before, but the experimenter's head was fully occluded. As in the previous study, infants were more likely to select the target toy of the hand action than of the box action. Again both of Mahajan and Woodward's (2009) 'imitation' studies seem to corroborate both the positive and the negative results reported by Woodward (1998, 1999) in the violation–of–expectation framework.[21]

Because they accept assumptions (4) and (5) of the shared-intentionality model, Hamlin *et al*. (2008) and Mahajan and Woodward (2009) are inclined to interpret their findings as evidence that infants selectively *imitate the goal* (or the goal–relevant aspects) of an agent's observed action. This interpretation presupposes that infants can determine what the agent's goal is. However, Hamlin *et al*. (2008) and Mahajan and Woodward (2009) seem to overlook what it takes both for an agent to form a goal and for an observer to determine the agent's goal. In simple cases, goal-directed actions are directed at some target. Whereas *targets* can be simple physical objects, *goals* are complex mental representations, involving three representational components. In the simplest cases, to form a goal is to mentally represent (i) an *actual* state of affairs S_1, e.g. a piece of food T (for target) resting on a plate at time t, (ii) a *possible* (non-actual) state affairs S_2, e.g. T's being in one's mouth at t+1, and (iii) an *action* that can bridge the gap between S_1 and S_2. What is common to S_1 and S_2 is T (the target); what distinguishes S_1 from S_2 is one of T's properties, i.e. its spatial position. The represented action (e.g. the agent's reaching-and-grasping T, followed by arm-flexion towards the agent's mouth) is a representation of the means by which the agent plans to achieve her goal by turning S_1 into S_2, i.e. by causing a change in T's spatial position.

So it is one thing to argue that the agent's choice between two targets is evidence that her action is goal-directed. It is quite another to claim that this evidence is sufficient to determine the agent's goal. In fact, in all the studies conducted within the Woodward choice-based paradigm, the agent's goal is considerably underdetermined, because the agent reaches for the target in a grasping posture but does not remove the target, thus leaving the initial state of affairs unchanged (cf. Section 5 for further discussion). An alternative way of interpreting the findings of both Hamlin *et al*. (2008) and Mahajan and Woodward (2009) is that the model's *preference* for one of the two toys cues the infant's own *preference*. On this interpretation, the model's preference will only cue the infant's if the agent selects a means-action that is in accordance with the representation of her action as being goal-directed (e.g. if she reaches for the target with her hand in a grasping posture, not with the back of her hand).

---

[21] The occlusion of the experimenter's head was meant to match Luo and Baillargeon's (2005) design, to be further discussed in Section 3.

3. Choice–Based Evidence Against the Shared–Intentionality Model

On the face of it, the shared-intentionality model derives support from both Woodward's (1998, 1999) positive and negative results. In accordance with the shared-intentionality model, the positive results have been interpreted as showing that preverbal human infants have a direct and immediate understanding of an agent's goal when the agent e.g. repeatedly reaches for one of two toys with her hand in a grasping posture. The negative results also seem to corroborate the model by showing that preverbal human infants do not represent the motions of an inanimate rod or a mechanical claw as goal-directed (and thus cannot *share* their goals), because they are neither perceptually nor motorically human-like. Nor do infants seem to represent the dropping of the back of a human hand onto a toy as goal-directed, because it is not a familiar motoric means of performing a grasping action.

However, not all developmental psychologists subscribe to the shared-intentionality model of goal-ascription. An alternative is the 'teleological stance' model. Unlike advocates of the shared-intentionality model, advocates of the teleological stance model make no assumption about the primacy of infants' familiarity with actions performed by human agents in shaping their ability to ascribe goals. Instead, at the heart of Csibra and Gergely's (1998) and Gergely and Csibra's (2003) teleological stance model is the fundamental assumption that humans are endowed with a ternary computational system that parses an agent's action into three basic components: the agent's *goal*, the *means* selected by the agent for achieving her goal and the situational *constraint* (e.g. an environmental obstacle) faced by the agent. Advocates of the teleological stance model have offered evidence showing that the above ternary system enables preverbal human infants to compute any one of the three components from the other two. They can compute an agent's means-action from a prior representation of both the agent's goal and the constraints faced by the agent. They can compute an agent's goal from a prior representation of both the constraints faced by the agent and the agent's selected means-action. Finally, they can compute the constraints faced by the agent from a prior representation of both the agent's goal and the agent's means-action (cf. Gergely and Csibra, 2003).

Advocates of the teleological stance model have emphasized the contribution of at least two sorts of further cues to infants' goal-ascription. While Premack (1990) argued that the agent's motion must be *self-propelled* (Premack, 1990), Csibra *et al.* (1999) and Biro *et al.* (2007) have argued that this is only a weak cue of goal–directedness. In addition, the agent's action must cause a *novel outcome* or bring about a change from one state of affairs to another (Kiràly *et al.*, 2003; Biro and Leslie, 2007). And the agent must also give evidence that it can produce *equifinal variations* of its behavior: that is, given that the outcome (i.e. final step) of the agent's action is constant, the agent must give evidence of its ability to adapt its behavior in response to environmental changes or changes in situational constraints (Kiràly *et al.*, 2003; Biro and Leslie, 2007). For example, given that an agent expresses a

behavioral preference for one of a pair of toys, the agent will accordingly vary the trajectory of its motion depending on changes of the toy's spatial location.[22]

Two of Woodward's (1998, 1999) basic negative results have been revisited by advocates of the teleological stance model. In the habituation trials of Kiràly *et al.*'s (2003) choice-based experiment, 10-, 8- and 6-month-olds were presented with two toys, one on the right of a stage, the other on the left. Then they saw the back of a human hand repeatedly approach the toy on the left and displace it from its initial position to a new position (towards the back of the stage). In the test trials, the spatial positions of the toys were switched and infants saw the back-of-the-hand action displace either the new toy or the old toy towards the back of the stage. Kiràly *et al.* (2003) report that 10- and 8-month-olds looked reliably longer at the back-of-the-hand action upon the new toy (old side) than upon the old toy (new side). Using exactly the same design, Jovanovic *et al.* (2007) report that 6-month-olds also looked reliably longer at the back-of-the-hand action upon the new toy (old side) than upon the old toy (new side). Woodward (1999) interpreted the failure of 6- and 9-month-olds to ascribe a goal to an agent performing a back-of-the-hand action onto an object in terms of the infants' unfamiliarity with this type of hand-action. But as Kiràly *et al.* (2003) point out, in Woodward's (1998, 1999) seminal choice-based paradigm, one fundamental outcome of a typical goal-directed grasping action is missing: namely, the target's change of spatial position (or displacement). After an agent reaches for and grasps a target, typically her hand does not rest on the target: rather the target is moved to a novel position. In effect, what Kiràly *et al.* (2003) and Jovanovic *et al.* (2007) did was to provide infants in both habituation and test trials with the missing ingredient of a normal grasping action from Woodward's (1998, 1999) experimental paradigm: the back-of-the-hand action *displaced* its target. When the missing ingredient was restored, 9- and 6-month-olds gave evidence of representing the back-of-the-hand action as goal-directed.

Biro and Leslie (2007) report the results of a set of four experiments that cast new light on Woodward's (1998, 1999) negative result with the inanimate rod: these experiments explore the putative impact of adding two cues of goal-directedness to the action of an inanimate rod on a designated target: a salient outcome or change of the target's spatial position, and the rod's displaying equifinal variations of behavior. The familiarization conditions in all four experiments included a presentation of two toys: a ball on the left and a teddy bear on the right. In the familiarization trials of the first experiment, 12-, 9- and 6-month-olds were divided into two groups: half the infants of each age were assigned to the 'poking hand' condition and half to the 'poking tube' condition. In the former, infants saw the experimenter's hand

---

[22] Because Csibra and Gergely (1998), Gergely (2007) and Gergely and Csibra (2003) stress the fact that the teleological representation of an agent's action is reality-based (i.e. non-mentalistic), they emphasize the role of an agent's goal-state, i.e. the physical outcome of the agent's action, in the content of the teleological representation.

(with her index finger extended) poke one and the same toy (i.e. the ball) ten times. In the latter condition, they saw an inanimate wooden tube (or rod) poke one and the same toy (i.e. the ball) ten times. After four familiarization trials, the toys' positions were switched, and in the test trials, half the infants in the poking hand condition saw the hand poke the same toy again at a different location and half saw the hand poke the new toy at the old location. Half the infants in the poking tube condition saw the tube poke the old toy at a new location and half saw the tube poke the new toy at the old location. Infants in all three age groups in the poking hand condition looked reliably longer at the new toy event than at the old toy event. The results were less clear-cut for infants in the poking tube condition (especially for the 9-month-olds). Perhaps the poking tube action was less interesting than the poking hand action. Perhaps the equifinal variations of behavior were less effectively displayed by the poking tube than by the poking hand.

In the second experiment, the tube could lift the toys by means of pieces of Velcro attached to them. In the familiarization condition of the second experiment, infants in the same three age groups as in the first experiment were first presented with an inanimate rod that demonstrated self-propelledness by freely moving around and changing its path in a random fashion. Then the infants saw the tube lift up and then replace one of the toys. In the test trials, the toys' positions were switched and infants saw the tube touch (without lifting) either the same toy or a new toy. Only the 12-month-olds looked reliably longer at the new toy (old location) event. The 9-month-olds also looked longer at this event, but the difference was not significant. The 6-month-olds looked equally at the two events.

The stimuli in the third experiment were the same as those in the second experiment, with one important exception: in the test trials, the tube not only touched but also lifted up either the same toy as in the familiarization trials or the new toy. In the third experiment, Biro and Leslie (2007) only tested 9- and 6-month olds: 9-month-olds looked reliably longer at new toy events, but 6-month-olds looked equally at both the new and the old toy events.

The familiarization trials in the fourth experiment were like those in the third experiment with one important exception: now Velcro pieces were attached to the top, front and back of both toys. As a result, the tube was able to display new cues of equifinal behavior by approaching the target toy from three different angles and picking it up by sticking to three different parts of the designated toy. The test trials were similar to those in the third experiment: this time, 6-month-olds looked reliably longer at new toy events. Thus, 6-month-olds seem to represent the action of an inanimate rod as goal-directed if provided with information about the novelty of the outcome of the rod's action on the object in both the familiarization and the test trials and also with strong cues of equifinal variations of the rod's behavior in the familiarization condition. Importantly, both the target's change of position and the cues of the rod's equifinal variations of behavior were missing in Woodward's (1998) experimental design.

Luo and Baillargeon (2005) also re-examined Woodward's (1998, 1999) negative results. In their first experiment, which involved a box covered with green paper, they compared two conditions: the experimental and the control conditions. The familiarization trials in both conditions were preceded by an 'orientation event' during which 5-month-olds saw the box give evidence of being self-propelled by moving back and forth at the center of the apparatus floor. In both conditions, infants received four familiarization trials. During the familiarization trials in the experimental condition, the infants were presented with a cylinder on the left and a cone on the right and then they saw the box repeatedly move to the right and contact the cone. During the two test trials in the experimental condition, the cone and the cylinder switched positions and the infants saw the box move towards and contact either the cylinder (new target and old side event) or the cone (old target and new side event). During the familiarization trials in the control condition, the infants were presented with only a single object, i.e. the cone on the right, and they saw the box move towards and contact the cone. During the two test trials in the control condition, the cone and the cylinder were (as in the test trials of the experimental condition) respectively on the left and the right and the infants saw the box move towards and contact either the cylinder (new target and old side event) or the cone (old target and new side event). In the experimental condition, 5-month-olds looked reliably longer at the new target event than at the old target event, but in the control condition, they looked equally at both events.

The result of Luo and Baillargeon's (2005) first experiment are difficult to reconcile with the basic assumptions of the shared-intentionality model, which predicts that 5-month-olds should not represent the behavior of a box as goal-directed on the grounds that a box is perceptually and motorically very unlike a human agent. Arguably, the orientation event played an important role in providing infants with cues that the box was a self-propelled agent—cues that were absent from Woodward's (1998) experiment with either the inanimate rod or the mechanical claw. If so, then the question still arises: why did Mahajan and Woodward (2009) find that 7-month-olds' choice of a toy was cued by the choice of a human hand, but not of a box? Arguably, two distinguishing features may account for the different experimental results. First, Mahajan and Woodward (2009) familiarized the infants with the targets, but not with seeing repeated actions of the box directed towards one of a single pair of targets, as did Luo and Baillargeon (2005). Secondly, the cues of the box's self-propelledness contained in a separate orientation event might have been richer in Luo and Baillargeon's (2005) than in Mahajan and Woodward's (2009) experiment.

Importantly, the contrast between the experimental and the control conditions in Luo and Baillargeon's (2005) first experiment corroborates the role of choice information in the infants' representation of the displacements of the self-propelled box as goal-directed: the fact that, in the control condition, 5-month-olds were familiarized with a self-propelled box that repeatedly contacted the same object (i.e. the cone) when there was no competing target was not sufficient to cause them to expect the box to continue to contact the same object in the presence

of a competing target. In the control condition, where no choice information was presented in the familiarization trials, the question arises whether infants represent the action of the box as directed towards the goal of contacting the cone and whether, lacking choice information, they did not rule out the possibility that the box could change its goal in the presence of a new target. I shall return to this issue at the end of the next section, in which I review some further developmental evidence against the model, based on experiments focusing on efficiency rather than choice.

4. Single–Target–Efficiency–Based Evidence Against the Shared–Intentionality Model

Advocates of the teleological stance model have designed another experimental paradigm for exploring the conditions in which preverbal human infants represent an agent's action as goal-directed. I call this paradigm the single-target-efficiency paradigm. In this paradigm, exemplified by Gergely *et al*.'s (1995) and Csibra *et al*.'s (1999) experiments, preverbal human infants are provided with cues of efficiency about the means selected by an agent for performing a goal-directed action on a single target.

12- and 9-month-olds were divided into two groups: the 'rational action' (or experimental) group and the 'non-rational action' (or control) group. In the habituation trials for both groups, infants first saw a large circle on the left and a small circle on the right: first, they saw the large circle expand, contract and regain its original size. Then they saw the small circle perform a similar expansion-contraction sequence. The difference between the habituation trials for each group was that, in the rational action group, infants saw a rectangle located midway between the two circles, whereas in the non-rational action group, infants saw the rectangle on the right of the small circle. Finally, they saw the smaller circle move in a parabolic trajectory from right to left and land next to the larger circle. In the rational action condition, the presence of the rectangle midway between the two circles could be perceived as constraining the leftwards motion of the small circle; but in the non-rational action condition, the rectangle was located to the right, not the left, of the small circle. In the test trials, the rectangle was removed and infants in both conditions saw the small circle either travel through the same parabolic trajectory as before (old action) or move in a straight line (new action), until it reached the large circle. Infants of both ages in the rational action group looked reliably longer at the old action than at the new action. Infants in the non-rational action group did not.

On the one hand, it seems as if the removal of the situational constraint (i.e. the rectangle) was sufficient to trigger the infants' expectation that the agent (i.e. the small circle) would select the shortest (most efficient) path to achieve its goal (i.e. reach the large circle). Even though infants in the experimental condition had already seen the small circle move in a parabolic trajectory when a rectangle was located on its left, they showed more surprise when, in the absence of the rectangle,

it re-enacted the same trajectory than when it moved in a straight line. On the other hand, only if they had been habituated to seeing the agent's parabolic trajectory in the constraining presence of the rectangle did infants expect the agent to select the shortest (most efficient) path.[23]

Kamewari *et al*. (2005) have further explored the relative contributions of efficiency cues and the agent's human–like perceptual appearance to infants' capacity to reason about goals. They performed three experiments with videos displaying the actions of a human agent, a humanoid robotic agent and a box, in each of which 6.5-month-olds were divided into an experimental and a control group. In the habituation trials for each experiment, infants in the experimental group saw an agent take a detour around an obstacle (a box) before reaching his target (another box); but in the control group, infants saw the same agent take a detour without a constraining obstacle. In test trials, infants in both groups saw the agent either take a detour or move in a straight line, in the absence of any obstacle. When the agent was either a human or a humanoid robot, infants in the experimental condition looked reliably longer at the detour path event than at the straight path event, but infants in the control condition did not. When the agent was the box, infants of neither condition did. Kamewari *et al*. (2005, p. 319) conclude that 'around 6 months old, initial understanding of goal-directed action is teleological, being simultaneously influenced by certain perceptual features of actions or knowledge about the agents'.

Interestingly, as Csibra (2008) has noticed, when the agent was the box, but not when the agent was either a human or a humanoid robot, infants in the control group looked significantly longer at the straight path event than at the detour path event. Arguably, as Csibra (2008) further emphasizes, in the habituation trials in Kamewari *et al*.'s (2005) experiment, the box reproduced exactly the same motions on each trial, displaying no equifinal variation in its behavior—as in Gergely *et al*. (1995) and Csibra *et al*. (1999)—and fewer cues of self-propelledness than in Luo and Baillargeon's (2005) orientation events. This lack of equifinal variation and poor cues of self-propelledness may have caused the surprise of 6.5-month-olds when they saw the box change path and move in a rectilinear efficient path in the test condition.

To test the potential role of cues of equifinal behavioral variations in goal-ascription within the single-target-efficiency paradigm, Csibra (2008) compared 6.5-month-olds' responses to two conditions: the single route and the variable route conditions. In both conditions, he divided the infants into an experimental and a control group. In the single route condition, infants from the experimental group were habituated (as in Kamewari *et al*., 2005) to seeing a moving box make a detour around an obstacle and move towards a target (another box). Infants from the control group were habituated to seeing an agent box make a detour without any obstacle and move towards another non-moving box. In the variable

---

[23] Interestingly, Rochat *et al*. (2008) report that macaque monkeys also look reliably longer at the less efficient of a pair of familiar grasping actions (experiment 1), but not if the displayed action is unfamiliar (experiment 3).

route condition, infants from the experimental group were habituated (unlike in Kamewari *et al*., 2005) to seeing a moving box make a variable detour, i.e. by turning sometimes left, sometimes right, around an obstacle, and move towards a target (i.e. another non-moving box). In the test trials of both conditions, infants saw the moving box either make a detour or go straight towards the target box without any obstacle. In the single route condition, 6.5-month-olds from either the experimental or the control group did not look reliably longer at the detour path event than the single path straight path event. In the variable route condition, infants from the experimental group did, but infants from the control group did not.

Further experiments in the single-target-efficiency paradigm by Southgate *et al*. (2008) corroborate the crucial role of efficiency cues in infants' goal-ascriptions. Southgate *et al*. (2008) divided 6- to 8-month-olds into an experimental and a control group. Infants in the experimental group were familiarized with seeing a video showing a human hand perform a two-step goal-directed action during which the hand removed a box lying in its path before it retrieved a target. In the control group, infants were familiarized with seeing a video showing a human hand inefficiently remove a box that did not lie in its path before retrieving a target. In the test trials, infants in both groups saw a human hand perform either the same hand action as infants in the experimental group had seen during the familiarization trials, or a biologically impossible but more efficient action whereby it retrieved the target by snaking around the obstacles. Infants in the experimental group, but not infants in the control group, looked reliably longer at the biologically possible but less efficient hand-action than at the more efficient but biologically impossible hand-action. This last result suggests that cues of efficiency act as a potent guide to goal-ascription by preverbal human infants. Moreover it is also difficult to accommodate within the shared-intentionality model because it further suggests that infants' ability to make use of efficiency cues for goal-ascription may precede, and even override, their expectations about the contrast between biologically possible and impossible human actions, based on perceptual and motor familiarity with human agents (in accordance with the first three assumptions of the model).

That the use of efficiency cues may precede and override choice information is also strongly suggested by results reported by Biro and Verschoor (2008), who compared the responses of 9- and 7-month-olds in three conditions. In the familiarization events of the 'efficient goal condition', infants saw a human hand reach through the side of the stage, remove the lid of one of a pair of boxes (each of which contained a toy) and grasp the toy inside. In the familiarization events of the 'non-efficient goal condition 1', infants saw the hand remove the lid of one of the boxes before grasping one of the toys, but this time the toys were in front of, not inside, the box. In the familiarization events of the 'non-efficient goal condition 2', no box was present and the hand performed a mime of lifting the lid of an absent box before grasping one of the two toys. In the test trials of all conditions, the positions of the toys were switched and infants saw the human hand grasp either the old toy (at a new location) or a new toy (at an old location). Biro and Verschoor (2008) report that in the efficient goal condition, both 9- and

7-month-olds looked reliably longer at the new toy (old location) event than at the old toy (new location) event; but they did not in either of the non-efficient conditions. If so, then efficiency cues override preference information.

Further evidence against the shared-intentionality model has recently been adduced by Luo (2011), who has adapted the control condition of Luo and Baillargeon's (2005) first experiment. In her first experiment, Luo (2011) replicated the contrasting results obtained by Luo and Baillargeon in their experimental and control conditions, but with 3-month-olds: in the experimental, but not the control condition, infants had choice information about a box repeatedly contacting one of two targets. Like 5-month-olds, 3-month-olds looked reliably longer at the box's new target action than the old target action in the experimental condition, but not in the control condition.

During the familiarization trials of her novel second experiment, Luo (2011) offered 3-month-olds cues of both self-propelledness and equifinal variations of the agent's behavior in the single-target-efficiency paradigm: after an orientation event in which a box gave evidence of being self-propelled, infants saw the box repeatedly contact a single target (without any competing target) by adapting the direction of its motion (from right to left and vice-versa) to the target's changing spatial position. In the test trials, infants saw the box contact either the same target (old target action) or a new target (new target action). 3-month-olds looked reliably longer at new rather than old target actions.

To sum up, by adding cues of equifinal variations of a box's behavior (which were missing from the experiment by Kamewari *et al.* 2005), Csibra (2008) showed that 6.5-month-olds ascribe a goal to a moving box that makes a variable detour around an obstacle lying on its path towards its target. Southgate *et al.* (2008) strongly suggest that the ability of 6- to 8-month-olds to make use of efficiency cues for goal-ascription precedes and even overrides their expectations about biologically possible human hand actions. Taken all together, the experimental results of Luo and Baillargeon (2005) and Luo (2011) show that 3-month-olds are disposed to ascribe a goal to a self-propelled box if either of two requirements is met: if the box displays either choice information (about one of two competing targets) or equifinal variations of behavior (by adapting the direction of its displacements to the target's changing positions). Clearly, given the absence of perceptual and motoric likeness between a moving box and either a human agent or her bodily parts, none of these goal-ascriptions to a self-propelled box could have been predicted by the shared-intentionality model.

As Luo and Baillargeon (2005) and Luo (2011) show, 5- and 3-month-olds do not look longer at a new target action performed by a self-propelled box if they are not provided with either choice information or cues of equifinal variations of the box's behavior. So the question arises: would the perceptual and motoric familiarity of preverbal human infants with human hand actions dispose them to represent a human hand action as goal-directed in the absence of both choice information and cues of equifinal variations of the hand's action? In fact, the preliminary question arises of how (in the control condition) infants represent the action of the

self-propelled box towards a single target when they lack both choice information and cues of equifinal variations of behavior.

According to Luo and Baillargeon's (2005, pp. 604-5) own interpretation, 'they (a) categorized the box as an agent, (b) interpreted the box's actions during famil-iarization as directed toward the goal of contacting the cone, and (c) recognized that the box might change its goal and hence its actions after the cylinder was introduced'. On this interpretation, infants do represent the box's action towards a target as goal-directed, but lacking choice information in the familiarization trials, they simply suspend their judgment on the box's goal when a second target is intro-duced in the test trials. In other words, preference information (provided to infants in the experimental group, but not in the control group) could be said to modulate goal-ascription. However, the results of Luo's (2011) second experiment somewhat weaken Luo and Baillargeon's (2005) interpretation, since by introducing cues of equifinal variations of the box's behavior without introducing any choice informa-tion, Luo (2011) caused 3-month-olds to look longer at a new target action than at an old target action performed by the box. According to an alternative interpretation, the very fact that infants do not look longer at a new target action performed by a self-propelled box if they are not provided with either choice information or cues of equifinal variations shows that they fail to represent the box's action as goal-directed.

Further evidence reported by Hernik and Southgate (submitted) is relevant to the question of how infants represent the action of a self-propelled box directed towards a single target, and to assessing Luo and Baillargeon's (2005) claim that in the absence of a competing target, infants do represent the action of a box towards a single target as goal-directed, but simply suspend judgment on the box's goal when a second target is introduced in the test trials. Hernik and Southgate divided 9-month-olds into an experimental and a control group. In the familiarization trials, infants in the experimental group saw a box take a detour around a wall in order to contact a target located behind the wall on its left side. Infants in the control group saw a box contact a target by making the same detour without any obstacle. Then during the 'orientation trials', infants from the experimental group saw two targets behind a wall, one of which was the same as in the familiarization trial, but now on the right side of the wall. Infants from the control group saw two targets (without any wall), one of which was the same, but not in the same spatial position, as in the familiarization trials. Finally, infants from the experimental group saw the box perform either a new target (old path) action or a new path (old target) action in the presence of the wall. Infants from the control group saw the box perform either a new target (old path) action or a new path (old target) action in the absence of the wall. Hernik and Southgate (submitted) report that infants in the experimental group, but not in the control group, looked reliably longer at new target (old path) than new path (old target) actions of the box.

These results suggest that cues of efficiency in the absence of choice information were sufficient to allow goal-attribution to the box by infants in the experimental group. But they also suggest that lacking both preference information and efficiency cues, infants in the control group might simply have failed to represent the

action of the box as goal-directed. If so, then contrary to Luo and Baillargeon's (2005) interpretation of their own first experiment, the fact that infants from the control group saw the box perform a simple action of contacting a target during the familiarization trials did not cause them to ascribe a goal to the box. It would follow that preference information (made available during the familiarization trials to infants in the experimental group, but not in the control group) cannot modulate goal-ascription for the simple reason that, lacking both efficiency cues and choice information, infants in the control group did not attribute a goal to the box.

On the face of it, a simple hand action directed towards a target should be perceptually and motorically far more familiar to human infants than the action of a box. The question arises whether human infants are disposed to represent a human hand action as goal-directed in the absence of both choice information and cues of equifinal variations of the hand's action. Clearly, the shared-intentionality model should lead us to expect a positive answer to the question. However, this expectation is not supported by results of experiments performed by Biro *et al*. (submitted) on 12-month-old human infants, whose responses Biro *et al*. compared in three conditions. In the familiarization trials of the 'efficient action' condition, infants saw a human hand first open a transparent box and grasp a teddy bear located inside the box. In the familiarization trials of the 'non-efficient action' condition, infants saw the human hand perform the same action as in the efficient action condition, but the teddy bear was in front of (not inside) the box. In the familiarization trials of the 'simple action' condition, there was no box and infants saw the hand simply grasp the teddy bear. In all three conditions, the grasping hand came out of an opening in a wall to the right of the scene and the teddy bear was located near the wall. Then in the pre-test condition, all infants saw the teddy bear now on the left and another toy on the right. Finally, in the test trials, all infants saw the human hand either move along the same path and grasp a new toy or move along a new path and grasp the same teddy bear. Biro *et al*. (submitted) report that in the efficient action condition, infants looked reliably longer at the new toy (old path) action than at the old toy (new path) event, whereas in the non-efficient condition or the simple action condition, they did not. In accordance with assumptions (1)-(3) of the shared-intentionality model, it might have been expected that, in the absence of choice information and cues of efficiency, preverbal human infants would still represent a simple hand action, unlike the action of a box, as goal-directed. But the results of Biro *et al*.'s (submitted) experiment contradict this expectation.

5. Why Should Performing Goal-Directed Actions Help Infants to Ascribe Goals To Others?

Arguably, the shared-intentionality model derives support from further experiments also within the choice-based paradigm, reported by Sommerville and Woodward

(2005). In the habituation condition of the first study, 12-month-olds were presented with a pair of toys sitting on different-colored cloths: e.g. a duck on a yellow cloth and a tortoise on a blue cloth. The infants were habituated to seeing a human agent repeatedly pull the yellow cloth and retrieve the duck. In the test trials, the toys' positions were switched so that the duck now sat on the blue cloth and the tortoise on the yellow cloth. Infants saw the human agent either retrieve the same toy (i.e. the duck) by pulling the new (blue) cloth or the new toy (i.e. the tortoise) by pulling the old (yellow) cloth. 12-month-olds looked reliably longer at new toy (old cloth) trials than at new cloth (old toy) trials, indicating that they interpreted the cloth pulling as a means selected by the agent to achieve her goal of retrieving one of the two toys. In the same conditions, however, 10-month-olds as a group did not look longer at the new toy (old cloth) trials than the new cloth (old toy) trials.

Sommerville and Woodward (2005) also report the results of a second study designed to investigate the impact of infants' first-person motor experience upon their ability to understand means-end goal-directed actions performed by others. Sommerville and Woodward (2005) found that ten-month-olds who were able to *produce* more well-organized means-end actions looked reliably longer at new toy (old cloth) trials than at new cloth (old toy) trials. Similarly, Sommerville *et al.* (2005) report that as a group, 10-month-olds, who had been habituated to seeing a human agent repeatedly retrieve one of two toys by pulling a cane with her right hand, did not look longer in the test trials when the human agent pulled the same cane with her right hand to retrieve a new toy than when she pulled a novel cane with her left hand to retrieve the old toy. However, they report that 10-month-olds who had been *actively trained* to use a cane to retrieve an out-of-reach toy did look reliably longer at new toy (old cane) trials than at new cane (old toy) trials.

Sommerville *et al.* (2005) report that, after being habituated to seeing a human hand in a grasping position repeatedly reach for, and rest on, one of two toys, 3-month-olds, unlike 5-month-olds, did not look longer at new toy (same side) trials than at old toy (new side) trials. 3-month-olds also happen to be unable to accurately reach and grasp objects. Sommerville *et al.* (2005) supplied 3-month-olds with Velcro mittens that enabled them to accurately grasp objects. They divided 3-month-olds into two groups: the reach-first group and the watch-first group. While infants in the reach-first group were allowed to reach for objects (with Velcro mittens) before being habituated to seeing a human hand in a grasping posture repeatedly reach for and rest on one of two toys, infants in the watch-first group were presented with habituation trials before being allowed to reach for and grasp objects. Sommerville *et al.* (2005) found that whereas infants in the watch-first group did not look longer at new toy (old side) events than at old toy (new side) events, infants in the reach-first group did so reliably. Sommerville *et al.* (2005) also report that infants in the reach-first condition looked significantly longer on the first habituation trial than infants in the watch-first condition, suggesting that their own preceding motor experience enhanced their attention to

another's goal-directed action. Furthermore, Gerson and Woodward (submitted) report that unlike 3-month-olds who were given the opportunity to execute (with Velcro mittens) grasping actions (active training), 3-month-olds who were given the opportunity to observe others grasp toys (observational training) did not look longer at new toy than old toy events (cf. Woodward *et al.*, 2009). As Rochat *et al.* (2008, p. 231) describe these findings, 'motor familiarity, but not perceptual familiarity, influences the capacity of 3-month-old infants to extract goals from observed action'.

Thus, the insistence of the shared-intentionality model on the role of infants' motor experience in goal-ascription seems vindicated. Moreover, as mentioned above, some of Woodward's (1998, 1999) negative behavioral results in human infants seem to mirror some of the negative results of single-cell recording in area F5 of the monkey. It is therefore not surprising that advocates of the shared-intentionality model have turned to mirroring (or MN activity) as a likely neural mechanism linking the execution and perception of action. As Hamlin *et al.* (2008) write, 'it has been proposed that [these] mirror representations exist early in life and play a role in development [...] Consistent with this proposal, recent experiments have documented relations between infants' ability to produce well-organized goal-directed actions and their comprehension of others' actions as goal-directed [...] These two abilities are correlated [...] and interventions that alter infants' action production affect their subsequent perception of others' actions'.[24]

Sommerville *et al.* (2005) provide convincing evidence that, unlike 3-month-olds who cannot accurately grasp objects, 3-month-olds who were enabled to accurately grasp a toy using Velcro mittens looked longer at new target (old side) than at old target (new side) actions. Thus, first-personal motor experience of reaching and grasping helps human infants to represent others' actions as goal-directed. This result seems to call for a mechanism able to couple the execution of grasping actions and the perception of grasping actions performed by others. However, as I will now argue, for two converging reasons MN activity is not likely to be the mechanism that accounts for why motor experience promotes infants' understanding of others' goal-directed actions, as exemplified in Sommerville *et al.* (2005).

Consider first the surprising result reported by Biro *et al.* (submitted) in their simple action condition (discussed at the end of Section 4). MN activity is elicited in non-human primates by the perception of simple hand actions directed towards a target. So MN activity should be elicited in human infants by the perception of a simple hand action directed towards a single target. If so, then it would seem to be predicted by the shared-intentionality model that when infants are familiarized with seeing simple hand actions repeatedly directed towards a single target (e.g. a teddy bear), they should represent the hand action as being goal-directed. In other words, infants in the simple action condition of Biro *et al.* (submitted) should look longer at new target than at old target hand actions. But Biro *et al.* (submitted)

---

[24] Cf. Del Giudice *et al.*, 2009.

report that they do not. So on the assumption that MN activity is being elicited in 12-month-olds in the simple action condition of Biro *et al*. (submitted), their negative result suggests that MN activity does not underlie infants' goal-ascription as tested in the violation–of–expectation paradigm by comparing new target and new path events.

I now turn to Sommerville *et al*.'s (2005) finding about the role of infants' motor experience in facilitating their representation of others' actions as goal-directed. Arguably, MN activity could be involved in the intervention stage of Sommerville *et al*.'s (2005) experiments when they enable infants to execute grasping by providing them with Velcro mittens. Indeed, MN activity has been shown to be elicited in monkeys by both the execution and the perception of hand actions directed towards a target. However, it is highly unlikely that MN activity is also at work in either the habituation/familiarization trials or the test trials of Sommerville *et al*.'s (2005) experiment. In accordance with Woodward's (1998) choice-based paradigm, what Sommerville *et al*. (2005) take as evidence for the representation of an agent's action as goal-directed is the fact that after being habituated to (or familiarized with) seeing an agent repeatedly choose one of two targets, infants are more surprised by a new target event than by a new path event. MN activity could be the common neural mechanism that accounts for why motor experience enhances infants' representation of others' actions as goal-directed only if it were at work during both the execution of grasping actions by infants and their perception of a human hand that repeatedly chooses one of two targets in the habituation/familiarization trials. But there is no evidence for the assumption that the perception of an agent's choice between two competing targets should elicit MN activity in either the agent's or the observer's brain.

If MN activity is not the mechanism coupling infants' execution of grasping actions and their perception of another's hand repeatedly choosing one of two targets, then what explains the role of infants' motor experience in goal attribution? Arguably, what the execution of acts of grasping (made possible by the use of Velcro mittens) may provide 3-month-olds with is *outcome* information, missing from their perception of the actions performed by others in Woodward's choice-based paradigm, in which the hand reaches its target in a grasping posture without lifting it. Moreover, as reported by Gerson and Woodward (submitted), only infants involved in active training, not in observational training, looked longer at new toy events than old toy events. So the question is: what relevant information is being carried by the execution of an act of grasping a toy that is not carried by observing the same act executed by another? Arguably, whereas observing an act of grasping provides the observer with purely visual information about the target's displacement, executing an act of grasping a toy also provides an infant with proprioceptive, haptic and perhaps even efference copy information about the outcome of the action. If so, then the evidence for the role of motor experience in infants' capacity to reason about goals is also evidence for asymmetry, rather than symmetry, between being an agent and being the observer of an action performed by another.

Concluding Remarks

In Sections 3-4 of this paper, I have reviewed developmental evidence that is hard to square with assumptions (1)-(3) of the shared-intentionality model, which claims that the early capacity of human infants to reason about an agent's goals is shaped by the human-like perceptual appearance of the agent and the fact that the agent's motor act falls within the infant's motor repertoire. In Sections 1 and 2, I have argued on conceptual and empirical grounds that assumptions (4) and (5) of the shared-intentionality model overlook both the complexity intrinsic to the mental representation of a goal and the gap between having and ascribing a goal. In Section 5, I have further argued that the developmental evidence showing that infants' motor experience enhances their ability to reason about goals should not be interpreted as evidence for the role of MN activity in goal-ascription by human infants.

As I argued in Sections 3-4, the evidence does not seem to vindicate the crucial role assigned to the perceptual and motoric interpersonal likeness between infants and the observed agents of goal-directed actions by advocates of the shared-intentionality model. However, there is another dimension of interpersonal similarity between an infant and an agent of a goal-directed action that seems to be overlooked by advocates of the shared-intentionality model of goal-attribution. Suppose that, as predicted by the shared-intentionality model, when preverbal human infants see an agent efficiently perform a goal directed-action, they come to share the agent's goal. If so, then, like an efficient agent, they must have the ability to form goals, even if they still lack the capacity to perform the action efficiently. (It is one thing to lack the executive capacity for efficient action. It is another to lack the capacity to represent and process information about the efficiency of an action performed by another.) In that case, they must be able to conceptually represent an action as an efficient means towards bridging the gap between two states of affairs (cf. Section 2). This conceptual ability would still fall short of possession of the concept GOAL (cf. Section 1). However, by giving evidence that they can not only form goals but also reason about goals and ascribe a goal to an agent (i.e. form beliefs or judgments about the goals of others), it follows from my arguments in Section 1 that infants also give evidence that they possess a (rudimentary) concept GOAL.

Institut Jean Nicod
ENS/EHESS/CNRS

References

Anscombe, G. E. 1957: *Intention*. Oxford: Blackwell.

Biro, S., Csibra, G. and Gergely, G. 2007: The role of behavioral cues in understanding animacy, agency and goal-directed actions in infancy. In C. von Hofsten and

K. Rosander (eds), *Progress in Brain Research: From action to cognition, Vol. 164.* Amsterdam: Elsevier, 303−22.

Biro, S. and Leslie, A. 2007: Infants' perception of goal-directed actions: development through cue-based bootstrapping. *Developmental Science*, 10, 379−98.

Biro, S. and Verschoor, S. 2008: The role of efficiency of action in infants' goal attribution. Poster presented at ICIS, Vancouver, 2008.

Biro, S., Coenen, L. and Verschoor, S. Submitted: Evidence for a unitary goal concept: 12-month-old infants are able to transfer goal-directed interpretations between situations with different types of information regarding the goal.

Bratman, M. E. 1992: Shared cooperative activity. *The Philosophical Review*, 101, 327−41.

Csibra, G. 2007: Action mirroring and action understanding. In P. Haggard, Y. Rossetti & M. Kawato (eds), *Sensorimotor Foundations of Higher Cognition. Attention and Performance XXII.* Oxford: Oxford University Press.

Csibra, G. 2008: Goal attribution to inanimate agents by 6.5-month-old infants. *Cognition*, 107, 705−17.

Csibra G. and Gergely, G. 1998: The teleological origins of mentalistic action explanations: a developmental hypothesis. *Developmental Science*, 1, 255−59.

Csibra, G., Gergely, G., Bíró, S., Koós, O. and Brockbank, M. 1999: Goal attribution without agency cues: the perception of 'pure reason' in infancy. *Cognition*, 72, 237−67.

Decety, J. and Grèzes, J. 2006: The power of simulation: imagining one's own and other's behavior. *Brain Research*, 1079, 4−14.

Decety, J. and Sommerville, J. A. 2003: Shared representations between self and other: a social cognitive neuroscience view. *Trends in Cognitive Sciences*, 7(12), 527−33.

Del Giudice, M., Manera, V. and Keysers, C. 2009: Programmed to learn? The ontogeny of mirror neurons. *Developmental Science*, 12, 350−63.

Ferrari P.F., Rozzi, S. and Fogassi, L. 2005: Mirror neurons responding to observation of actions made with tools in monkey ventral premotor cortex. *European Journal of Neuroscience*, 17, 1703−14.

Gallese, V. 2001: The 'Shared Manifold' hypothesis, from mirror neurons to empathy. *Journal of Consciousness Studies*, 8(5−7), 33−50.

Gallese, V., Keysers, C. and Rizzolatti, G. 2004: A unifying view of social cognition. *Trends in Cognitive Sciences*, 8, 396−403.

Gergely, G. 2007: Learning 'about' versus 'from' other minds: human pedagogy and its implications. In P. Carruthers, S. Laurence and S. Stich (eds), *The Innate Mind: Foundations and the Future.* Oxford: Oxford University Press, 170−98.

Gergely, G. and Csibra, G. 2003: Teleological reasoning about actions: the naïve theory of rational actions. *Trends in Cognitive Sciences*, 7, 287−92.

Gergely, G., Nadasdy, Z., Csibra, G. and Biro, S. 1995: Taking the intentional stance at 12 months of age. *Cognition*, 56, 165–93.

Gerson, S. and Woodward, A. Submitted: What's in a mitten? The effects of active versus passive experience on infants' attribution of goals.

Gilbert, M. 1989: *On Social Facts*. London: Routledge.

Goldman, A.I. 2009: Mirroring, simulating and mindreading. *Mind & Language*, 24, 235–52.

Hamlin, J. K., Hallinan, E. V. and Woodward, A. L. 2008: Do as I do: 7-month-old infants selectively reproduce others' goals. *Developmental Science*, 11, 487–494.

Hernik, M. and Southgate, V. Submitted: Does preference information generate, rather than modulate, goal attribution in 9-month-old infants?

Jackson, P. L., Rainville, P. and Decety, J. 2006: To what extent do we share the pain of others? Insight from the neural bases of pain empathy. *Pain*, 125, 5–9.

Jacob, P. 2008: What do mirror neurons contribute to human social cognition? *Mind & Language*, 23, 190–223.

Jacob, P. 2009: The tuning-fork model of social cognition. *Consciousness and Cognition*, 18, 229–43.

Jacob, P. and Jeannerod, M. 2003: *Ways of Seeing*. Oxford: Oxford University Press.

Jeannerod, M. 1999: To act or not to act: perspectives on the representation of action. *Quarterly Journal of Experimental Psychology*, 52A, 1–29.

Jeannerod, M. 2006: *Motor Cognition, What Actions Tell the Self*. Oxford: Oxford University Press.

Jovanovic, B., Király, I., Elsner, B., Gergely, G., Prinz, W. and Aschersleben, G. 2007: the role of effects for infants' perception of action goals. *Psychologia*, 50(4), 273–90.

Kamewari, K., Kato, M., Kanda, T., Ishiguro, H. and Hiraki, K. 2005: Six-and-a-half-month-old children positively attribute goals to human action and to humanoid-robot motion. *Cognitive Development*, 20, 303–20.

Kiràly, I., Jovanovic, B., Prinz, W. Aschersleben, G. and Gergely, G. 2003: The early origins of goal attribution in infancy. *Consciousness and Cognition*, 12, 752–69.

Luo, Y. 2011: Three-month-old infants attribute goals to a non-human agent. *Developmental Science*, 14, 453–60.

Luo, Y. and Baillargeon, R. 2005: Can a self-propelled box have a goal? Psychological reasoning in 5-month-old infants. *Psychological Science*, 16(8) 601–8.

Mahajan, N. and Woodward, A. 2009: Seven-month-old infants selectively reproduce the goals of animate but not inanimate agents. *Infancy*, 14, 667–79.

Meltzoff, A. N. 2005: Imitation and other minds: The 'like me' hypothesis. In S. Hurley and N. Charter (eds), *Perspectives on Imitation: From Neuroscience to Social science*, Vol. 2. Cambridge, MA: MIT Press, 55–77.

Meltzoff, A. N. 2007: 'Like me': a foundation for social cognition. *Developmental Science*, 10, 126–34.

Pacherie, E. and Dokic, J. 2006: From mirror neurons to joint actions. *Cognitive Systems Research*, 7, 101–12.

Premack, D. 1990: The infant's theory of self-propelled objects. *Cognition*, 36, 1–16.

Rakoczy, H. and Tomasello, M. 2007: The ontogeny of social ontology: steps to shared intentionality and status functions. In S. L. Tsohatzidis (ed.), *Intentional Acts and Institutional Facts: Essays on John Searle's Social Ontology*. Berlin: Springer Verlag, 113–37.

Rizzolatti, G. and Craighero, L. 2004: The mirror neuron system. *Annual Review of Neuroscience*, 27, 169–92.

Rizzolatti, G., Fadiga, L., Gallese, V. and Fogassi, L. 1996: Premotor cortex and the recognition of motor activities. *Cognition Brain Resources*, 3, 131–41.

Rizzolatti, G., Fogassi, L. and Gallese, V. 2001: Neurophysiological mechanisms underlying the understanding and imitation of action. *Nature Reviews Neuroscience*, 2, 661–70.

Rizzolatti, G. and Sinigaglia, C. 2008: *Mirrors in the Brain: How Our Minds Share Actions and Emotions*. Oxford: Oxford University Press.

Rizzolatti, G. and Sinigaglia, C. 2010: The functional role of the parieto-frontal mirror circuit: interpretations and misinterpretations. *Nature Reviews Neuroscience*, 11(4) 264–74.

Rochat, M. J., Serra, E., Fadiga, L. and Gallese, V. 2008: The evolution of social cognition: goal familiarity shapes monkeys' action understanding. *Current Biology*, 18, 227–32.

Searle, J. R. 1983: *Intentionality. An Essay in the Philosophy of Mind*. Cambridge: Cambridge University press.

Searle, J. R. 1990: Collective intentions and actions. In P. R. Cohen, J. Morgan and M. E. Pollack (eds), *Intentions in Communication*. Cambridge, MA: MIT Press.

Searle, J. R. 1995: *The Construction of Social Reality*. London: Allen Lane.

Singer, T. and Lamm, C. 2009: The social neuroscience of empathy. *Annals of the New York Academy of Sciences,* 1156, 81–96.

Sommerville, J. A. and Woodward, A. L. 2005: Pulling out the intentional structure of human action: the relation between action production and processing in infancy. *Cognition*, 95, 1–30.

Sommerville, J. A., Woodward, A. L. and Needham, A. 2005: Action experience alters 3-month-old infants' perception of others' actions. *Cognition*, 96, B1–B11.

Southgate, V., Johnson, M. H. and Csibra, G. 2008: Infants attribute goals to even biologically impossible actions. *Cognition*, 107, 1059–69.

Tomasello, M. 2008: *Origins of Human Communication*. Cambridge, MA: MIT Press.

Tomasello, M., Carpenter, M., Call, J., Behne, T. and Moll, H. 2005: Understanding and sharing intentions: the origins of cultural cognition. *Behavioral and Brain Sciences*, 28, 675–91.

Tomasello, M. and Carpenter M. 2007: Shared intentionality. *Developmental Science*, 10, 121–25.

Tuomela, R. 1995: *The Importance of Us*. Stanford, CA: Stanford University Press.

Umiltà, M.A., Kohler, E., Gallese, V., Fogassi, L., Fadiga, L., Keysers, C. and Rizzo-latti, G. 2001: 'I know what you are doing': a neurophysiological study. *Neuron*, 32, 91–101.

Umiltà, M.A., Escola, L., Intskirveli, I., Grammont, F., Rochat, M., Caruana, F., Jezzini, A., Gallese, V. and Rizzolatti, G. 2008: When pliers become fingers in the monkey motor system. *Proceedings of the National Academy of Sciences*, 105(6) 2209–13.

Woodward, A. L. 1998: Infants selectively encode the goal object of an actor's reach. *Cognition*, 69, 1–34.

Woodward, A. L. 1999: Infants' ability to distinguish between purposeful and nonpur-poseful behaviors. *Infant Behavior and Development*, 22, 145–60.

Woodward, A. L., Sommerville, J. A., Gerson, S., Henderson, A. M. E. and Buresh, J. 2009: The emergence of intention attribution in infancy. In B. Ross (ed.), *The Psychology of Learning and Motivation*, Vol. 51. Burlington. MA: Academic Press, 187–222.