5 Collective intentionality and the roots of human societal life

Hannes Rakoczy

What can we learn from animals? In this chapter, I would like to pursue the more specific question: What can we learn from other animals about what kind of social animals we are, and how we become so? Obviously, human sociality is quite unique. Our cooperative, societal, and institutional forms of life clearly set us apart from the rest of the animal kingdom. From an ontogenetic point of view, I will inquire into the potential cognitive underpinnings of such unique sociality and its development. The ontogeny of different forms of intentionality in early childhood will be traced, with a comparative eye on common primate and uniquely human aspects. The picture that emerges will be this: We share with other animals, in particular great apes, basic forms of individual intentionality, and probably even simple forms of individual second-order intentionality that develop in human ontogeny in the course of the first one and a half years. What lies at the heart of uniquely human cognition, though, and what lays the foundation for uniquely human sociality, is the ability to enter into collective "WE"-intentionality, which develops from the second year on.

INDIVIDUAL INTENTIONALITY

Intentionality in the broad philosophical sense of "aboutness" is the mark of the mental (Brentano, 1874/1973; Dennett & Haugeland, 1987; Searle, 1983). To be capable of mentality is to be able to entertain intentional, contentful attitudes (beliefs, desires, intentions, etc.) towards the world and to be guided by these in reasoning and rational action.

A central assumption in the Cartesian tradition was that there was a fundamental divide between different kinds of creatures that explains the radical differences between their forms of life: On the one hand, there were humans with their monopoly on intentionality, which enabled them to speak languages and live in culture and societies. And on the other, there were the rest of the animals—mere soulless automata, brutes without culture.

Needless to say, this simple picture is not taken seriously any more. Without Cartesian ontological substance dualism, all animals—humans

included—are physical beings and thus automata at some (physical) level of description. But on another, folk psychological level of description, humans are correctly characterized in intentional terms—and so are some other animals, as they display systematic and differentiated behaviour that from a folk psychological standpoint or an "intentional stance" (Dennett, 1987) provides grounds for ascribing at least simple forms of intentionality (see Bermudez, 2003; Hurley, 2003).

Let me mention two relevant examples from a developmental psychological perspective that help to bring out the analogy between the intentionality developed early in human ontogeny and intentionality in other species, namely object cognition, and planned action. In developmental psychology since Piaget's (1952) seminal work, thinking about an objective world—in its simplest form: thinking about objects existing as "out there"—and acting intentionally and in planned ways have been stressed as the two major milestones in the transition from purely sensorimotor dealings with the world to intentionality proper. All thinking requires a minimal notion of objectivity: The objects thought about exist independently from the perceiver and enduringly out there in the world. Regarding human ontogeny, Piaget has described infants' development from initial undifferentiated sensation without any notion of persisting objects ("out of sight, out of mind") to what he called "object permanence"—the appreciation that objects continue existing objectively whether perceived or not. In their actions infants begin to display object permanence from (at latest) the end of their first year; they begin to search for occluded and hidden objects they previously perceived. Furthermore, infants from around 1 year not only track objects as chunks of matter continuously existing in space and time; they also individuate objects as objects of certain kinds, e.g., this chair, that table, that rabbit. Recent findings suggest that by 1 year of age infants begin to apply our common sense metaphysical framework of objects as enduring substances, individuated under sortal (kind) concepts—and thus share the rudiments of our adult conceptual architecture of objective thought (van de Walle, Carey, & Prevor, 2000; Xu & Baker, 2005; Xu & Carey, 1996).

Many other animals are on a par with infants; many primate species, and dogs, for example, reach the highest levels of Piagetian object permanence, levels typically reached by infants in the second year (reviewed in Tomasello & Call, 1997). Recent research suggests that some monkeys and great apes also individuate objects qua objects of certain kinds much in the same ways as human 1-year-olds do (Mendes, Rakoczy, & Call, 2008; Phillips & Santos, 2007; Santos, Sulkowski, Spaepen, & Hauser, 2002).

The second milestone in the development of intentionality in human ontogeny stressed by Piaget is the emergence of intentional, planned action. While much behaviour may be voluntary right from the start, the first clear instances of intentional instrumental action, i.e., actions done purposefully and in a planned way in order to achieve some end in mind, appears in human ontogeny towards the end of the first year: Infants organize their behaviour in

means–ends structures and indicate an awareness of the relations between means and ends. In a classic example, infants remove barriers in order to reach a desired object or pull a cloth, on which a desired object is placed, towards them in order to be able to grasp it. And they persist until they achieve their end, varying their means if necessary (Piaget, 1952; Willats, 1985, 1999). These phenomena are also widespread in the nonhuman animal kingdom. Many species, notably primates, show instrumental problem-solving of remarkable complexity—Köhler's (1926) apes perhaps being the most famous examples.

In sum, many animals share with us the bare bones of simple individual intentionality.[1] Like human infants from around 1, many animals are capable of the most basic form of objective thought: perceiving and cognizing about objects. And many animals don't just behave, but perform intentional instrumental acts in planned ways. Not to mention the remarkable cognitive abilities recent research has found in many species concerning, for example, causal reasoning, self-recognition, tools use, memory, simple numerical cognition, etc.

SECOND-ORDER INTENTIONALITY

Whatever it is that explains the uniqueness of human cultural, linguistic, and societal life forms in contrast to all other species' forms of life, it cannot be that humans have intentionality, while other animals lack intentionality of any kind. So what is the crucial dividing line? One prominent candidate in comparative research in the last decades was not individual intentionality as such, but the ability to understand others and oneself qua intentional beings—second-order intentionality, also called "theory of mind" after Premack and Woodruff's seminal paper "Does the Chimpanzee have a Theory of Mind?" (1978). So-called "theory of mind" research became a booming field in developmental and comparative psychology when philosophers and psychologists joined efforts to find suitable operationalizations of second-order intentionality (Bennett, 1978; Dennett, 1978; Harman, 1978; Wimmer & Perner, 1983). The ability to ascribe epistemic subjectivity to others and to oneself in the past, i.e., to attribute intentional attitudes that represent reality as being a certain way—which aim at truth but may be false (paradigmatically, false beliefs), emerged as the accepted milestone for full-fledged second-order intentionality. Empirically, it turned out that around 4 years of age in human ontogeny a social-cognitive "revolution" occurs; children begin to manifest a suite of new behaviours. They ascribe false beliefs to others (and themselves in the past) and explain and predict their actions accordingly (see Wellman, Cross, & Watson, 2001, for a meta-analysis). They distinguish appearances from reality (Flavell, Flavell, & Green, 1987) as well as the different observers' conflicting perspectives on the same situation (Flavell, Flavell, Green, & Wilcox, 1981; Perner, 1991). And they begin to intentionally deceive others, i.e., lead them to have false beliefs (Sodian, 1991).

Clearly, second-order intentionality of this kind is crucial to many characteristically human activities and achievements such as self-consciousness and reflective thinking, full-fledged communication (on Gricean analyses), and complex conventional activities (Lewis, 1969). And it seems quite clear and with (almost) consensus in the field that no other species, not even chimpanzees, attain these sophisticated levels of second-order intentionality (see, e.g., Call & Tomasello, 1999).

Soon, however, simpler forms of second-order intentionality like understanding perception and intention received attention, shifting the focus from full-fledged epistemic subjectivity (in particular, false beliefs) to simpler intentional attitudes. Recourse to John R. Searle's taxonomy of intentionality might help to clarify this issue. Searle (1983), following Anscombe (1957), distinguishes two kinds of intentional attitudes that have close analogues in different kinds of speech acts. First, there are cognitive attitudes with "mind-to-world" direction of fit. Their job is to bring the mind in accordance with the world. Beliefs and knowledge are the paradigm cases, but perception falls into this category as well. Second, there are conative or "pro" attitudes (Davidson, 1963) with "world-to-mind" direction of fit whose job is to bring the world into line with the content of the attitude, e.g., desires, wishes, hopes, and intentions. Now, while beliefs and desires are the paradigm cases on both sides, there are specific attitudes on each side on the "periphery" towards the world that, according to Searle, are the biologically and ontogenetically primary ones: perception (on the mind-to-world side) and intentions (on the world-to-mind side).

What came into focus was the following possibility: Just as the first intentional attitudes that develop in human infants and other animals concern perception (of an objective world) and intentional action, the first form of second-order intentionality to develop should be an understanding of perception and action. And with this arose the further possibility that the divide between humans and other animals might go, ontogenetically speaking, even deeper. It might be that while the development of simple individual intentionality in humans and nonhuman primates runs in parallel, humans surpass all other animals in developing even simple forms of second-order intentionality.

The development of simple forms of second-order intentionality in human infancy has been described as the "9-month-revolution" in analogy to the social-cognitive revolution at 4 years (Tomasello, 1995b, 1999). This cognitive revolution in social understanding manifests itself in several distinct—but cognitively related—behaviours that first emerge around this time: After engaging in "dyadic" behaviours with either persons or objects for some time during the first year, children for the first time begin to engage in "triadic" behaviours that involve a referential triangle between child, another person, and an object/event that is jointly perceived/attended to or acted upon. Infants at this age begin to flexibly and reliably look where adults are looking (gaze following), understand what others do and don't see

(perceptual perspective taking), use adults as social reference points to disambiguate novel events (social referencing), and act on objects in the way they have seen adults act on them (imitative learning)—revealing an understanding of the adults' attitude/directedness towards the outside events (reviewed in Carpendale & Lewis, 2006; Carpenter, Nagell, & Tomasello, 1998; Tomasello, Carpenter, Call, Behne, & Moll, 2005). At the same age, infants also begin to use communicative gestures like pointing to direct adult attention and behaviour to entities and make protocomments on them (Liszkowski, Carpenter, Henning, Striano, & Tomasello, 2004), trying to influence the adult's attitude/directedness towards the events (Tomasello, 1995b). The fact that all these skills emerge in developmental synchrony and a correlated fashion (Carpenter et al., 1998) suggests a common underlying cognitive basis—an emerging understanding of oneself and others as intentional agents.

The importance of these cognitive abilities for developing cultural and linguistic forms of life is obvious. Learning a language presupposes a rudimentary notion of other speakers' perceptual perspectives and a rudimentary understanding of what they are up to. And cultural development essentially involves imitation as a form of reliably reproducing what one takes others to do intentionally.

Even great apes do not naturally and spontaneously point out things to one another (Povinelli & O'Neill, 2000) nor do they imitate each other (Tomasello, 1996). And most studies during the 1980s and 1990s failed to find convincing evidence for perceptual perspective taking (Povinelli & Eddy, 1996) or any other kind of social-cognitive abilities present in infants after the 9-month-revolution (reviewed in Povinelli & O'Neill, 2000; Tomasello & Call, 1997). Thus, one proposal became widely accepted in the field: The roots of uniquely human life forms lie in even the simplest forms of second-order intentionality that human infants develop early in ontogeny, but which no other species develops (e.g., Povinelli & Prince, 1998; Tomasello, 1999). These forms of social understanding enable humans to acquire language and to enter into culture, and their absence in other species explains why they do not develop cultural forms of life.

This simple picture, however, has been called into question by recent research, which suggests that chimpanzees might also be capable of at least simple second-order intentionality. First, a series of studies by Hare et al. (2000) and Hare et al. (2001) found that chimps understand something about others' perception. In a food competition situation, a subordinate and a dominant chimpanzee were placed in separate rooms on opposite sides of a third room. In the crucial conditions, food was placed in the third room so that the subordinate could see two pieces of food hidden while the dominant only saw one (his line of sight to the second being blocked by a barrier). The basic finding was that the subordinates did indeed take into account what the dominants could and could not see. Knowing that the dominants would take all the food they could see, the subordinates went for the food, visible

only to them, more often than they went for the food that both they and the dominant could see. Several control procedures and conditions (one using a transparent barrier that the subordinate apparently understood did not block the dominant's visual access to the food) effectively ruled out simpler explanations in terms of mere behaviour reading.

Second, a study by Call, Hare, Carpenter, and Tomasello (2004) suggests chimpanzees understand something about intentional action. Chimpanzees were presented with a human who had food in his hands and then behaved in different ways, marked as either unwilling or unable to give them the food. There were three conditions in which the experimenter was "unwilling" in different ways, e.g., just staring at the ape, eating the food, teasing the ape with the food. These conditions were each paired with two "unable conditions", e.g., trying to get the food out of a jar, dropping it accidentally, etc. In each group of matched conditions the surface topography of the experimenter's behaviour (body movements and gaze direction) were kept as similar as possible. The main finding was that chimpanzees were more impatient—banged on the cage more, left the area sooner—when the human was "unwilling" than when the human was "unable" (trying but failing), even though they got the food in neither case. The chimps in this study behaved as human infants did in a comparison study from 9 months of age (Behne, Carpenter, Call, & Tomasello, 2005).

The upshot of these two studies is that the rudiments of second-order intentionality that develop in human ontogeny at around 1 year are probably not so uniquely human after all; our common cognitive primate heritage runs deeper than previously thought.

COLLECTIVE INTENTIONALITY

Against this background, a broader alternative picture emerges. Basic forms of individual intentionality, both first and second order, develop in human ontogeny in the first one and a half years, but contrary to previous assumptions they develop to some degree in other primate species as well, and so cannot be the sole foundation for uniquely human forms of life. What is uniquely human, however, and a likely foundation of specifically human forms of life is the ability, developing from the second year in human ontogeny, to enter into collective (or "We") intentionality (Tomasello & Rakoczy, 2003; Tomasello et al., 2005).

In collective intentionality two or more subjects share an intentional "we" attitude which is not straightforwardly reducible to individual intentional attitudes.[2] When you and I meet and agree to take a walk together, to use an example from Margaret Gilbert (1990), we form and then pursue the joint We-intention "*We* walk together", which is not reducible to the sum of my individual intention "I walk" plus your analogous one. When I pursue my individual intention to walk and you pursue yours, we might end up walking

beside each other, but not together. When we pursue our We-intention, in contrast, each individual does walk, of course, but acts as part of a joint action (Searle, 2005).

As in the case of individual intentionality, different kinds of collective intentional attitudes can be distinguished: collective beliefs, collective desires, etc. The central cases of collective intentionality for present purposes are arguably the basic ones, namely collective intentions and actions—which constitute the class of cooperative acts and lie at the heart of societal and institutional life.

Clearly, collective intentionality presupposes first-order individual intentionality (for us to walk together I have to be able to walk intentionally myself), and second-order intentionality (for us to walk together I have to have track your walking intentionally and how I can adapt to it). But clearly individual intentionality, while being necessary, is not sufficient for collective intentionality. And so the present proposal is that while humans share with other animals simple forms of individual intentionality of first- and second-order, only humans have the ability to build on these to enter into collective we-intentionality.[3]

The relation between individual and collective intentionality is dialectical. On the one hand, human infants are cognitively equipped to understand each other as agents, as potential cooperators, which allows them to enter into collective intentionality and culture. But on the other hand, once children enter into collective intentionality and culture, acquire conventional practices, and above all a language,[4] this in turn shapes and transforms their individual cognitive development by supplying them with new means for thinking, much as Vygotsky and Mead have stressed (Tomasello & Rakoczy, 2003).

Before we turn to the empirical phenomena, some further taxonomic distinctions within the class of collective intentional affairs are relevant (see Figure 5.1, after Searle, 1995). Walking together is an example of a cooperative activity that does not essentially involve the conventional use of objects or any assignment of functions. Though such cooperative activities constitute the most basic form of collective intentionality, their cognitive structure is already quite complex: The individual participants have to understand each other as intentional actors, have to form and pursue a joint intention, and in the course of the joint act they have to be mutually responsive to each others' intentions and acts, often involving division of labour and complementary roles (Bratman, 1992). Crucially, even simple joint activities involve a normative dimension of commitment (Bratman, 1992; Gilbert, 1990).

An important subclass of collective intentionality involves the conventional use of objects and the collective ascription of functions to these objects, e.g., using tools to build something together or using pieces of wood to play chess together.[5] Two kinds of function can be distinguished here with two corresponding degrees of conventionality. *Causal usage functions* are functions we ascribe to objects when we collectively use them instrumentally, i.e., as tools,

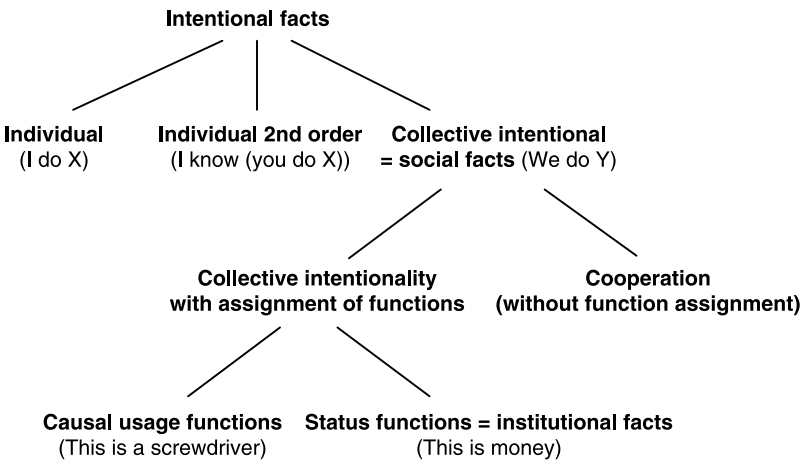

Figure 5.1 Taxonomy of individual and collective intentionality (after Searle, 1995).

and when we design and create objects as tools. The objects fulfil the function partly due to their physical causal make-up, e.g., a knife due to its sharpness. Such causal usage functions are conventional in a weak sense; nothing per se makes a certain object a tool, but we can assign a function to the object simply by making use of its intrinsic physical make-up for our instrumental purposes.

Status functions, in contrast, are conventional in a stronger sense. They are assigned to objects solely as a matter of collective practice; the objects do not fulfil the function due to their intrinsic properties. A slip of paper, for example, is money and a piece of wood is a queen in chess, but one could have decided to pay with wood and play with paper. An object has a certain status function only in virtue of the collective intentional treatment of it as having this status function. Status functions are brought into existence, are constituted by collective intentionality. "X counts as a Y in context C" is the formula that expresses status function creation.

Collective intentionality along with the creation of status functions is what lies at the heart of institutional reality. Status functions create institutional facts, like "This is a queen" or "this is money"; these are observer-dependent facts that only exist in the eyes of the beholders who collectively create them—in contrast to brute facts "out there" like "This is a piece of wood". Institutional reality as a system of status functions pervades our normal adult social life—we live as much in an institutional as in a natural world: We go to work, earn money to pay the rent, own property, are citizens, husbands, or wives, and we utter sounds with semantic status functions (meaning) all day, i.e., speak a language.

Specific normative dimensions are involved in the different forms of collective "We" intentionality. In cooperation, we commit ourselves to pursuing

a joint action and are therefore responsible for trying our best in this pursuit. The assignment of causal usage functions introduces the notions of good functioning and malfunctioning and the notions of appropriate and inappropriate uses of tools. Status functions, finally, involve a specific kind of rule, namely constitutive rules. Whereas regulative rules regulate an already existing activity (e.g., rules regarding on which side to drive regulate driving, which already exists before the rule), constitutive rules bring into existence the very activity they apply to (Rawls, 1955; Searle, 1969, 1995). Formally, "X counts as a Y in context C" specifies a constitutive rule: that X is a Y in the relevant context; and that it is a Y in the relevant context confers normative powers on the objects and carries normative implications (that it ought to be treated as a Y). A piece of wood is a queen in the context of chess and that means it has the power to move in certain ways, ought to be used accordingly, and ought not to be used as firewood in this context, for example.

In sum, collective intentionality involves two or more subjects who share an irreducible "we" attitude, paradigmatically a "we"-intention. Some forms of collective intentionality involve the collective assignment of functions to objects. The strongest form of such functions, status functions, are those that get collectively assigned to objects merely by virtue of convention, when objects are collectively treated as having that function ("counting as something"). Constitutive rules underlie status functions, create institutional reality, and bring with them normative implications—that the objects be treated according to the rules in the relevant context. With this taxonomy in hand, let us now turn to the development of the different forms of collective intentionality in human ontogeny from the second year on, and from a comparative point of view to the question how this development contrasts with that of other species.

Cooperation

Natural observations of human children suggest that during the second year they begin to engage in collaborative games with complementary roles and turn-taking structure and in collaborative instrumental activities with clearly differentiated roles (Brownell & Carriger, 1990; Eckerman & Didow, 1996; Ratner & Bruner, 1978; Ross & Lollis, 1987). In interpreting such observations, however, a fundamental conceptual challenge emerges: how to distinguish instances of cooperation proper from merely socially coordinated behaviour. This challenge is also apparent in interpreting natural observations of complex social behaviours in other animals. Prima facie, social hunting in chimpanzees looks very much like cooperation proper with division of labour (e.g., Boesch & Boesch, 1989). However, on closer inspection many commentators argue that what looks like real cooperation turns out merely to be complicated social coordination in which each participant plans and acts individually, while the impression of collectivity emerges from the contingency between the individual's behaviours due to external constraints (one individual starts

hunting at a certain place, then the next individual starts hunting, but cannot take the same place, then the third individual has to take even another place, etc.; Tomasello & Call, 1997; Tomasello et al., 2005).

To illustrate the challenge, consider Gilbert's (1990) walking example. Imagine two people taking a walk together. And now imagine two people that happen to walk beside each other at the same pace on a pavement, e.g., heading towards the same bus station. Both cases might look virtually the same (in both cases the two participants walk at the same pace, coordinate in the sense of looking where the other is currently walking, and wait when only one can get by, etc.). What we would usually do to find out whether this is a joint walk or just a coincidental parallel walk is to ask the participants what they are doing and consider their answers, e.g., "We're taking a walk" or "I'm walking to the station". So, the challenge is to distinguish the two cases in prelinguistic children or nonlinguistic animals.

Fortunately, the two cases may look very much alike, but not identical. In particular, different courses of action are to be expected in the case of obstacles to the pursuit of the joint goal.[6] When partner A cannot keep up with the pace of B, for example, B is committed to adapting her pace, and will usually do so (Gilbert, 1990). Or if she doesn't, then A will usually not only be surprised but offended ("But didn't we want to walk together?"). Similarly, if B stops walking for no obvious reason, A will usually reengage her ("Hey, come on!"). These responses are not to be expected in the case of mere coordination (the just coincidental parallel walking), because they are indicative of the specific commitment inherent in joint actions.

Recent experimental studies have employed these specific differences to test for young children's and chimpanzee's ability to engage in cooperation proper. Warneken, Chen, and Tomasello (2006) directly compared children (18 and 24 months old) and human-raised chimpanzees on experimental tasks that required collaboration with an experimenter. Some of these tasks had an instrumental problem-solving structure (toys for the children/food for the chimps could be obtained through an apparatus that had to be operated together), others were ends-in-themselves, collaborative games. During the collaborative acts, the experimenter was programmed to interrupt the shared activity at a predetermined point (by ceasing to perform his role) to see if the subjects would attempt to reengage his participation.

Children engaged proficiently in both kinds of collaborative tasks and during the interruption periods they communicated appropriately with the experimenter, e.g., pointed to his side of the apparatus, and tried to reengage him in the collaborative activity. Chimpanzees, in contrast, engaged with the experimenter in a coordinated way only in the instrumental tasks where a food reward was to be gained, but not in the noninstrumental game acts. Furthermore, even in the instrumental tasks, when the partner interrupted his engagement, they never communicated with him or tried to reengage him in the collaborative activity, rather they tried on their own or disengaged.

Another crucial difference between cooperation proper and mere social

coordination is that the former often involves division of labour, the assignment of complementary roles to different participants. With regard to the grasp of such role structures of joint activities, Carpenter, Tomasello, and Striano (2005) tested 18-month-old children in the context of novel collaborative games with an experimenter. First the experimenter played one role and the child learned to play the complementary one. In a second phase the children then spontaneously performed role-reversal imitation, i.e., acted out the other role when appropriate and thus indicated an awareness of the complementary role structure of the game. Three human-raised chimpanzees, in contrast, though they had acquired some imitative abilities through enculturation, revealed little competence for such role-reversal imitation (Tomasello & Carpenter, 2005).

Taken together these studies suggest that during the second year of human ontogeny children develop a nascent ability to engage in cooperative activities as the basic form of collective intentionality. On a simple level, they form and pursue shared "we" intentions with others, with a rudimentary awareness of the commitments and role structures characteristic of cooperative enterprises. The behaviour of chimpanzees, in contrast, does not necessarily warrant the ascription of collective intentionality proper, but might plausibly be characterized as complex social coordination only.

Collective function assignment and the foundations of institutional reality

Objects such as tables, nut-crackers, carburettors, or pens have causal usage functions and they have them essentially in virtue of the collective assignment of these functions to them. A stone might be very useful for cracking nuts, might afford nut cracking. But a regular stone is not "for" cracking nuts, and is not "broken" when it fails to crack nuts—in contrast to established nut-crackers (which might be stones).

Infants in their first year and many other animals do use tools individually. But does this imply they have any grasp of the collective assignment of usage functions to objects? Probably not. What has to be distinguished is using and understanding tools in a wide instrumental sense ("this stone affords nut-cracking") and understanding function assignment in a more narrow, normative-teleological sense ("this is a nut-cracker, that's how we use it and that's what it's *for*. If it doesn't crack, it's broken").

These two forms of understanding can easily be distinguished in older children with verbal interview methodology. And preschool children reveal rich knowledge about teleological ascription of functions to artifacts when talking about them (e.g., German & Johnson, 2002; Kelemen, 1999). But how do we know about younger children? One source of evidence comes from naturalistic observations: Children in their second year begin to use teleological normative vocabulary in describing malfunctioning artifacts as "broken" (for example, Kagan, 1981). Another source of evidence comes from imitation combined

with so-called "functional fixedness" (Duncker, 1945), the phenomenon that we find it hard to use artifacts in ways that deviate from their normal use. Imitation of conventional acts with artifacts alone from around age 1 might already be viewed as an index of understanding artifact functions. But imitation alone clearly cannot distinguish such a rich understanding from a simpler, affordance-based one. More plausibly, functional fixedness, however, can be interpreted as an index of understanding the conventional usage of artifacts, of the ways "we use them". Recent experimental studies have revealed that children from 2 years show just this kind of fixedness after imitating novel acts with artifacts (Casler & Kelemen, 2005). They are reluctant to use an artifact for purposes other than the one it was introduced for or to use a different artifact for a given purpose in place of the one introduced for it. Nonhuman primates, in contrast, seem not to be good imitators at all, and we have no evidence for imitative learning involving functional fixedness in any other species. Thus, while many animal species use tools in complex ways, only human children seem to develop an appreciation of the creation of artifacts as objects with collectively assigned usage functions.

Let us now turn to collective intentionality with the assignment of status functions. This form of collective intentionality, as we saw above, lies at the heart of institutional reality without which human society would be virtually inconceivable. And it is here that the dividing line between human sociality and that of other species can be seen most clearly (Searle, 2005).

Money and political leadership are obvious examples of status functions, but from an ontogenetic point of view, it is equally obvious that young children early in development do not have much of a grasp of such phenomena. I would like to suggest that playing games serves as a potential cradle for children's entry into collective intentionality with status function creation (see Rakoczy, 2006, 2007; Rakoczy & Tomasello, 2007). In fact, adult rule games such as chess are among the paradigmatic examples for practices involving status functions, e.g., "This piece of wood counts as a king in the context of chess". Of course, 2-year-olds don't play chess, but from the second year children do begin to engage with others in playing simple rule games, and in particular, games of pretence.

Take the example of some great pretenders, the toddler twins Adam and Eve. Eve takes a tennis ball, puts it to her nose and says "Hm, how delicious this apple is". She then offers it to Adam ("Here, have some"), who starts to make peeling movements on the ball with a toy knife. He then puts the ball to his mouth and makes enthusiastic chewing movements and "Yum" sounds. Eve joins in and finally they close the episode by saying "All gone, eaten up". Though this is not an instance of playing an established game with fixed rules, it is an instance of collectively playing a game with the assignment of transient status functions, making up ad hoc constitutive rules on the spot. "This ball *counts as* an 'apple' in our pretence context" is the central status function assignment and others follow—"It counts as peeled now" and then "It counts as eaten up now".

These assignments bring with them a normative structure of the joint activity. Some pretence acts are inferentially licensed in the game, others are not. Pretending to peel the ball/apple, pretending to eat it or to bake a cake with it are licensed, pretending to drive it or pretending to fax it are not (Walton, 1990).

From a comparative point of view, pretend play is quite clearly a uniquely human phenomenon. Although there are anecdotes of pretence-like behaviour in some human-raised animals (reviewed in Mitchell, 2002), these are difficult to interpret and generally it is quite clear that no other species reliably engages in pretend play as we know it (for reviews of precursors to pretend play in great apes, see Gómez & Martin-Andrade, 2002, 2005). Ontogenetically, children usually start to engage in simple pretend play in their second year. Early pretence is probably acquired by imitation (El'Konin, 1966; Rakoczy, Tomasello, & Striano, 2005), usually first occurs with caregivers (Haight & Miller, 1992; Slade, 1987), and centres around replica objects (e.g., toy cups and plates) and simple corresponding pretence actions (e.g., pretending to eat and drink) (e.g., Elder & Pederson, 1978).

But when they begin to engage in joint pretending, what do young children understand about the logical structure of pretence? In particular, to what extent do they grasp the normative and inferential structure created in shared pretence scenarios? In an elegant set of studies, Harris and Kavanaugh (1993) showed that children from 2 years do seem to grasp the normative structure created through joint pretence stipulations and engage in inferentially appropriate pretence acts. When the experimenter pretended to pour tea into a cup, for example, children pretended to drink from the cup. When the experimenter pretended to spill tea on the table, children pretended to clean the table.

Based on these studies, we recently tested young children in even more stringently designed situations (Rakoczy & Tomasello, 2006; Rakoczy, Tomasello, & Striano, 2004). Children (old 1-year-olds, young 2-, and young 3-year-olds) saw pairs of superficially analogous incomplete as-if-behaviours with objects: pretending to do an action and unsuccessfully trying to do the same action, e.g., pouring from a container into a cup. In both cases the actor would make pouring movements with a novel container over a cup, but without actual pouring happening. In one case, he would mark it with signs of playfulness and sound effects as if pretending to pour, in the other he would mark it with signs of surprise and frustration as if really trying to pour. Importantly, the container really did contain water and thus could really be used to pour. The logic is the following: In both cases the child sees superficially analogous movements, but they constitute radically different intentional actions. In the trying case the model wants to properly perform the action but fails. If the child understands the intentional structure of the model's act, this licenses the inference "If I want to do the same, other means should be used". In the pretence case, in contrast, the model performs an intentional pretence act involving the assignment of a status function ("This act counts as pouring"

and "The cup counts as full now"). If the child understands this as such, it licenses the child's entering into a joint we-pretence organized around this status function ("We pretend that the cup is full now") with appropriate inferential pretence acts, e.g., pretend to drink from the cup that the model had pretended to pour into.

2- and 3-year-olds (and to some lesser degree old 1-year-olds) showed this systematic and differential response pattern: After "trying" models, they themselves really performed the action or tried to, but by different means, e.g., using a tool to open the container first and then pouring. After "pretence" models, in contrast, they performed appropriate inferential pretence acts, e.g., pretended to drink from the cup into which the model had pretended to pour. This systematic pattern of responses is best interpreted, I suggest, as follows: In both cases children discern the intentional structure of the model's behaviour and respond accordingly. In the trying case, children perform an appropriate individual instrumental act. In the pretence case, however, they join into a collective we-pretence centred round the status function introduced by the model's overture.

That is, young children respect the inferential structure that comes along with collective intentionality and status function assignment, as indicated in their own actions. But what do they understand about the normativity that status functions introduce? Are they really following a rule or are they just acting in accordance with a rule? Do they indicate an awareness of the normative structure more directly and explicitly than in their own acts? Would they not only act correctly themselves, but criticize others for incorrect acts? This is crucial as critique, beyond mere surprise, in response to incorrect acts is the hallmark of appreciating normative structure; mere surprise is the appropriate response when there are acts deviant from purely statistical regularities.

In a recent study I addressed this issue (Rakoczy, in press). 2- and 3-year-olds were engaged in games of pretence with status functions assignment to objects. For example, with a pile of clothespegs, one clothespeg was pretended to be a knife, all others were pretended to be carrots. The child and the adult pretended to peel the carrots with the knife and cook and eat the carrots. Then at some point a third character (a puppet) came, joined into the pretence ("Oh, may I join your game?"), and in the target condition performed pretence acts that were normatively inappropriate in the light of the status functions of the objects, e.g., she pretended to eat the clothespeg that was the knife. In the control condition, the puppet pretended appropriately. Three-year-olds (and to a lesser degree 2-year-olds) frequently protested explicitly against such violations of the constitutive rules of the pretence game, e.g., "No, that's not a carrot, that's our knife!" in the experimental condition, but were content in the control condition.

In sum, in joint games of make-believe young children from 2 actively and knowingly participate in collective intentionality with status function creation—as indicated both in their own competent inferential actions and in their normative responses to other's mistakes.

Similar patterns were also found regarding young children's playing of simple rule games. In a recent set of studies (Rakoczy, Warneken, & Tomasello, 2008) we applied a logic similar to that in the pretence study mentioned earlier. Children of 2 and 3 years of age were engaged in a joint game with a partner, when at some point a third character, a puppet, came, wanted to join the game, but violated the constitutive game rules. Again, in all three studies the children protested in response to the puppet's mistakes frequently (but hardly did so in a control condition where the very same behaviour was performed, but the context was changed such that it was now not a mistake). That is, the 2- and 3-year-olds understood the normative structure of the game rules, i.e., the constitutive rules, also appreciating the context relativity of this normative structure.

In embryonic and isolated form we thus have here the basic structure of institutional reality in the games of 2- and 3-year-olds. Of course this is a long way from money, marriage, and universities, but the seeds are there, and so joint pretending and playing other games quite plausibly can be considered the central cradle for and the entrance gate to institutional life.[7]

SUMMARY AND CONCLUSION

Humans live complex social lives and so do countless other species. Humans, though, are not only social beings, as are other animals, but participate in cooperative, societal, and institutional ways of life that are quite clearly unique on this planet. Traditionally, individual intentionality was supposed to be the monopoly of humans that explained why only humans talk and have culture. But today in light of the impressive parallels discovered in the early development of individual intentionality in humans and in other species, there is little plausibility to such a story. For a while second-order individual intentionality ("theory of mind") seemed a likely successor candidate for explaining uniquely human sociality. And while complex second-order intentionality, i.e., ascribing epistemic attitudes to each other etc., is surely a uniquely human achievement and an important foundation for many uniquely human ways of life (self-consciousness, full-fledged communication, etc.), recent research with apes suggests that at least simple forms of second-order intentionality aren't as uniquely human as once thought.

What thus emerges as the likely core of the uniquely human cognition which underlies uniquely human sociality is the ability to enter into irreducibly collective "we" intentionality. Different forms of this ability develop from the second year on: the ability to engage in shared "we" activities with others, opening up the space of cooperation, and the ability to participate in collective practices with the assignment of functions, in particular status functions, opening up the space of institutional reality.

The human animal is essentially a collective animal, and ontogenetically it becomes so from very early, from the second year of life. This presumably

is the point where humans begin to leave all other animals behind and grow into a community of sharing "we" intentions and practices, an essentially normatively structured public space that is constitutive of full-fledged personhood.

ACKNOWLEDGEMENT

This work was supported by a "Dilthey Fellowship" of the Volkswagen Foundation and the Fritz Thyssen Foundation.

NOTES

1 I am speaking of "simple" intentionality here, as arguably many forms of full-fledged human individual intentionality are essentially dependent on language, a point I will return to later (see Bermudez, 2003, for a proposed taxonomy of simple nonlinguistic intentionality in contrast to linguistically mediated intentionality).
2 For the central works in recent analytical philosophy on this, see Bratman (1992), Gilbert (1990), Searle (1990, 1995, 2005), Tuomela (1995), and Tuomela & Miller (1988). For an overview, see Tollefsen (2004).
3 In other words, the "Machiavellian Intelligence Hypothesis" (Byrne & Whiten, 1988) with its emphasis on the social cognition underlying individual manipulation of conspecifics (i.e., a form of individual second-order intentionality) might well be adequate for nonhuman primates, but falls short of adequately describing the inherently collective dimension of human social cognition.
4 I will here hardly touch upon the development of language and its relation to collective intentionality—as this would easily go beyond the scope of the present chapter. On the one hand, language as a conventional practice is itself an instance of collective intentionality and thus in some sense secondary to collective intentionality. On the other hand, language is in some sense the fundamental collective activity without which many other collective practices would not be conceivable.
5 Strictly speaking, functions are not only assigned to objects, but to actions as well (and, in fact, actions are logically the primary case—the status of objects is dependent on relevant actions one can do with the objects). Language is the paradigmatic example: Emitting such and such sounds in the right context according to the right rules counts as speaking. But I will here focus on the case of object functions, first because regarding objects the general forms of collective intentionality can best be illustrated. And second, because ontogenetically it is plausible that children come to understand function assignment to objects before they understand it in the case of actions.
6 In analogy to the case of individual intentional acts: Persistence in the presence of obstacles is usually a good criterion to distinguish goal-directed action proper from just seemingly goal-directed behaviour.
7 Of course, language is the first instance of collective intentionality involving status functions into which young children enter in rudimentary form from 1 year on. However, and this is one of the reasons why I haven't touched upon language in this context, arguably young language learners do not have to have any understanding whatsoever of the logical status of constitutive rules and the creation of status functions. Children up to the age of at least 4 or 5 just do not view language sounds as phonologically or syntactically defined events or objects (brute facts) that

additionally are assigned meaning (institutional facts). They hear through the sounds, directly perceiving them as meaningful (as we all normally do when we do not take any kind of metalinguistic stance). The situation is different in the case of games, however, because there status functions are assigned to physical objects that children surely see as such. This is especially clear in the case of pretence (e.g., "this tennis ball counts as the apple in our pretence") where children have to at least implicitly distinguish the brute fact about the object (it's a ball) from its status function ("apple") in the game.