
Intentions and Intentionality Foundations of Social Cognition

© 2001 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

Set in Sabon by The MIT Press.

Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Intentions and intentionality: foundations of social cognition / edited by Bertram F. Malle, Louis J. Moses, and Dare A. Baldwin

p. cm.

"A Bradford book."

Includes bibliographical references and index.

ISBN 0-262-13386-5 (hc: alk. paper)

I. Intentionalism. I. Malle, Bertram F. II. Moses, Louis J. III. Baldwin, Dare A.

BF 619.5.I58 2001

153.8—dc21

00-064590

edited by Bertram F. Malle, Louis J. Moses, and
Dare A. Baldwin

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

young as 12 months display appreciations of object directedness and action connectedness reveals and helps shape their developing intentional understandings.

Acknowledgment

Preparation of this chapter was supported by NICHD grant HD 34004.

Notes

1. For more on this, see Baldwin and Baird 1999 and Baird and Baldwin, this volume.
2. In addition, findings for the needed control condition were not straightforward. In the control condition of Gergely et al. (1995), infants were again habituated to a circle taking the same indirect path as in the experimental condition described in the text, but in this case even in habituation there was *no* barrier to circumvent. The control condition test events compared direct and indirect paths with *no* barriers just as in the experimental condition. In the control condition, 12-month-olds dishabituated to *both* the direct and indirect test events. Dishabituation to the indirect test event in both the experimental and control conditions raises the suspicion that the indirect test event was somehow just generally attention-eliciting.

7

How Infants Make Sense of Intentional Action

Amanda L. Woodward, Jessica A. Sommerville, and José J. Guajardo

A glance out the window reveals a scene populated by two kinds of entities: inanimate objects (such as cars, trees, and fire hydrants) and animate beings (in particular, people). Adults discriminate these two kinds of entities readily and have very different expectations about how members of each class will behave. Adults understand not only the physical regularities that govern the motions of inanimate objects but also the psychological underpinnings of human action. This "folk psychology" allows us to make sense of human behavior in terms of the goals and plans that drive it, the beliefs that inform it, and the emotions that color it, among other things. At the core of this folk theory is the idea that human action, unlike object motion, is driven by intentions. Adult understanding of intentions is embedded in rich knowledge about mental states and behavior. This enables adults to detect intentional actions on the basis of behavioral evidence and to reason about the particular intention behind an action.

One of the most enduring questions in developmental psychology is how children come to understand the distinction between inanimate objects and animate beings. A key part of this distinction is how children come to understand intentional action. How do children first detect intentional actions, and how does their understanding of intentions develop? Folk physics has roots in infancy (Baillargeon 1995; Spelke, Breinlinger, Macomber, and Jacobson 1992). In our work, we ask whether folk psychology also has roots this early in life. That is, do infants have ways of making sense of intentional action that are continuous with adult understandings?

A number of theorists have considered this question in recent years. There have been two major, and conflicting, views of the development of intentional understanding. One set of theorists has focused on the social abilities

that appear during the toddler years. At about 12 months, babies make impressive strides as interactors. They begin to produce and respond to communicative gestures such as points, to actively follow other people's gaze, to engage in social referencing, to play games such as peek-a-boo, to imitate the goal-directed behavior of others, and to understand words. (For a review, see Tomasello 1995.) Several theorists have proposed that the onset of these abilities signals an understanding of other people as intentional agents (Bretherton 1991; Carpenter, Nagell, and Tomasello 1998; Tomasello 1995). Because these behaviors are absent in infants younger than 9–12 months, it is sometimes further concluded that "there is no joint attention or any other indication that infants at this age understand others as intentional agents" (Tomasello 1995, p. 108). Thus, under these accounts, the birth of intentional understanding occurs at around 12 months. Other theorists have interpreted infants' naturally occurring social behaviors as evidence that the birth of intentional understanding occurs still later, at 18–24 months (Barresi and Moore 1996).

At the other extreme, several theorists have proposed that infants are innately endowed with abstract and elaborate systems for interpreting intentional action, including notions of goal-directedness, perceptual contact, affinity, reciprocity, the ability to learn, enduring preferences, and rationality (Baron-Cohen 1995; Gergely, Nadasdy, Csibra, and Biro 1995; Premack 1990). These notions would be activated by the presence of perceptual cues, such as self-propelled or biological motion. To illustrate, Premack (1990) suggests that when infants see that an object is self-propelled they then infer that it moves intentionally. Premack also hypothesizes that, as a part of understanding the intentions of the object, infants may also infer that it will prefer its own kind and seek to reciprocate the actions of other objects. While proposed to be different from full-fledged adult systems of reasoning, these innate systems would provide a critical substrate for further development, on analogy with proposals for innate core knowledge in other domains, such as physics (see, e.g., Spelke et al. 1992; Spelke and Newport 1998). The existence of early abstract expectations about intentional action has not generally been evaluated by direct empirical tests (but see Gergely et al. 1995). Instead, theorists have argued that the concept of intentional action is so complex and so important for survival that strong innate constraints are required to explain its ontogeny.

In our work, we question both of these accounts. First, we take seriously the possibility that infants, before they acquire the communicative tool box of the 12–24-month-old, understand some aspects of intentional action. Second, we take seriously the possibility that intentional understanding develops a piece at a time based on experience with particular actions and actors, rather than being innately specified in some abstract form.

Action as Goal-Directed

Mature conceptions of intentional action are multi-faceted. Adult folk psychology explains behavior in terms of an actor's goals, perceptions, emotions, beliefs, knowledge, preferences, and personality traits, among other factors (Heider 1958; Wellman 1990). In investigating the potential infant precursors to this mature system of knowledge, therefore, the first decision we faced was where to start. We began by exploring a foundational component of folk psychology: the assumption that human action is goal-directed. Both adults and preschoolers understand human behavior not as an undifferentiated series of motions through space but rather as actions directed toward goals (Heider 1958; Lillard and Flavell 1990). This insight provides one important basis for understanding intentional action.

Recent research indicates that an understanding of goal-directed action is present very early in childhood, as young as 14 months. Fourteen-month-olds are more likely to imitate actions that appear to be purposeful than behaviors that seem to be accidental (Carpenter, Akhtar, and Tomasello 1998), and older toddlers show similar attention to goal-directed actions over other kinds of behaviors in word learning (Tomasello and Barton 1994). When 18-month-olds see a person slip and fail to complete an intended action, they imitate the intended action and not the actual movements that the actor made (Meltzoff 1995). Thus, toddlers selectively attend to and remember the elements of an event that are relevant to the actor's goals.

Our first question, then, was whether infants share this propensity to pay attention to the goals of an actor. We began with the familiar goal-directed action depicted in figure 1. Infants saw a person reach into a curtained stage, move her arm through a distinctive path, and grasp one of two toys that were mounted on the stage. There were at least two aspects of this event




Figure 1
Sample habituation and test events for the human grasp (based on Woodward 1998).

that infants could attend to and remember: the salient path taken by the actor's arm and the relation between the actor and the object that was her goal. Adults would likely describe the event in terms of the latter ("She grasped the bear") rather than the former ("She moved her arm through 10 inches of space from the bottom right corner of the stage to the far left side"). Adults can notice either aspect of the event, but because reaching is understood as goal-directed, the goal-related features seem more central than others.

To assess infants' construal of the event, we used the visual-habituation paradigm. A well-established finding is that once infants habituate to one stimulus they will look longer at another stimulus that seems new to them. We drew on this response in tapping infants' representations of the grasp

event. To determine which aspect of the event infants weighted most heavily in their representations, we measured the strength of the novelty response to a change in each aspect. After the infant habituated to one event, we reversed the positions of the toys and presented test events in which there was a change in either the relation between the actor and the goal (new-goal trials) or the path of reach (new-path trials). Figure 1 provides an example of these events.¹ Six-month-olds and 9-month-olds showed a stronger novelty response (i.e., looked longer) on new-goal trials than on new-path trials (Woodward 1998). That is, like toddlers, young infants selectively attended to and remembered the features of the event that were relevant to the actor's goal.

As is the case for toddlers (Meltzoff 1995; Meltzoff and Brooks, this volume), infants' propensity to attend to goals seems to be specific to human actors. Infants did not selectively attend to the relation between actor and object for events involving a range of inanimate "actors" (Woodward 1998). For example, when 6-month-olds saw a mechanical claw grasp the toy (figure 2), they showed somewhat greater recovery on new-path trials than on new-goal trials. The claw events were similar on several dimensions to the hand events: there was motion through space ending in contact with the toy, the claw was covered in cloth that matched the actor's sleeve and hand, and the claw grabbed hold of the toy. Nevertheless, infants construed the hand events and the claw events differently.

Subsequent analyses revealed that this pattern of findings was not a by-product of infants' interest in hands as compared to their interest in claws and other inanimate actors. Infants across studies and conditions had their attention drawn to the toy that was contacted, whether it was contacted by a hand or by an inanimate object. However, infants who saw a hand differed from those who saw an inanimate actor in the features they weighted most heavily in their representation of the event.

This set of findings is the first to indicate that infants under a year of age understand certain human actions as goal-directed, in that they construe these actions primarily in terms of the relation between the actor and the goal. (For further evidence, see Wellman and Phillips, this volume.) By themselves, these findings leave open the question of exactly how infants understand this relation. Infants' understanding of action as goal-directed is likely quite different from adults'. Adults understand goals as mentally represented




Figure 2
Sample habituation and test events for the inanimate grasp (based on Woodward 1998).

entities that are embedded in the rest of a person's mental life. A person's goals relate to that person's beliefs about the world, beliefs about his or her abilities, desires, and preferences. In view of the literature on theory of mind in preschoolers, it is very unlikely that infants have a full-fledged understanding of this sort. However, infants may understand agent-object relations in other ways that are continuous with adult understandings.

One possibility is that, in addition to understanding that certain actions imply a relation between a person and an object, infants understand these actions as being "directed at" the object. This possibility, considered by Wellman and Phillips in the present volume, is also consistent with our data. There are at least two ways in which infants could understand action as directed toward objects. For one, infants might understand this aspect of

action at a purely behavioral level. Csibra and Gergely (1998) have developed an account of such a system of knowledge that could serve as a precursor to later mentalistic understandings of goals. In addition, infants may understand something about the internal aspects of the relation between an actor and the actor's goal—for example, that the phenomenological experience of wanting something accompanies the actions that are deployed to get it. Infants could understand this aspect of goal-directed action without yet understanding very much about mental life. Further empirical work is required to explore these possibilities.

These issues aside, our findings indicate that infants are on the right track, in that they are attending to just those aspects of actions that are relevant to goals in the adult sense. These findings have implications for the theories we outlined earlier. They weigh against the claim that infants younger than 12–24 months lack any understanding of intentional action or of relations between agents and objects (Meltzoff 1995; Tomasello 1995). Six-month-olds construe one intentional action, grasping, in terms of the relation between the actor and the goal. In addition, 6-month-olds have begun to draw the line between animate and inanimate entities, interpreting motions of the former, but not the latter in terms of the relation between agent and object. Therefore, by the time they are 6 months old, infants can detect certain instances of intentional action, and they attend to the aspects of the action that are critical for understanding the specific intention behind it.

Our subsequent studies revealed important limitations on infants' early propensity to attend to human goals. These limitations are relevant to theories at the other extreme, that is, theories that posit infants to be endowed with rich and abstract notions of intentional action.

The Specificity of Infants' Understanding of Goal-Directed Action

As we mentioned earlier, some theorists (Baron-Cohen 1995; Gergely et al. 1995; Premack 1990) have proposed that the ability to understand action as intentional is rooted in rich, innately specified systems of reasoning, rather than being acquired through experience. Because they are not rooted in knowledge about particular actions, these systems of reasoning are argued to be very general, potentially applying to any event the infant sees.

To limit the cases in which infants interpret an event as intentional, theorists further propose that there is a perceptual "trigger" that activates these innate systems of knowledge. For example, Premack (1990) has proposed that infants are born with a system for interpreting action as intentional, and that it is triggered whenever they see self-propelled motion. (See also Baron-Cohen 1995.) Other triggers that have been proposed are the biological patterns of motion associated with animals and people (Baron-Cohen 1995) and apparently rational motion toward a goal (Gergely et al. 1995).

These features would generally serve to identify the actions of people. However, they will always be at best an approximate cue. Many apparently self-propelled motions are not intentional, even those produced by people. On Premack's theory, infants would over-attribute intentionality to inanimate objects that move with no apparent external force (such as drifting leaves) and to human behaviors that are not purposeful (such as sneezes or stumbles). In our later studies, we evaluated these predictions. Our findings indicate that, contrary to these proposals, infants' initial understandings of agents and actions are rooted in their knowledge about specific agents and specific actions.

How Do Infants Identify Agents?

Do babies attribute the potential for goal-directed action to anything that presents the right perceptual trigger? In all theories that posit a trigger, the trigger is argued to be a characteristic pattern of motion—most often, self-propelled motion.² Baron-Cohen (1995, p. 34) proposes that even though "the visual input might look as shapeless as an amoeba, as weird as a giraffe, or as minimal as a stick insect," nevertheless "because of their self-propelled motion, all these are instantly interpretable as agents with goals and desires."

There are two levels at which this proposal can be evaluated. First, is it really the case that properties of motion are the sole basis for infants' judgments of agency? Second, regardless of the particular features infants use, is the idea of a hard-wired trigger accurate? Our findings address each of these questions.

On the one hand, this emphasis on patterns of motion seems intuitively correct. After all, one critical feature of animate entities is that they can move on their own. On the other hand, it seems equally intuitive that other

features bear on whether an object is identified as animate. Imagine you are walking through the woods and see a long thin object lying motionless across the path. If the object has the texture of snakeskin, you will likely react quite differently than if it has the texture of tree bark. Similarly, imagine you enter a room to find a blue plastic disk moving through the air, apparently without an external source of energy and in an irregular "biological" manner. You would likely be much less surprised to discover that the disk was attached to a mobile by means of a thin wire than to discover that it was a kind of flying animal.

Our earlier findings indicate that by 6 months infants distinguish between people and inanimate objects such as mechanical claws, but they leave open the issue of the means by which infants make this distinction. With respect to self-propelled motion, the hand and claw events were equally ambiguous. From only the information in the display, it was not clear whether either the hand or the claw was self-propelled, because the infant could only see the end of the arm or claw. The display did not specify whether the hand and claw started up on their own or were made to move by another object. The claw clearly differed from the hand on several other dimensions, including patterns of motion (biological versus mechanical), parts (e.g., fingers), overall shape, and texture (skin versus plastic).

In an ongoing set of studies (Guajardo and Woodward, in progress), we are testing infants' sensitivity to texture. Will infants differentiate between objects that have identical motion properties but different textures? The setup and the procedure of the study are similar to those depicted in figure 1. However, the grasping action is performed by an actor wearing a metallic-gold-colored evening glove. The glove covers all traces of skin but preserves the characteristic patterns of motion, parts, and overall shape of the actor's hand. In the first study, the actor reached in through a curtain so that infants could only see her arm (clothed in a magenta sleeve) and her hand (in the gold glove). We tested infants at two ages: 7 months and 12 months. Control groups of 7-month-olds and 12-month-olds watched the events performed by an ungloved hand, just as in earlier studies. The control groups at both ages showed the same patterns seen in previous work: they responded strongly to a change in the relation between actor and goal, and less strongly to a change in the path taken by the actor's arm. In contrast, infants at both ages who saw the gloved hand did not show this pattern. The younger group

looked marginally longer on new-path test trials; the older group looked about equally as long on the two kinds of test trials. At both ages, infants' overall level of attention did not vary as a function of whether they saw the glove. That is, the difference between the findings in the two conditions does not seem to be due to infants' finding the glove either extremely interesting or extremely aversive. Infants were equally attentive to the gloved hand, but they did not seem to treat it as an agent. Contrary to the predictions of several theories, then, texture is an important cue in infants' determination that an entity is capable of goal-directed action. This finding concurs with work by Smith and Heise (1992), who found that texture is a powerful source of information for young children in categorization tasks involving animals and artifacts.

One possible conclusion from these findings is that the idea of a hard-wired trigger is correct, but the trigger is textural rather than motion based. Perhaps infants consider an entity to be an agent so long as it has skin. However, our later findings argue against this conclusion. For one thing, as we will discuss below, infants do not treat all motions of a naked human hand as goal-directed. In addition, we found we could alter babies' interpretation of the gloved hand by giving them more information about it. We reasoned that the original events gave infants very little evidence that the hand was a part of a person. Most of the actor (including her face) was hidden from view, and the actor never spoke or interacted with the infant. Thus, the actor in the display lacked many of the features that infants probably associate with people. This fact, coupled with the absence of skin texture, might have undermined infants' understanding that the hand was a part of a person. In the next study, we again showed 12-month-olds the gloved hand grasping the toy, but this time accompanied by other cues that the hand was a part of a person. We showed infants the actor from the waist up. At the start of each trial, she made eye contact with the baby, said "Hi," and then said "Look" as she turned to look at and grasp one of the toys. After habituation, as in previous studies, the toys' positions were reversed, and infants saw the actor reach for the same toy, now on the other side of the stage (new-path test events), or for the other toy (new-toy test events). The actor wore the same gold gloves and sweater as before. This time, however, 12-month-olds interpreted the movement of the golden hand as goal-directed, looking reliably longer on new-goal trials than on new-path trials.

These findings run counter to the idea that infants rely on the presence or absence of a single trigger feature in determining whether they are seeing an agent. The patterns of motion associated with humans were not sufficient to convince infants that the gloved hand was an agent. Texture is also important. However, texture is not the only feature that matters for infants. When given more information about the gloved hand (that is, when they could see it attached to a person), infants readily interpreted its motions as goal-directed. More generally, babies do not seem to focus on only one feature, but instead respond to several different kinds of cues to agency. These findings suggest that, rather than attributing agency to an overly broad class of items that share a single feature, infants begin by focusing on one specific kind of entity, the person, which has many typical features.

How Do Infants Identify Goal-Directed Actions?

Nativist theories predict overgenerality not only in the entities identified as agents but also in the actions identified as intentional. On Premack's (1990) account, any movement of a self-propelled agent would be considered as intentional, including "bumps between people, as well as falls or tumbles" (p. 13). On the account of Gergely et al. (1995), infants would consider as intentional any motion through an apparently rational path toward a potential goal. On these accounts, because infants would initially identify an overly broad set of behaviors as intentional, development would be a process of paring down the range of events considered as intentional.

We tested this possibility using the same paradigm as in earlier work with a normal, ungloved human hand, this time varying the type of contact between the hand and the object (Woodward 1999). In one condition, the actor reached toward and grasped one of two toys. In the other, the actor lowered her arm toward the toy, letting her hand fall the last few inches so that it landed, palm up, on top of it (figure 3). Aside from this difference, the two events were similar on several dimensions. In both events, the actor's arm moved through the same paths and the actor's hand hid roughly the same portion of the toy from view. The two events took the same amount of time and were equally effective at drawing infants' attention to the toy that was contacted by the hand. Critically, each event involved self-propelled, biological motion toward the object, and each involved a human actor.




Figure 3
Sample habituation events for different kinds of contact between actor and object (based on Woodward 1999).

As in previous studies, infants were habituated to one event and were then shown test events that altered either the path of motion taken by the actor's arm or the identity of the object that was contacted. We first tested 9-month-olds. Despite the similarities between the grasp and back-of-hand events, infants encoded the two events differently. For grasps, infants responded more strongly to a change in goal object than to a change in path, as in earlier studies. That is, infants represented this event primarily in terms of the relation between the actor and the goal. The infants who saw the actor touch the toy with the back of her hand did not do this; they looked equally at the two kinds of test events. This difference was not due to infants' finding the back-of-hand events much more interesting or much less interesting than the grasp events. The overall amount of looking in the two conditions was the same. In a follow-up study, 5-month-olds showed similar though somewhat weaker patterns. As early as we have seen evidence that infants construe grasping in terms of the relation between actor and object, therefore, we have also seen evidence that infants distinguish between grasping and other manual actions. Counter to the predictions of several nativist theories, infants do not construe all self-propelled motions or even all human motions as goal-directed.

These findings argue against the conclusion that infants initially identify an overly broad set of behaviors as goal-directed. They further suggest two possible means by which infants might identify particular actions as goal-

directed. One possibility is that infants understand some features of action (for example, smooth movements, articulated motion with respect to an object, and palm orientation) as evidence of goal-directedness. The grasp and back-of-hand events differed on each of these dimensions. By 18 months, babies use this kind of evidence to interpret novel behaviors (Meltzoff 1995; Tomasello and Barton 1994). A second possibility is that infants' ability to determine whether behaviors are goal-directed initially derives from their familiarity with particular actions. Grasping is ubiquitous in infants' environments and is also an action they have experienced from the agent's point of view. On the basis of these experiences, infants may have learned about the goal-directed nature of grasping. Since they are unlikely to have encountered the back-of-hand event before, they may not have a ready way to make sense of it.

Another finding from our lab is consistent with the second of these possibilities. In an ongoing series of studies, we are testing infants' understanding of pointing. Like grasping, pointing involves smooth, coordinated motion toward an object and (in the events for this study) contact with the object. In addition, like grasping, pointing indicates a relation between an actor and an object for adults. Adults understand a good deal about the specific nature of the relations between objects and people who grasp them or point at them. Grasps are understood as indicating a desire on the part of an actor, whereas points indicate an attentional state and perhaps the intention to communicate with someone about the object. Even if infants did not understand the exact nature of these actor-object relations, in order to understand the point gesture in even a rudimentary way they would have to know that there is some relation between the actor and the object of her point.

We used the same methodology as in previous work to study whether 9-month-olds and 12-month-olds represent pointing in terms of the relation between the person who points and the referent object (Woodward and Guajardo, in progress). We showed infants events similar to the grasp events in our first studies, except that now the actor pointed to and touched the toy instead of grasping it. Some infants saw only a pointing hand; others saw the actor's upper body and face too. In this case, the actor began each trial by making eye contact with the infant, saying "Hi," and then saying "Look" as she looked at and pointed to the object. We thought that these

behaviors might provide further evidence that the point was directed toward the object. However, we found that the results in the two conditions did not differ. As in previous studies, infants, having been habituated to a single event, then saw test events in which either the path taken by the actor's arm or the identity of the referent object had changed. Despite the presence of behavioral cues that the point was directed at the object, 9-month-olds did not seem to construe it in this way. They looked equally long at the new-referent and the new-path test events. In contrast, 12-month-olds showed a greater novelty response on new-referent trials than on new-path trials. These findings are noteworthy because it is between 9 and 12 months that babies begin to follow other people's points (Butterworth and Grover 1988; Schaffer 1984) and to use the point gesture to direct other people's attention (Bates et al. 1979).

There are several reasons why infants may have a harder time discovering the relation between a person who points and the referent of the point than discovering the relation between a person who grasps and the object of the grasp. As has already been noted, grasping and pointing are manifestations of different kinds of underlying intentions, and these may more easily be understood for grasping than for pointing. Moreover, in everyday life grasping is accompanied by concrete cues to the actor's intentions. Grasping often involves physical consequences for the object (e.g., it might be moved closer to the actor), which could help infants to understand the actor's goals (e.g., obtaining the object). In contrast, pointing may not have clear physical consequences for the referent object. In addition, outside the laboratory pointing most often occurs at a distance from the object. The demands posed by relating entities separated in space may make it difficult for infants to learn about points.

In summary, our findings indicate that the early development of intentional understanding is not a process of paring down initially overgeneral notions but instead a process of building up initially undergeneral ones. Infants begin by understanding particular actions as goal-directed, and with time the range of actions they understand in this way increases. Just as infants' notions of actors seem to focus on people in particular, rather than on the broad class of anything that moves on its own, so infants' notions of goal-directed action seem to focus on particular actions.

A Developmental Challenge: How Are Early, Specific Notions Enriched?

The above conclusions raise this question: How do infants move beyond these early, specific beginnings to a more general understanding of intentional action? In recent work (Woodward and Sommerville 2000), we have investigated one means by which infants' understanding of intentional action may be enriched.

Adults are not limited to understanding just a few actions, or even to understanding only actions with certain behavioral features, as intentional. Instead, adults can interpret action in context, and thus can understand completely novel actions as intentional. At the heart of this ability is the understanding that distinct actions can often be related to the same overarching goal (Schank and Abelson 1977; Searle 1983). For example, on observing someone grasp a refrigerator door handle, pull open the door, grasp a gallon jug of milk, and carry the jug to a waiting glass, adults readily interpret the sequence as "getting a drink of milk" and understand the separate actions as relating to this overarching goal. This ability to interpret actions in sequence provides adults with a way to infer the intention behind novel or ambiguous actions. Running a thumbnail around the top of the milk jug might not in itself have a clear goal, but in an informative behavioral context it can be understood as a means to opening the jug and thus as related to the goal of getting a drink of milk.

If infants could link actions to an overarching goal, therefore, this would provide them with one way to interpret novel actions. (For a discussion of this problem and another approach to it, see Baird and Baldwin, this volume.) To test this possibility, we created a simple analogue of the above situation for 12-month-olds. First, we devised an action that we hoped would be ambiguous to babies (figure 4). On a stage sat two clear boxes, each a different color of translucent plastic and each containing a different toy. Infants saw an actor reach into the stage area, rest her hand on top of one of the boxes, and hook her thumb under the lid. For reasons that will become clear, we called this the *single-action* condition. This action could be construed as being directed at the toy inside the box or at the box itself. In order to determine how infants construed it, we tested whether they showed a greater novelty response to a change in the relation between the




Figure 4
Sample habituation and test events for the single-action condition (based on Woodward and Sommerville 2000)

actor and the box or to a change in the relation between the actor and the toy. After habituation we switched the positions of the toys and then showed infants two kinds of test trials in alternation. On new-box trials, infants saw the actor perform the same action on the other box, which contained the toy associated with the reach during habituation. On new-toy trials, infants saw the actor perform the action on the same box as in habituation, which now contained the other toy. Twelve-month-olds looked equally long at the two test events, indicating that the goal of this action was ambiguous to them.

For a second group of infants, we embedded this ambiguous action in a sequence that culminated in an action infants understand as goal-directed—namely, grasping. We called this the *embedded-action* condition. After touching the box lid, the actor opened the box and grasped the toy within it. Infants saw this event during habituation. After habituation, we showed

infants the same two test events as in the single-action condition. The actor touched the top of the box but did not open it, and there was either a change in the relation between the actor and the box or a change in the relation between the actor and the toy. In this condition, infants looked longer on new-toy trials than on new-box trials. That is, infants now interpreted touching the lid as being directed at the toy within the box. In sum, infants used the second action (grasping the toy) to interpret the first (touching the lid).

From the results so far, it is not clear how infants related the two actions. One simple strategy would be to assume that any actions that occur in sequence are related to the same goal. Young infants have the wherewithal to employ this strategy, since they are adept at detecting temporal patterns (Saffran, Aslin, and Newport 1996). However, this strategy runs the risk of relating actions that co-occur serendipitously. Imagine someone reaching into a cookie jar and pausing in mid-reach to pick up a ringing phone. Adults would not think that the reach to the phone and the reach to the cookie were directed toward the same goal. If infants relied solely on co-occurrence, they would make this error. In relating actions to overarching goals, adults draw on their knowledge of the constraints in a situation, including physical constraints (e.g., whether an action is a possible means to obtaining the goal), psychological constraints (e.g., what the actor knows), and social constraints (e.g., conventionally appropriate behaviors). In our study, there was a physical relation between the two actions that could have led infants to relate them. In order to obtain the toy, the actor had to first touch (and remove) the box lid. The question was whether infants used this relation to link the actions.

To address this question, in the next study we showed infants events in which the temporal relation between opening the box and grasping the toy was identical to that in the embedded-action condition, but in which the causal relation between these two actions was disrupted. Rather than sitting inside the boxes, the toys sat outside and in front of them. During habituation, infants saw the same event sequence as in the previous study. The actor opened the lid and grasped the toy that now stood in front of the box. The test events were like those in the first study. The toys' positions were reversed, the boxes remained closed, and the actor either reached to a new box or to the box behind a new toy. If infants relate actions based only on temporal sequencing, we would expect the same findings as in the previous

study. What we found was almost the opposite. There was a marginal trend toward looking longer on the new-box trials than on new-toy trials. Thus, disrupting the causal connection between the two actions interfered with infants' propensity to relate them. Like adults, therefore, infants draw on their knowledge of the causal constraints in a situation in relating actions to overarching goals.

These findings indicate a mechanism that would allow infants to move beyond highly specific ideas about goal-directed action. The ability to use familiar actions to interpret ambiguous actions would enable infants to understand a much greater proportion of the actions that they witness in day-to-day life than they would if they understood only the goals behind certain isolated actions. Importantly, the ability to consider the causal constraints in a situation in relating actions would enable infants to determine which components of an event sequence are relevant to a particular goal. When considering which actions are related to a goal, infants can limit their search to actions that appear to be causally related to obtaining the goal (for example, opening a box when the goal object is inside it) and exclude those that are not causally related to obtaining the goal (for example, opening a box when the goal object is not inside it). One important question for future research concerns the range of causal constraints that infants can use in this way. Children may initially be limited to understanding some of the physical constraints in a context, and only later come to understand the role of psychological and social constraints.

These findings also point to a potential general advance in infants' understanding of intentional action. They indicate that 12-month-olds understand actions as means to an end. When seeking evidence for the onset of intentional control of infants' own behavior, Piaget (1952) focused on means-end problem solving. He reasoned that when infants produce a clearly independent means to obtain a goal (for example, pulling a cloth in order to grasp a toy placed at one end of it) this is evidence that the intention of getting the toy is represented independent of the particular actions that are commonly associated with the toy. Carpenter, Nagell, and Tomasello (1998) have drawn on Piaget's account to propose that a critical step in infants' developing understanding of intentional action is the realization that actions are means to an end. Our findings suggest that this element of intentional understanding is in place by 12 months.

We do not yet know whether infants younger than 12 months understand this aspect of intentional action. Some researchers (e.g., Carpenter, Nagell, and Tomasello) have interpreted Piaget's observations about the development of means-ends problem solving as evidence that they do not. On the other hand, Baillargeon, Graber, DeVos, and Black (1990) report that 5½-month-olds apparently understand means-ends relations when this understanding is measured by visual attention rather than action. It is possible, therefore, that infants younger than 12 months can draw on this knowledge in interpreting goal-directed actions.

Interpreting actions as means to an end is an important advance because it suggests that infants understand goals as separable from the particular actions that they drive. For example, the act of touching the box lid does not express a single, unvarying intention. It could be driven by any one of several underlying goals, such as obtaining the box, exploring the texture of the box, indicating the box to a conversation partner, or opening the box to obtain its contents. The infants in our study seemed to understand this fact: They interpreted the action differently depending on the context in which it occurred. Similarly, the goal of obtaining a toy inside a box could be attained by a number of means, such as opening the box oneself, enlisting the aid of one's mother, or holding the box upside down so that the lid opens. Any one of these actions will work toward obtaining the toy in some situations, but not in others. The understanding that goals and actions exist independently may be a step toward understanding goals as mentally represented entities—that is, understanding that goals exist in people's minds, not in their hands.

Conclusion

We began with two accounts of how early intentional understanding might arise. On one account, the seeds of intentional understanding appear at the end of the first year and are tied to acts of communication and social life. On a second account, the seeds of intentional understanding are present from birth, by virtue of an innate system for interpreting self-propelled motion. Our findings suggest a third alternative: that early in life infants come to understand certain actions as goal-directed, and this early, specific knowledge provides a foundation for later developments in intentional

understanding. This knowledge is evident several months before infants begin to engage in the joint attention behaviors taken to be critical evidence by the first account, and is tied to the specifics of actors and actions in a way that is not predicted by the second.

Premack and others suggest that the earliest notions of intentionality have their origins in the infant's inherited starting equipment. Our findings indicate, in contrast, that initial understandings of intentional action may well derive from experience, because infants' notions of agents and actions seem to be grounded in the particular details of their experience. Infants' notion of agent seems to focus on the person, the type of agent that they most commonly encounter. In addition, the first action that we find infants to understand as goal-directed—grasping—is frequent in everyday life and is an action with which infants literally have firsthand experience. There is circumstantial evidence for a relation between infants' experience as graspers and their understanding of grasping in others as goal-directed. Six-month-olds show strong attention to the goal-related properties of reaches (Woodward 1998), 5-month-olds show this pattern more weakly (Woodward 1998 1999), and, in pilot work, 3-month-olds have not shown this pattern at all. This period, between 3 and 6 months, is the period in which infants become expert reachers.

The ability to understand certain isolated actions as goal-directed is an important initial foothold for infants, but this ability alone would leave much of human behavior uninterpretable. Moreover, to the extent that young infants are limited in this way, their understanding of intentions would be quite different from the mature concept of intentions as independent of particular actions. By 12 months, infants are breaking free of these limitations. They have widened the range of actions understood as goal-directed, can relate distinct actions to overarching goals, and can interpret new actions on the basis of the context in which they occur. Here there are points of contact between our findings and descriptions of burgeoning social competence in 12-month-olds. First, as would be expected from these descriptions, we find that infants begin to understand pointing in terms of the actor-object relation between 9 and 12 months. Second, the finding that 12-month-olds can relate actions to overarching goals could help to explain the onset of behaviors such as imitation of novel goal-directed actions, social referencing, and word learning, since these all require relating dis-

tinct actions to the same goal (e.g., gaze, pointing, emotional expressions, and utterances).

Our findings have just begun to sketch an outline of early intentional understanding. They leave many critical questions unanswered, but nevertheless they give a strong indication of the kind of account that will best address them. Infants' first attempts to make sense of intentional action seem not to be driven by abstract, innate expectations, but rather by what they have learned and by the development of general abilities such as means-end reasoning. In this way, our account dovetails with those of Baird and Baldwin and those of Wellman and Phillips (this volume). There is good reason to believe that infants' understanding of intentional action emerges in large part from general properties of learning and early experience. In the case of intentional action, the information available to infants is rich from the beginning, including countless hours of observing the actions of others as well as the firsthand experience of gaining control over their own actions. These experiences could well provide the initial basis for making sense of intentional action.

Acknowledgments

The research described in this chapter was supported by grants from the John Merck Fund, the Robert R. McCormick Tribune Foundation, and the National Institutes of Health (FIRST grant No. HD35707-01) to the first author. We are grateful to Catharine Seibold and Anneliese Hahn for their assistance in completing the studies, and to the parents and infants who participated. We thank Dare Baldwin and Bertram Malle for insightful comments on an earlier version.

Notes

1. In all our studies, the side of the stage to which the actor reaches during habituation and the type of test trial given first are counterbalanced.
2. The proposal that motion provides the basis for the agent/nonagent distinction is a common starting point for a range of theories, which vary in the extent to which they posit elaborate innate structure (e.g., Gergely et al. 1995; Leslie 1995; Mandler 1992).