Meltzoff, A. N. (2007). Infants' causal learning: Intervention, observation, imitation.
In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 37-47).
Oxford: Oxford University Press.

2

Infants' Causal Learning

Intervention, Observation, Imitation

Andrew N. Meltzoff

Infants' Understanding of Interventions by Self and Other

Causal learning by children combines both observation and action. These two sources of information have not been well integrated in developmental theory. Following Michotte (1963), some developmental scientists argue that young infants are exquisitely tuned observers, and that their perceptual understanding of causality far outstrips their ability to use this information to manipulate the world. Following Piaget (1954), others argue that young infants learn little by pure observation—self-produced motor action is critical; cognitive development generally, and causal reasoning in particular, is charted as a progressive combination of action schemes.

Bayes net approaches provide a way of using both observation and action (in the form of "interventions"), combining them to generate veridical representations of the causal structure in the world. In fact, on some interpretations (Woodward, 2003), the link to intervention is crucial for observed patterns of covariation to be understood as fully causal. The concept of an intervention may help us move beyond a debate about the primacy of perception (Michotte) versus action (Piaget) to theories that map observations and actions to the same abstract causal representations.

For developmental scientists, one striking feature of the philosophical notion of an intervention is that it is abstract—an intervention can be performed by the self or by another person (or even by a "natural experiment" not involving an agent). We can learn not only through our own interventions on the world, but also by watching the interventions of others. This intriguing idea is incompatible with many classical views of infancy, which explicitly deny the equivalence between observing others and acting oneself. In classical developmental views, we observe others from the outside as a series of movements in space, but we feel ourselves from the inside as yearnings, intentions, and freely willed plans. The way we represent self versus other is fundamentally different. This results in a disconnect between learning by doing (self-action) and learning by watching (other's action). A prime developmental achievement is to bring these two modes of learning into line.

There are many ways of testing the psychological linkage between observed and executed interventions. I have used infant imitation, which has several virtues. First, imitation is natural to humans, even babies. Second, in imitating novel acts, infants fashion their interventions based on observing interventions performed by others. Third, it is widely acknowledged that humans are far more proficient imitators than other primates (Meltzoff, 1996; Povinelli, 2000; Tomasello & Call, 1997), and therefore we may be getting at distinctively human cognition by examining human imitation and its development. Fourth, computational models including Bayesian approaches have been applied successfully to both human and robotic imitation (e.g., Demiris & Meltzoff, in press; Meltzoff & Moore, 1983; Rao, Shon, & Meltzoff, 2007).

Historically, there are two principal theories of how infants come to imitate the acts of others: Skinnerian and Piagetian theory. I argue that neither of these can encompass the modern empirical work on infant imitation. The new data are more compatible with the view that there is a fundamental equivalence between the perception and performance of goal-directed acts—an abstract mapping connecting acts seen and acts done—that was not envisioned in the classical frameworks.

Skinner (1953) proposed that young infants cannot imitate the acts of others without specific training. When a young infant sees a mother perform an act such as shaking a rattle to make a sound, the infant does not know what movements to recruit to copy this act. Rather, the mother needs to shape the child's response through operant conditioning. Mom shakes the rattle, and then the infant responds with random motor acts. Mom selectively reinforces those acts that are similar to shaking the rattle. Over time, the mother's shaking comes to serve as a discriminitive cue (a bell or a light would do as well) that elicits the reinforced act (the baby's rattle shaking). To the outside observer, the infant is imitating, but this is *not* because the baby is able to translate the acts seen into acts done. The parent essentially teaches the infant what to do and when to do it through operant conditioning.

This is not an entirely hypothetical example. In fact, Skinner (1953) has shown that pigeons can be conditioned to peck a key when they see other pigeons peck: If Pigeon 1 (P-1) pecks at a key to obtain food and an observer Pigeon 2 (P-2) is reinforced for pecking on seeing this event, then P-2 will eventually be shaped to peck when seeing P-1 pecking. But, P-2 did not learn this intervention on the basis of observing the other animal. All that has happened is that the behavior of P-1 has become a cue for eliciting a conditioned response in P-2. It follows that the observer pigeon could be conditioned to perform a nonimitative act just as easily. Skinner (1953) endorses this implication: "The similarity of stimulus and response in imitation has no special function. We could easily establish behavior in which the 'imitator' does exactly the opposite of the 'imitatee' " (p. 121).

It is known that human infants as young as 3 to 6 months old can be operantly conditioned quite readily (e.g., Rovee-Collier, 1990). This means that they can learn the contingency between their own actions and results in the world. But, the capacity for operant conditioning does not mean that the infant can learn these action-outcome relations from observing the acts of others. In other words, the fact that infants can learn an intervention through their *own* trial and error (learning by doing) does not mean that they can learn to perform the intervention on the basis of observing the interventions of others (learning by watching). The latter would be imitation. The former is just a special case of operant conditioning in which a friendly demon (a clever mother or experimenter) has arranged it so the discriminative cue matches the reinforced response. The moral is that if we want to know whether infants can learn an intervention through observation, then we need to know the infant's reinforcement history or, failing that, use a novel act for which prior shaping is unlikely.

Piagetian theory (1962) came to similar conclusions as Skinner, albeit for entirely different reasons. Piaget also thought that young infants could not imitate spontaneously. In Piaget's case, it was not that infants needed to be conditioned to learn to imitate, but rather that they needed to reach a certain stage of cognitive sophistication. Piaget realized that translating a seen intervention into one executed by the self was nontrivial, and he claimed it was beyond the capacity of infants in the first half year of life. He hypothesized that infants were "egocentric," even "solipsistic." The youngest infants could not learn novel acts from observing others (whether these acts were complex means-ends relationships or simple body acts) because learning at first occurred through self-action independently of other people (what Piaget called *practical intelligence*).

The Piagetian concept of infantile egocentrism was most famously illustrated in his predictions about facial imitation. Infants can see you make a facial movement, but they cannot see their own faces. If the infant is young enough, he or she will never have seen his or her face in a mirror. How could the infant link the observed facial acts of others with personal unseen bodily acts? According to Piaget, this "invisible imitation" was impossible because self and other were known in such different terms; there was no abstract framework for connecting observation and performance. Piaget (1962) put it this way: "The intellectual mechanism of the child will not allow him to imitate movements he sees made by others when the corresponding movements of his own body are known to him only tactually or kinesthetically (as, for instance, putting out his tongue) . . . since the child cannot see his own face, there will be no imitation of movements of the face [before approximately 1 year old]" (p. 19).

Thus, Piaget shared Skinner's view that actions could be observed and performed, but that the *observation* of an act did not engender the *production* of a matching act without a long path of prior learning. Neither Skinner nor Piaget thought that imitation was a mechanism for early learning; rather, imitation itself needed to be learned, and a good deal of theoretical effort was put into explaining how babies could eventually associate the observation of others' actions with manipulations performed by the self.

Newborn Imitation: Innate Mapping Between Observation and Execution

In part because of Skinner's and Piaget's theories about a gulf between the observation and the execution of human acts, I designed a series of tests of facial imitation in young infants. Contrary to classical theories, the results show that newborns imitate facial gestures. The work suggests an abstract notion of goal-directed action that cuts across the observed acts of others and one's own freely willed actions.

In an early study, Meltzoff and Moore (1977) tested facial imitation in 2- to 3-week-old infants. The results showed that they could imitate four different adult gestures: lip protrusion, mouth opening, tongue protrusion, and finger movement. The mapping between observation and execution was quite specific: Infants confused neither actions nor body parts. They differentially responded to tongue protrusion with tongue and not lip protrusion, revealing an innate body scheme that maps from observed body parts to their own body, despite never having seen their own face. Similarly, they responded accurately to lip protrusion versus lip opening, showing that different patterns of action can be extracted and imitated when the specific body part is controlled.

As my psychology colleagues quickly pointed out, these infants may not have been young enough to answer the objections of Skinner and Piaget. In their 2 weeks of life, they might have learned the relevant associations. Perhaps mothers conditioned their children to stick out their tongues whenever they saw this gesture. The definitive test involved newborns who averaged 32 hours old at the time of the test. The oldest infant was 72 hours old, and the youngest was just 42 minutes old. The newborns accurately imitated (Meltzoff & Moore, 1983, 1989). Apparently, facial imitation is innate. This suggests a fundamental equivalence between the perception and production of acts that is built into the mind of the human baby.

Goal-Directedness in Early Imitation

Does facial imitation involve a "goal-directed" act? In this chapter, I discuss goal-directed acts that cause something to happen in the world. These simple bodily acts do not do that. Nonetheless, I think that early imitation is goal directed.

A characteristic of goal-directed action is that it converges toward the endpoint along flexible routes. This has been demonstrated in early imitation. Accurate imitation does not pop out fully formed. Infants have to work on it. They make errors and gradually correct their motor attempts to achieve a more accurate match to the observed target (Meltzoff & Moore, 1994). This error correction occurs even though the adult gives no feedback to the child (no smiles or encouragement) and, most important, even though the child observes the others' act but not their own.

The goal directedness of the response is also illustrated in the "creative errors" infants make. One study showed infants the novel gesture of poking out the tongue at 45° off midline (from the side of the mouth) (Meltzoff & Moore, 1994). The predominant pattern was to poke the tongue into the inside of their check and then gradually adjust. However, some infants adopted a novel approach. They poked out their tongues and simultaneously turned their heads to the side, thus creating a new version of "tongue to the side" (Meltzoff & Moore, 1997). This head movement was not something the adult demonstrated but was the infants' construction of how to combine a tongue protrusion and an off-midline direction. Although the literal muscle movements were very different, the end-state orientation of the tongue was similar, and in this sense it can be seen as an act organized by a goal.

The Innate Representation of Human Action

One way of accounting for these results is to hypothesize that infants innately represent the perception and performance of elementary human acts using the same mental code. There is thus something like an act space or primitive body scheme that allows the infant to unify the visual and motor information into one common "supramodal" framework (Meltzoff & Moore, 1997).

The nature of the supramodal framework can be further dissected. Three pieces of data suggest that the supramodal system is not simply a Gibsonian resonance device that directly turns observations into like movements—a perception-production transducer. First, the voluntary nature of the response indicates that the infant need not produce what is given to perception. The observations of others' acts can be stored and accessed after a delay. At minimum, there is an intermediary representation and not simply an automatic transduction. Second, as we have seen, infants correct their imitative efforts (and make creative errors). Information about one's acts has to be available for comparison to the representation of the adult's act, but the representation of the observed act is not confused with or modified by one's own multiple motor attempts. Third, infants show special interest in being imitated themselves; they recognize when their behavior is being copied (Meltzoff, 2007). Such recognition implies that there is a representation of their bodily acts.

This takes us beyond the simple transducer story. The data suggest a differentiation in the supramodal system. The representation of the observations are tagged to keep them differentiable from the representation of one's own motor acts. The cognitive act is to compare these two representations—in one case to match one's own acts to the other (imitative correction) and in the other case to detect being matched oneself (recognizing being imitated). The mental code may be abstract enough to unite perception and production, but the representations deriving from observation and self-action are not confused. They retain some source information (e.g., tongue-beyond-lips [observed] and tongue-beyond-lips [produced]).

I would argue that this fundamental equivalence (with differentiation) between self and other is a starting point for social cognition, not an endpoint reached after months of postnatal learning à la Piaget. The chief goal for the remainder of the chapter is to flesh out the thesis that these innately registered equivalences between observed and self-generated actions provide a substrate for infants' learning causal relations from others' interventions.

Learning Interventions From Observation: Making Things Happen

Adults manipulate objects to cause other things to happen in the world. Infants carefully observe adult's causally directed acts and begin reproducing what they see as soon as they become capable of handling objects.

One study tested whether 14-month-olds could learn an intervention purely from observation. To ensure that a new causal relation was being learned, a novel act was used (Meltzoff, 1988). The adult put a flat box on the table, looked down at it, and then bent from the waist, touching it with his head, which caused the top panel to light up. (This was an early blicket detector that was activated by human heads.) Later, when infants were given the box themselves, 67% of them leaned forward from the waist and touched the panel with their own foreheads. Many kept their eyes open, staring at the top of the box, and smiled when the light came on. Control infants showed that the baseline probability of infants touching the panel with their foreheads was literally 0%. Not a single infant did so in the absence of seeing the intervention. In a recent study, I changed the head-touch apparatus to incorporate a remote effect. When the adult touched the box with his forehead, this caused a remote box to light up. The remote box was 2 feet away. When the infants were given their turn, they touched the adult's box with their foreheads and immediately turned to stare at the remote box, waiting for it to activate (Meltzoff & Blumenthal, 2006). Carpenter, Nagell, & Tomasello (1998) reported related effects. Taken together, the experiments show that infants can learn novel interventions based purely on observation.

A Privileged Role for Manipulations Performed by Self

These head-touch studies show that infants can learn an intervention by watching others. Is anything added if infants perform the intervention themselves (Kushnir & Gopnik, 2005; Meltzoff, 2006)?

I conducted a relevant study with 14-month-olds (Meltzoff, 2006). Infants were randomly assigned to two groups. Infants in Group 1 watched the adult perform manipulations of two novel objects. The experimenter shook one object to cause it to make a sound; he held another one from a string and bounced it up and down on the tabletop. Infants observed these acts and then were sent home without manipulating the objects themselves. Infants in Group 2 were treated similarly but were immediately given the objects before being sent home. Virtually all of them imitated the actions they saw and thus had manipulatory experience as well as observational information.

The critical test came the next day when both groups returned to the laboratory, and the objects were put before them. The adult gave no hint what to do. Infants who had been given the opportunity for immediate imitation performed significantly more of the target acts on Day 2. Something appears to be gained if infants perform the action themselves directly after observing it. Infant performance is boosted if they quickly convert an observed manipulation into a self-produced manipulation. In line with the work on facial imitation, it appears that the actions of self and other are coded in commensurable terms, but that the self-produced acts are tagged distinctively from acts that were merely observed; converting observation into a self-action makes it memorable.

I hasten to add that infants can remember causal events without taking concurrent action. We know this because the first group of children, who only observed on the first day (by experimental design), imitated from memory on the next day. Evidently, the observed intervention can be stored and used to generate one's own manipulations after a delay. But, it is equally interesting that memory for the causal act is stronger if the act is first performed by the self before the delay.

Inferring an Intervention Based on Unsuccessful Action Patterns

Learning Actions Versus Learning Outcomes

We have seen that infants who see an adult use unusual means to accomplish an intervention do not simply reproduce the result (making the light come on) using any motor acts at their disposal (e.g., their hands), but instead faithfully copy the whole behavioral envelope. Based on this research, one might wonder whether means and ends are differentiable aspects of an intervention, or whether infants achieve causal results by reenacting the precise actions used by the adults.

This makes a difference to theories because it could be that (a) infants faithfully copy the adult's actions, and sharing body types and the laws of physics, the causal results naturally follow; or (b) infants represent the causal results and strive to achieve them by their own invented means. This is a tricky distinction to test empirically because if infants copy our actions, then they are likely to achieve our causal results "for free."

The way I investigated this question was to have infants observe an unsuccessful intervention. I wanted to test whether infants can read through our failed attempts and infer the intervention we intended to achieve. Because the adult's actions were unsuccessful, infants could not copy the adult's actions and achieve the desired result.

Inferred Interventions

I showed 18-month-olds unsuccessful interventions (Meltzoff, 1995). For example, the adult used a stick tool in an attempt to push a button to make a sound but "accidentally" under- or overshot the target. Or, the adult grasped the ends of a dumbbell-shaped object and attempted to yank it in two, but his hands slid off as he yanked, and thus the goal was not achieved. To an adult, it was easy to decode the actor's intended intervention. The measure of how infants interpreted the event was what they chose to reenact. In this case the "correct answer" was not to imitate the manipulation that was seen (the unsuccessful attempt), but to perform the intervention the adult "meant to do."

The study compared infants' tendency to perform the target act in several situations: (a) after they saw the successful intervention demonstrated, (b) after they saw the unsuccessful attempt to perform the intervention, and (c) after the intervention was neither shown nor attempted (control). The results showed that 18-month-olds can infer interventions from adult attempts to perform them. Infants who saw the unsuccessful attempts and infants who saw the successful interventions both performed the goal acts at a significantly higher rate than the controls. Evidently, infants can understand our goals even if we use means that are insufficient to fulfill them.

In further work, 18-month-olds were shown similar displays but were handed a trick toy that prevented them from performing the intervention (Meltzoff, 2006). For example, the dumbbell-shaped object was surreptitiously glued shut. If infants attempted to pull it apart, then their hands slipped off the ends, duplicating the adult's behavior. The question was whether this satisfied infants. It did not. They did not terminate their behavior. They varied the way they yanked on the dumbbell, systematically changing their interventions to find one that worked. They also appealed to their mothers and the adult for help. About 90% of the infants looked up at an adult within 2 seconds after failing to pull apart the trick toy, and many vocalized while staring at the adult's face. Why were they appealing for help? They had matched the adult's surface behavior, but evidently they were striving toward something else—the adult's intended intervention.

Inventing New Means to Achieve an Inferred Intervention

If infants are inferring the adult's goal, then they should also be able to achieve it using a variety of means. I tested this. As before, an adult grasped the ends of a gigantic dumbbell and attempted to yank it apart, but his hands slid off. The dumbbell was then presented to the infants. Infants did not even try to copy the adult's exact movements. Rather, they put their tiny hands on the inside faces of the cubes and pushed outward, or stood upright and used both hands to pull upward, and so on. They used different means than the experimenter, but these acts were directed toward the same causal result. This fits with the hypothesis that the infants had inferred the goal of the intervention, differentiating it from the surface behavior that was observed.

Work by Want and Harris (2001) goes further and shows that older children, 3-year-olds, benefit from observing others using multiple means to achieve a goal. They benefit more from watching an adult modify a failed attempt into a successful act than from watching the demonstration of successes alone. Other work also underscores the importance of goals in imitation (e.g., Gattis, Bekkering, & Wohlschläger, 2002; Gleissner, Meltzoff, & Bekkering, 2000; Williamson & Markman, 2005).

Agents and Goals: Infants Infer Interventions for Agents

In the adult commonsense framework, the acts of people can be goal directed and intentional, but the motions of inanimate devices are not; they are governed by physics, not psychology. Do infants interpret the world in this way? Meltzoff (1995) designed an inanimate device made of plastic and wood. The device had short poles for arms and mechanical pincers for hands. It did not look human, but it traced the same spatiotemporal path the human actor traced and manipulated objects much as the human actor did. When the pincers slipped off the ends of a dumbbell, infants did not infer the intervention as they did with the human agent. The infants were no more (or less) likely to pull the toy apart after seeing the unsuccessful attempt of the inanimate device than infants in the baseline condition. However, if the inanimate device successfully completed this act, then infants did perform the successful intervention.

Evidently, infants can understand and duplicate a successful intervention displayed by the inanimate device but do not read meaning into the device's unsuccessful "attempts." This makes sense because successes lead to a visible change in the object. Failures leave the object intact and therefore must be interpreted at a deeper level, in terms of the intended interventions of the agent. Perhaps infants do not interpret inanimate devices as psychological agents with goals and intentions; thus, no intervention is inferred.

In summary, the research shows that infants distinguish between what the adult meant to do and what he actually did. They ascribe goals to human acts; indeed, they can infer an intended intervention from a pattern of behavior (multiple unsuccessful attempts) even when the intervention was not performed. The acts of persons—but not the motions of mechanical devices—are understood within an agentive framework involving goals and intentions.

A Natural Experiment: The Primacy of People in Infants' Notion of Interventions

The Involvement of People Causes Infants to Interpret the Same Scene Differently

As we have seen, infants interpret the acts of people in special ways. This suggests a way of testing Woodward's idea of a natural experiment in which a causal event occurs without an agent as the source of the change. I showed 18-month-old infants an intervention and varied whether a person was involved in producing the result.

Infants saw the dumbbell-shaped object in three successive states. The three views were separated from each other by raising a black screen, so that the infants saw three snapshot views of an event that unfolded over time. What varied is the causal story of how it got to be that way. After infants saw the three displays, they were given the dumbbell. The question was whether they produced the target behavior, which was to pull the object apart.

Group 1 was a baseline control condition to assess infants' spontaneous tendency to manipulate the object. For this control group, infants simply saw three identical states—the assembled object sitting in place with no person present. As expected, infants did not pull the object apart spontaneously: They mouthed it, banged it, and slid it across the table, but they did not spontaneously discover pulling it apart in the absence of seeing this intervention. For Group 2, the three snapshots revealed the affordances of the object but did not specify the involvement of a person. The views were the following: (a) object assembled, no person present; (b) object disassembled, no person present; (c) object assembled again, no person present. For Group 3, the snapshots revealed an agent as a potential cause. The views were the following: (a) object assembled, in person's hands; (b) object disassembled, in person's hands; (c) object assembled, in person's hands.

Infants in Group 2 did not pull apart the toy; in fact, they did not differ from the baseline controls. In contrast, infants in Group 3 pulled the object apart significantly more often than those in Group 1 or Group 2. Thus, the involvement of a person as a potential cause led infants to interpret the same scene differently. In the case of the natural experiment in which the object was seen in its pre- and posttransformed state (Group 2), infants observed but did not try to re-create the event. However, if a person held the object, although sitting stony-faced and displaying no effort at acting, infants did so.

These results are especially interesting when combined with the Meltzoff (1995) intention-reading study. In that study, the dumbbell remains untransformed, and the person is trying to perform an intervention. In the current study, the person is present and shows no intent, but the results of the object transformation are shown (three static states: assembled, apart, reassembled). In the former case, there is human effort and no object transform, and this suffices for infants to infer the intervention. In the latter case, there is an outcome state and no effort, but if the person is present, this can be interpreted as a potential cause for what happened. In both cases, it provides enough for infants to interpret the observations as relevant to their own actions and for them to fill in the blanks and produce a manipulation that was never directly observed but only inferred.

Agentless Transformations: Magic

In a further study, the dumbbell object was magically pulled apart and reassembled in front of the child's eyes but appeared to do so autonomously. This provides object transformation data in full view.

The object was placed on a black box, and inside the box there were magnets. The magnets were moved, and thus the dumbbell came apart and was reassembled without this being caused by a human agent. The results were that 15-month-olds did not pull the toy apart at any higher than baseline levels. Interestingly, half of the infants picked up the object and placed it back on the box several times, as if situating the object on the magic spot would cause the result. Infants saw the intervention and wanted it to repeat but, in the absence of a human cause, drew no implication for their own causal actions. Evidently, they thought the object transformation would happen *to* the object if it was spatially positioned, rather than thinking *they* could cause the transform through their own manipulation.

Learning to Use a Tool

In the developmental and animal psychology literatures, one of the most celebrated examples of causal reasoning is the case of tool use. We know a lot about the ability of chimpanzees to use tools—starting from Köhler's (1927) observations of Sultan moving crates below an overhead banana to reach it and extending to Jane Goodall's (1968) reports of termite fishing on the Gombe Stream Reserve. Although it was once argued that tool use was uniquely human, it is now widely acknowledged that other animals are successful tool users, including the gold standard of using a stick to obtain an out-of-reach target. The debate concerns whether animals use tools based on trial and error or based on insight about the causal relations involved (Povinelli, 2000; Tomasello & Call, 1997).

For the purposes of this chapter, I am interested in exploring tool use from a different perspective. Instead of asking whether animals and infants use tools when left on their own to "figure it out," I wish to examine learning through observation—in particular, seeing an expert use a stick to obtain an out-of-reach goal. The extant data are mixed. Tomasello and Call (1997) suggest that wild chimpanzees do not readily learn how to use a tool from observation, but that some enculturated chimps may; Povinelli (2000) remains skeptical of the latter.

The literature concerning human infants is similar. There is good evidence that infants can eventually learn to use sticks as tools when left to their own devices (Bates, Carlson-Luden, & Bretherton, 1980; Brown, 1990; Piaget, 1954) but much sparser evidence concerning learning from the interventions of others. Of course, it is well known that adults and older children learn how to use a wide variety of tools and complex machinery by watching experts; the debate concerns younger ages.

To test for observational learning of tool use, one needs a few conceptual distinctions. To begin, one needs to distinguish imitation from stimulus enhancement. The latter refers to the fact that the infants' attention may simply be drawn to a tool by virtue of the adult handling it. With their attention drawn to the stick, infants may increase their random play with the object, thereby increasing the probability that they will learn through trial and error that it can be used as a tool. The child is not learning a new causal relation based on what they see the other do. Rather, the child is learning that the stick is interesting—stimulus enhancement—and thereby is more likely to pick it up, with the rest following by chance or trial and error.

In the developmental literature, there have been surprisingly few well-controlled tests of learning to use complex tools through observation. Nagell, Olguin, and Tomasello (1993) performed a relevant experiment comparing chimps and human infants. They reported that the 18-month-old children failed to learn how to use a rake (to obtain a distant object) from observation, but that 24-month-olds could do so.

I tested younger infants. The sample consisted of 120 infants evenly distributed at 16, 18, 20, and 22 months of age (Meltzoff, 2006). Within each age group, infants were randomly assigned to one of three test conditions: (a) learning by observation, in which the adult modeled the correct use of the rake to obtain the out-of-reach goal; (b) Control 1 (baseline), in which infants saw no modeling and were simply given the rake; and (c) Control 2 (stimulus enhancement), in which infants saw the adult use the rake to touch the goal, thereby drawing attention to the rake and to the fact that it could make spatial contact with the goal (correct use of the rake was not shown).

The tool was a 17-inch long rake. It was placed horizontally in front of the infant, with approximately a 2-foot spatial gap between it and the goal object. The goal was a highly desirable rubber giraffe. Infants had 1 minute to solve the problem. Preliminary studies in our lab suggested that infants performed better when they observed the model from a first-person perspective—when the adult and infant were side by side, rather than facing each other across the table. This may be important because previous studies have not modeled tool use from this perspective (e.g., in the Nagell et al. 1993 study the adult faced the infant, so the modeling entailed using the tool to pull the object away from the infant and toward the adult). Viewing the goal-directed act of the model from the same perspective as one's own may facilitate learning from observation.

Infants showed great enthusiasm for obtaining the goal (stretching out their arms, vocalizing, looking at the adult, etc.). In the two control groups, there was no significant difference in the successful use of the rake as a function of age. Across all 120 subjects, only 7.5% (6 of 80) of the infants solved the problem spontaneously; in contrast, fully 50% (20 of 40) of the infants succeeded after they saw the adult show them how to use the tool, $p < .001$. The older infants (20- and 22-month-olds) profited far more from observation (70% succeeded) than did the younger infants (30% succeeded), $p < .05$.

Infants learn from observation but not automatically. There appears to be an interaction between the infants' initial cognitive level and what they gain from observing others. The young infants learn, but they do not exceed spontaneous rates by the same degree that the older infants do. I would predict that still younger infants would not learn how to use the rake from observation. I say this because of the nature of the failures. After watching the expert adult, the younger infants pounce on the rake and wield it with great confidence. However, once they move the rake to the quarry, they are not be able to "think through" the causal relations—that the business end of the rake has to be behind the goal-object and the tines pointed downward before the rake could be pulled in. (Their reaction reminds me of undergraduates who get halfway through a difficult conceptual distinction and then, face fallen, find themselves lost, unable to bring things to conclusion. The "uh-oh, what-do-I-do-next" expression seems to be invariant across age.) One possibility that arises from this work is that infants have to be "on the cusp" of solving the problem themselves to get the boost from seeing how someone else solves it (see Gopnik & Meltzoff, 1986, 1997, for related findings). The older infants would be more intelligent consumers of the observed interventions. I am exploring this possibility through further research.

Conclusions

The work described in this chapter has implications for both psychology and philosophy.

Psychology

The power of imitation has always been underestimated in psychology. Skinner underestimated imitation because he thought it was simply a variant of operant conditioning in which the infants' response had been shaped up. Just as infants could be trained to perform Behavior X when they saw a red light, so they could be trained to perform Behavior X in response to Behavior X. There was nothing special about the match between self and other. Skinner thought that the opposite behavior would do just as well as a cue. I doubt it. I think you would be in for a long series of training sessions if you tried to teach a baby to open his or her hand every time the baby saw you close yours. The intrinsic connection would interfere with learning the arbitrary association.

Chomsky underestimated imitation because it was a learning mechanism. To say children learn through imitation means that they are sculpted by experience. Chomsky relegated experience to "parameter setting" or the "triggering" of innately structured systems. It is difficult to see how these concepts can explain the imitation of novel acts like head-touch. Infants duplicate this act, but it is unlikely to be biologically specified and simply triggered. Chomsky may (or may not) be correct about the domain of grammar, but in the domain of action, observing others' novel acts has a powerful effect of sculpting infants' own actions. Parents do not need to slavishly condition their child for the child to begin to act like those around the dinner table. The babies are observing and learning. Moreover, research suggests that "auditory observation" may be more powerful in language acquisition than traditionally assumed, particularly for the acquisition of culturally specific phonology. Kuhl (2004) reports that infant phonology, as indexed by both brain measures and perceptual measures, is influenced by the sounds infants hear in their culture; furthermore, studies show that young infants reproduce speech sounds they hear through imitation (Kuhl & Meltzoff, 1996).

Piaget underestimated imitation because he thought that infants were born with "heterogeneous spaces"—a "visual space" that was initially independent of their "motor space." A major task of the first 2 years of life was to unify these spaces so infants could learn from watching, not just from doing. Piaget predicted that facial imitation was impossible until about 1 year of age and deferred imitation (imitation from the memory of observed, now absent, events) impossible before about 18 months of age. My research shows facial imitation at birth and deferred imitation soon thereafter.

These theorists missed the idea that there is a fundamental equivalence between observing and performing goal-directed motor acts. It is not that seeing and doing need to be linked by associative learning or conditioning. Imitation is innate. Infants can even imitate facial gestures they have never seen themselves perform. Infants have an abstract mental code, we call it a supramodal code, that unites acts seen and acts done within the same framework.

The innate equivalence between elementary acts of self and other has implications for learning about cause-effect relations. Instead of relying exclusively on the contingencies between your acts and the consequences in the world, you can learn through observing the actions of others—actions that you immediately recognize as "like my own." If acts performed by another make something happen, perhaps they will make the same thing happen when I do them. Such learning could not get off the ground if the observed acts were not recognized to be the same as my own acts. That much is nature's share.

Philosophy

Woodward (chapter 1, this volume) describes three levels of causal understanding:

1. A *purely egocentric* causal view: The subject understands the relationships between personal actions and the resulting effects but is unable to grasp that the same relationships can occur when the self is not the cause.
2. An *agent causal* view: The subject understands that the causal relationships that exist between personal actions and effects also apply to the actions of other people.

3. A *fully causal* view: The subject understands that the same causal relationships that the subject exploits in intervening can also be used by other agents and can exist in nature even when no other agents are involved.

The egocentric infant described by Piaget's theory (1952, 1954) closely resembles what Woodward called the *egocentric causal view*. This egocentric organism is capable of being conditioned because he or she can grasp the relation between bodily movements and effects in the world but cannot learn from watching the causal actions of others. The modern empirical results suggest that the egocentric infant is a fiction. Laboratory rats and other animals may conform to this description, but the human infant does not.

There is evidence, however, that up to about 18 months of age, the human infant is not fully causal in Woodward's sense. Several experiments suggest that the human infants learn interventions differently from a person than from an inanimate device (inferred intervention studies) and draw only limited inferences when no agent is present (natural experiment studies). Based on the current research, it may be that Woodward's (chapter 1, this volume) characterization of an "agentive view" is a reasonable description of the prelinguistic toddler. How and when an infant develops into a fully causal agent is a central question for developmental cognitive science (Gopnik et al., 2004; Meltzoff, 2006).

Summary

The perception of others' actions and production of self-action are mapped onto commensurate representations starting from birth. This allows infants not only to learn interventions through their own manipulations but also to multiply greatly their learning opportunities by observing the manipulations of others and profiting from them. For example, in the novel head-touch case, infants immediately knew how to activate the object 24 hours after seeing the adult do so, without ever having handled the object themselves. Importantly, infants do not seem to confuse acts of self and other. On the one hand, they correct their behavior (showing a retention of the observed target that is differentiable from the self's motor efforts). On the other hand, they treat their own acts in a privileged manner that suggests some sort of mental tagging that helps track whether an act was of external or internal origins.

Infants imitate but do not blindly copy everything they see. First, they make creative errors. Second, they skip over the literal behavior they see and choose to duplicate inferred interventions—what the adult meant to do, not what the adult did do. Third, when causal relations are difficult, as in the rake case for younger infants, observation alone does not seem to guarantee success; older infants glean more from the modeling than do younger ones.

Starting at birth, there seems to be a delicate interplay between learning by observation and learning by doing. The two are not quarantined from each other as Michotte (with an emphasis on observation over motor experience) or Piaget (with an emphasis on motor experience over pure perception) might have supposed. Instead, there seems to be a reciprocal exchange between these two modes of learning. What infants observe influences what they do (novel head-touch imitation), and what they can do changes their attention to the model and how they interpret it (tool use from observation).

ACKNOWLEDGMENTS This work was supported by the National Institutes of Health (HD-22514), the James S. McDonnell Foundation, and the Tamaki Foundation. I thank Alison Gopnik for her causal powers as editor, colleague, and friend. I also acknowledge helpful conversations with Jim Woodward, Laura Schulz, and other members of the McDonnell Foundation Causal Inference Workshop at the Stanford Center for Advanced Study in 2003. Thanks also to Craig Harris and Calle Fisher for help on the final stages of chapter preparation.

References

Bates, E., Carlson-Luden, V., & Bretherton, I. (1980). Perceptual aspects of tool using in infancy. *Infant Behavior & Development*, 3, 127–140.

Brown, A. L. (1990). Domain-specific principles affect learning and transfer in children. *Cognitive Science*, *14*, 107–133.

Carpenter, M., Nagell, K., & Tomasello, M. (1998). Social cognition, joint attention, and communicative competence from 9 to 15 months of age. *Monographs of the Society for Research in Child Development*, *63*(4, Serial No. 255).

Demiris, Y., & Meltzoff, A. N. (In press). The robot in the crib: A developmental analysis of imitation skills in infants and robots. *Infant and Child Development*.

Gattis, M., Bekkering, H., & Wolschläger, A. (2002). Goal-directed imitation. In A. N. Meltzoff & W. Prinz (Eds.), *The imitative mind: Development*, *evolution, and brain bases* (pp. 183–205). Cambridge, England: Cambridge University Press.

Gleissner, B., Meltzoff, A. N., & Bekkering, H. (2000). Children's coding of human action: Cognitive factors influencing imitation in 3-year-olds. *Developmental Science*, 3, 405–414.

Goodall, J. (1968). The behavior of free-living chimpanzees in the Gombe Stream Reserve. *Animal Behavior Monographs*, *1*, 3.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*, 3–32.

Gopnik, A., & Meltzoff, A. N. (1986). Relations between semantic and cognitive development in the one-word stage: The specificity hypothesis. *Child Development*, *57*, 1040–1053.

Gopnik, A., & Meltzoff, A. N. (1997). *Words, thoughts, and theories*. Cambridge, MA: MIT Press.

Köhler, W. (1927). *The mentality of apes* (E. Winter, Trans.) (2nd ed.). London: Routledge & Kegan Paul.

Kuhl, P. K. (2004). Early language acquisition: Cracking the speech code. *Nature Reviews Neuroscience*, *5*, 831–843.

Kuhl, P. K., & Meltzoff, A. N. (1996). Infant vocalizations in response to speech: Vocal imitation and developmental change. *Journal of the Acoustical Society of America*, *100*, 2425–2438.

Kushnir, T., & Gopnik, A. (2005). Young children infer causal strength from probabilities and interventions. *Psychological Science*, *16*, 678–683.

Meltzoff, A. N. (1988). Infant imitation after a 1-week delay: Long-term memory for novel acts and multiple stimuli. *Developmental Psychology*, *24*, 470–476.

Meltzoff, A. N. (1995). Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Developmental Psychology*, *31*, 838–850.

Meltzoff, A. N. (1996). The human infant as imitative generalist: A 20-year progress report on infant imitation with implications for comparative psychology. In C. M. Heyes & B. G. Galef (Eds.), *Social learning in animals: The roots of culture* (pp. 347–370). New York: Academic Press.

Meltzoff, A. N. (2006). The "Like Me" framework for recognizing and becoming an intentional agent. *Acta Psychologica*, doi:10.1016/j.actpsy.2006.09.005.

Meltzoff, A. N. (2007). "Like Me": A foundation for social cognition. *Developmental Science*, *10*, 126–134.

Meltzoff, A. N., & Blumenthal, E. J. (2006, April). Causal monitoring. Paper given at the meeting of the McDonnell Foundation Causal Learning Workshop. Pasadena, CA.

Meltzoff, A. N., & Moore, M. K. (1977). Imitation of facial and manual gestures by human neonates. *Science*, *198*, 75–78.

Meltzoff, A. N., & Moore, M. K. (1983). Newborn infants imitate adult facial gestures. *Child Development*, *54*, 702–709.

Meltzoff, A. N., & Moore, M. K. (1989). Imitation in newborn infants: Exploring the range of gestures imitated and the underlying mechanisms. *Developmental Psychology*, *25*, 954–962.

Meltzoff, A. N., & Moore, M. K. (1994). Imitation, memory, and the representation of persons. *Infant Behavior & Development*, *17*, 83–99.

Meltzoff, A. N., & Moore, M. K. (1997). Explaining facial imitation: A theoretical model. *Early Development and Parenting*, *6*, 179–192.

Michotte, A. (1963). *The perception of causality* (T. R. Miles & E. Miles, Trans.). London: Methuen & Co. Ltd.

Nagell, K., Olguin, R. S., & Tomasello, M. (1993). Processes of social learning in the tool use of chimpanzees (*Pan troglodytes*) and human children (*Homo sapiens*). *Journal of Comparative Psychology*, *107*, 174–186.

Piaget, J. (1952). *The origins of intelligence in children* (M. Cook, Trans.). New York: International Universities Press.

Piaget, J. (1954). *The construction of reality in the child* (M. Cook, Trans.). New York: Basic Books.

Piaget, J. (1962). *Play, dreams and imitation in childhood* (C. Attegno & F. M. Hodgson, Trans.). New York: Norton.

Povinelli, D. J. (2000). *Folk physics for apes: The chimpanzee's theory of how the world works*. New York: Oxford University Press.

Rao, R. P., Shon, A. P., & Meltzoff, A. N. (2007). A Bayesian model of imitation in infants and robots. In K. Dautenhahn & C. L. Nehaniv (Eds.), *Imitation and Social Learning in Robots, Humans and Animals: Behavioural, Social and Communicative Dimensions*. Cambridge, UK: Cambridge University Press.

Rovee-Collier, C. K. (1990). The "memory system" of prelinguistic infants. In A. Diamond (Ed.), *The development and neural bases of higher cognitive functions* (Vol. 608, pp. 517–542). New York: New York Academy of Sciences.

Skinner, B. F. (1953). *Science and human behavior*. New York: Macmillan.

Tomasello, M., & Call, J. (1997). *Primate cognition*. New York: Oxford University Press.

Want, S. C., & Harris, P. L. (2001). Learning from other people's mistakes: Causal understanding in learning to use a tool. *Child Development*, *72*, 431–443.

Williamson, R. A., & Markman, E. M. (2006). Precision of imitation as a function of preschoolers' understanding of the goal of the demonstration. *Developmental Psychology*, *42*, 723–731.

Woodward, J. (2003). *Making things happen: A theory of causal explanation*. New York: Oxford University Press.