Varieties of Understanding

New Perspectives from Philosophy, Psychology, and Theology

Edited by

STEPHEN R. GRIMM

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Oxford University Press is a department of the University of Oxford. It furthers the University's objective of excellence in research, scholarship, and education by publishing worldwide. Oxford is a registered trade mark of Oxford University Press in the UK and certain other countries.

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America.

Library of Congress Cataloging-in-Publication Data
[To come]
ISBN 978–0–19–086097–4

1 3 5 7 9 8 6 4 2

Printed by Sheridan Books, Inc., United States of America

Contents

Contributors

Christos Bechlivanidis, University College London
Kimberly A. Brink, University of Michigan
Elisabeth Camp, Rutgers University
Richard Foley, New York University
Jennifer Gosetti-Ferencei, Johns Hopkins University
Anthony Gottlieb, Oxford University
Stephen R. Grimm, Fordham University
Hyowon Gweon, Stanford University
Babak Hemmatian, Brown University
Frank Keil, Yale University
David Lagnado, University College London
Tania Lombrozo, Princeton University
Steven Sloman, Brown University
Ernest Sosa, Rutgers University
Terrence W. Tilley, Fordham University
Henry M. Wellman, University of Michigan
Daniel Wilkenfeld, University of Pittsburgh
Linda Zagzebski, University of Oklahoma
Jeffrey C. Zemla, Brown University

C9

9

Understanding Others to Learn and Help Others Learn

Inferences, Evaluation, and Communication in Early Childhood

Hyowon Gweon

C9.P1 Think about all the things you know: the names of things and places, the planets in the solar system, your favorite pasta recipe, or the meaning of "democracy." Now, think about *how* you've come to know all this. You might soon realize that much of your knowledge about the world comes from what you've learned from other people.

C9.P2 Humans are powerful learners. Within just a few years of life, young human learners construct a coherent system of intuitive theories about how the world works (Carey, 2011; Gopnik & Wellman, 2012; Schulz, 2012). What is even more striking is that this remarkable feat seems to happen almost effortlessly; children learn as they naturally play and interact with their environment. One important thing to note, however, is that humans, especially young children, rarely experience the world in isolation. Even though we learn a lot from our firsthand experiences with the *things around us*, these experiences are often heavily intertwined with the *people around us*. We observe other people as they navigate our surroundings, and as they show, tell, or teach us about faraway places, planets in the solar system, and the concept of democracy. Through these interactions, we learn things that go far beyond what we can directly experience. A fundamental challenge for theories of human learning is to provide a unified account for how we combine our own experience with the world and information provided by others in order to learn so much, so quickly, and so accurately, in a complex, noisy environment.

C9.P3 One remarkable and distinctive feature of human social learning is that it involves more than simply adopting others' behaviors through observation and imitation; it also involves adopting abstract knowledge that resides in others' minds, such as their causal and ontological beliefs about how the world works. This process is much more complex than observing what others do and reproducing their actions. Because one's knowledge cannot be directly copied from one mind to another, a "teacher" needs to generate a set of observable behaviors for a learner, and the "learner" needs to interpret what those behaviors mean to recover the teacher's knowledge. Even though the learners can only see the concrete actions of the teacher, what is actually shared (by the teacher) and acquired (by the learner) is the message encoded in these actions. Critically, in order for such encoding and decoding to occur successfully, both the teacher and the learner must reason and act with the other person in mind. Thus, the basic human capacity to *understand other minds*—in particular, the ability to interpret others' actions in terms of their underlying mental states—may be the cognitive foundation of distinctively human social learning.

C9.P4 Humans naturally understand others' behaviors in terms of their underlying mental states. When someone is running up the stairs in a train station, we readily reason about the actor's goals (e.g., catching the train), beliefs (e.g., he thinks the train has arrived), and even his expected utilities of successfully achieving his goal (e.g., the benefit of catching the train probably outweighs the cost of running). Even infants can represent and reason about these mental states (Woodward, 1998; Repacholi & Gopnik, 1997; Wellman, Cross, & Watson, 2001) and expect others to act in ways that maximize their utilities (Gergely & Csibra, 2003; Jara-Ettinger et al., 2016; Liu et al., 2017). Indeed, these social-cognitive capacities are critical for many different kinds of social interactions; the goal of this chapter is to focus on a particular class of these interactions (i.e., informative, communicative interactions between a learner and a teacher) and highlight how these capacities can make human social learning so powerful, flexible, and effective.

C9.P5 Critically, successful social learning involves more than the ability to represent others and their internal states. These representations must be incorporated in inferences that may also involve representations of the physical world, such as objects, number, or space. Although the importance of domain-general inferential abilities in early learning has been widely recognized in cognitive development literature (Gopnik, 2012; Schulz, 2012; Xu & Kushnir, 2013), this literature has remained relatively agnostic about the social context in which early learning most frequently occurs, or how social contexts can modulate the inferences learners draw from data. On the other hand, although theories of social learning have proposed learning mechanisms that operate specifically in social contexts (e.g., imitation, natural pedagogy; Meltzoff, 2007; Csibra & Gergely, 2009; Dean et al., 2012; Legare & Nielsen, 2015; Rendell et al., 2010), these domain-specific learning mechanisms are insufficient to fully capture the active nature of early social learning. Rather than passively accepting information by others, young children readily seek out information from others, combine it with their own prior knowledge, and even provide useful information for others to share what they know.

C9.P6 The studies reviewed here are at the intersection of these two perspectives, examining how children draw inferences from available data to support their decisions. While some focus on children's inferences and social evaluations as learners and others look at children as teachers, they are aimed at a larger common goal: To better understand how humans flexibly navigate these roles—recipients and providers of information—depending on the social context. To this end, they take a similar approach to study these interactions: Examining children's inferences and communicative behaviors in experimental settings, specifically in contexts in which one agent provides information, and the other learns from that information. These contexts range from casual interactions between a child and an adult to more explicitly pedagogical ones between a "teacher" and a "learner." In what follows, I first explain why pedagogical interactions can provide a particularly useful window into the cognitive mechanisms that underlie human social learning. Then, across three sections, I review studies that show how young children (1) draw rational inferences from information provided by others, (2) use such information to evaluate others' quality as teachers, and (3) provide information for others as teachers themselves. These studies together show how humans, starting early in life, actively capitalize on their understanding of others in order to learn and to help others learn.

C9.S1 1. Why Study Social Learning in Communicative, Pedagogical Contexts?

C9.P7 Social learning occurs in many forms. Sometimes we learn by simply observing others' actions even when they have no intent to communicate with us. However, social learning can be more effective when someone has a pedagogical goal to inform or teach us. A large body of work has shown that humans are responsive to pedagogy from early in life; for instance, when information about a particular object is provided along with ostensive, communicative cues (e.g., pointing, eye-gaze), even infants show a tendency to interpret that information as generalizable knowledge that applies not only to that particular object but also to a broader set of objects in the same category (Csibra & Gergely, 2009; 2011; Csibra & Shamsudheen, 2015). While such sensitivity might play an especially important role during the earliest months of learning, the studies reviewed here mostly focus on late infancy to preschool years. In these studies, the presence of ostensive cues is often held constant, rather than manipulated, because the goal is not to examine the influence of these cues but to study how children's understanding of other minds can modulate their inferences, evaluations, and communicative decisions in pedagogical contexts where such communicative cues are amply present.[1] Studying children's behaviors as learners and as teachers in these contexts can be particularly interesting and informative for a few reasons.

C9.P8 First, teacher-learner interactions are ubiquitous, especially in young children's lives. The fact that much of early learning unfolds in pedagogical contexts makes these interactions worth studying in their own right. Granted, there is some cultural variation; there are small-scale societies where formal schooling does not exist, and explicit, direct instruction from adults to children is relatively infrequent (MacDonald, 2007; Rogoff, 2003). Nevertheless, the presence of informal social learning opportunities that involve intentional communication of information (e.g., a parent pointing to an object and labeling it, or an older sibling demonstrating how to play a game) may be a universal feature of young children's social environment (Hewlett et al., 2011). As we will see later in this chapter, the ability to interpret information in pedagogical contexts may arise from basic cognitive capacities that are assumed to be universal, rather than from frequent exposure to the cultural practice of teaching.

C9.P9 Second, both learning and teaching require an integration of a host of cognitive faculties, which may all have different origins and developmental

C9.N1 [1] Communicative contexts can vary with respect to the strength of cooperative expectations that interlocutors hold about each other. In this chapter, I consider pedagogical interactions as special cases of communication where such expectations are very strong: the learner expects the teacher to communicate cooperatively, and the teacher expects the learner to interpret the evidence with that expectation (Grice, 1975; see also Goodman & Frank, 2016). Naturally, these interactions often involve the communicator's use of ostensive cues such as eye-gaze or pointing (Csibra & Gergely, 2009; 2011), but the presence of ostensive cues per se does not determine whether an interaction is considered pedagogical.

trajectories. Thus, studying how these complex behaviors emerge early in life can also inform theories of human cognition more generally. In particular, recent computational work on pedagogical reasoning and communication offers formal descriptions of the representations and inferences involved in these interactions (Shafto, Goodman, & Frank, 2012; Shafto, Goodman, & Griffiths, 2014; Goodman & Frank, 2016). This computational framework provides firm grounds for developing theory-driven hypotheses about developmental change in children's inferential and communicative abilities in social contexts (see Figure 9.1). Thus, researchers can go beyond testing qualitative predictions about children's successes and failures; they can test hypotheses about *how* children succeed and *why* they fail in precise, quantitative terms, and generate graded predictions about children's behaviors across tasks and across age groups. Designing model-based experiments also helps researchers be more explicit about how their experimental manipulations are related to the variables in the models.

C9.P10 Third, pedagogical contexts can be used as an effective methodological tool for discovering early competences for social learning and communication. In everyday social interactions, people's goals and epistemic states are often ambiguous or left unspecified. However, when two agents engage in a pedagogical interaction, their goals, motivations, and epistemic states are naturally made clear in the context. When children participate in pedagogical interactions as learners, they expect to learn from a knowledgeable

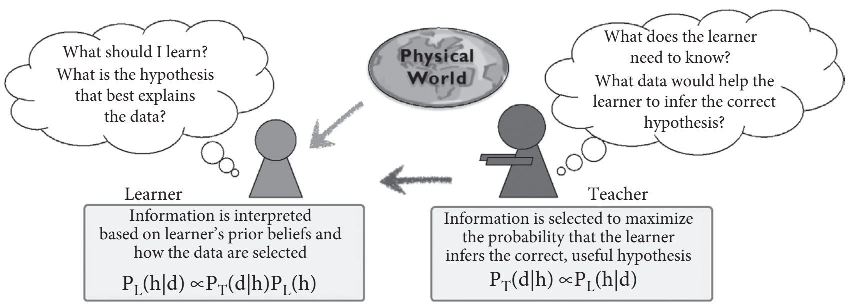

C9.F1 **Figure 9.1** Learning in the social context. Young children learn from their own interactions with the world. However, learning can be more effective when another agent ("teacher") selects and provides data for the learner. Successful learning requires not only the teacher's ability to select the right kind of data, but also the learner's ability to use this sampling process ($P_T(d|h)$) to recover the knowledge that the teacher intends to convey.

person who intends to teach them something new; when children participate as teachers, they expect that learners need their help to order to learn something new, and they are motivated to offer help. Even though these expectations are certainly not unique to pedagogical interactions, they are strengthened and amplified in these contexts, providing a useful test-bed for identifying early inferential and communicative competences that may fail to manifest in other kinds of social interactions .

C9.P11 In sum, young children's inference, evaluation, and communication in pedagogical contexts is both a theoretically interesting area of inquiry and a useful methodological approach that allows us to study how humans use their understanding of other minds to learn from others and help others learn.

C9.S2

2. Rational Inference in the Social Context

C9.S3

2.1. Young Children as Active Interpreters of Information Provided by Others

C9.P12 Observing and imitating others' behaviors is a simple yet useful way to learn from others (Dean et al., 2012; Legare & Nielsen, 2015; Meltzoff, 2007). For instance, if you see someone press a button to activate a toy, you would likely imitate that action and press the button when given a chance to play with the toy. However, if you fail to activate the toy, what would you do next? One study shows that early-developing inferential abilities can support learning even when imitation goes awry; even 16-month-old infants rationally decide what to do next in the face of a failed action, based on the pattern of others' actions and their outcomes (Gweon & Schulz, 2011).

C9.P13 In this study, infants watched two adult experimenters taking turns to press a button to activate a toy. The critical manipulation was the statistical dependency between the experimenters' actions and their outcomes during the demonstration. Some infants saw one experimenter succeed twice and the other experimenter fail twice (i.e., the outcomes of actions covaried with the agent), while others saw each experimenter succeed once and fail once (i.e., outcomes varied independent of the agent). Infants were then given a chance to try the toy, only to fail to activate the toy. Note that all infants saw the same two agents and equal numbers of successful and unsuccessful actions, and readily imitated their actions to press the button on the toy. However, their subsequent actions reflected different inferences about the underlying cause of their failures. When outcomes covary with the agent (i.e., one agent fails twice and the other succeeds twice), the evidence is consistent with the possibility that the failures are due to an incompetent agent rather than a faulty toy. Faced with their own failure, infants in this condition readily attributed the cause to the self (i.e., "I probably can't do it") and approached their parent for help. However, when outcomes vary independently of the agents (i.e., both agents fail once and succeed once), this pattern of evidence suggests that their failures are due to a faulty toy rather than something about the agents. Accordingly, infants in this condition attributed the cause of their own failure to the object ("this toy is faulty") and reached for a different toy. Beyond inferring *what to do now* (i.e., how to activate the toy), infants in this study readily interpreted the statistical patterns of others' actions and outcomes to figure out *what went wrong*, and used this causal attribution to decide *what to do next*.

C9.P14 When agents act on the physical world, their actions and their outcomes generate a pattern of data. These outcomes don't necessarily have to be clear successes and failures, as in Gweon & Schulz, 2011; they could be anything, such as objects that have been drawn from a box full of toys. There is now a body of work on preverbal infants' ability to draw probabilistic inferences based on statistical information (e.g., Aslin, Saffran, & Newport, 1998; Téglás et al., 2011; Xu & Garcia, 2008), and in particular, agents' goal-directed sampling of objects (e.g., Xu & Tenenbaum, 2007; Kushnir, Xu, & Wellman, 2010; Ma & Xu, 2013). For instance, when an agent selects a clearly non-representative sample from a population (e.g., a few blue toys from a box that contains mostly yellow toys and just a few blue toys), infants understand that the agent must have selected them for a reason (i.e., she must like the blue toys). Importantly, when these actions are directed toward someone with a communicative purpose, they can give rise to a rich understanding of w*hat the agent intends to communicate,* given *how the agent selected the information* (see Figure 9.2, left). Imagine someone pulls out such a non-representative sample (a few blue toys from a mostly yellow box) and squeezes each blue toy to show that they all squeak. This sample would be highly improbable under random sampling, and the data strongly suggest that she deliberately *selected* the blue toys. Why would she bother selecting just the blue toys to demonstrate this squeaking property? Presumably because the yellow toys do not squeak!

C9.P15 This intuitive conclusion is, in fact, the result of a sophisticated set of inferences about both the sample and the sampling process by which the

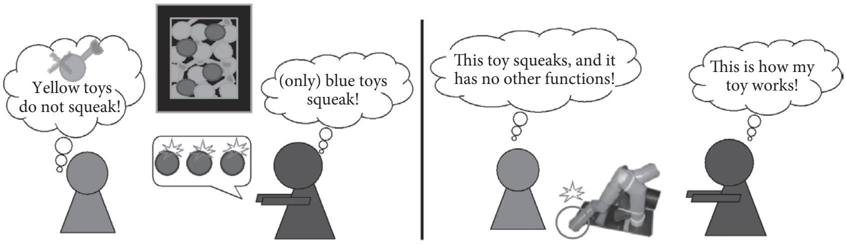

C9.F2 **Figure 9.2** Inferences in communicative contexts. Left: When an adult demonstrates the hidden property of a few blue toys (i.e., pressing them generates squeaky noise), infants use the statistical dependence between the sample and the population to infer the property of undemonstrated objects (yellow toys). When the toys are drawn from a box that contains mostly blue toys (sampling process is ambiguous), infants readily generalize the property to the yellow toys; if the sample is drawn from a box of mostly yellow toys (suggests deliberate sampling of blue toys), infants do not expect the yellow toys to have the same squeaking property (Gweon, Tenenbaum, & Schulz, 2010). Right: These inferences are even stronger in pedagogical contexts where a knowledgeable teacher is expected to help the learner acquire useful knowledge. If the teacher demonstrates just one function of a toy, the learner not only learns about that function but also further infers that the toy has no other functions (if it did, the teacher would have shown them, too).

sample was generated. Even preverbal infants can engage in such reasoning, using the statistical dependence between a population and a sample to guide their inferences about object properties (Gweon, Tenenbaum, & Schulz, 2010). In this study, infants first watched an adult draw three blue toys just like in the preceding scenario. Given a chance to explore the yellow toy (which never squeaked), infants squeezed the yellow toy far less often when the sample was likely to have been selectively drawn (i.e., the box contained mostly yellow toys) than randomly drawn (i.e., the box contained mostly blue toys). Remarkably, the strength of these inferences reflected the strength of the data infants observed; the less likely the samples, the less frequent were their squeezing behaviors. Another notable aspect of these findings is that infants' inferences did not require a lot of observations; they were based on minimal statistical data. Using the relative proportion of blue and yellow toys and just a few exemplars from the population, infants inferred properties of objects that were never sampled or demonstrated. Such sensitivity to sampling processes may emerge from a basic understanding of agents' goal-directed actions and their costs. When agents' behaviors violate the assumption that they should act in ways that minimize the costs (Gergely & Csibra, 2003; Jara-Ettinger et al., 2016), even infants seek rational explanations for such behaviors (i.e., "*why* did she go out of her way to deliberately select these objects?"), and use them to learn about the world.

C9.P16 Pedagogical contexts are an especially interesting and important domain in which deliberate, selective sampling processes affect learning. When a teacher—someone who is assumed to be knowledgeable—communicates information for a learner, it is implied that the teacher is deliberately selecting information that would be most helpful for the learner. Computational models of pedagogical reasoning (Shafto et al., 2012; 2014) have formalized these intuitions, characterizing "what the teacher communicates" and "what the learner learns" as a set of mutually dependent inferences. The teacher's selection of data depends on what the learner would infer given the data (i.e., *pedagogical sampling*), and the learner's inferences from the data depend on how the learner thinks the data were selected by the teacher. These two mutually constraining inferences help the learner draw powerful inferences that go beyond the face value of the evidence (see Figure 9.1).

C9.P17 A series of studies have used demonstrations with causal toys to show that even children as young as two years old expect teachers to engage in pedagogical sampling and interpret the evidence accordingly (Bonawitz et al., 2011; Shneidman et al., 2016; see also Butler & Markman, 2012). In the Pedagogical Condition in these studies, an adult experimenter claims to know everything about a complex-looking novel toy, and says, "Let me show you how it works!" She then demonstrates that pressing a lever on the toy makes an interesting sound. In this context, the observer (i.e., the learner) expects the demonstrator (i.e., the teacher) to be knowledgeable about the toy and to provide the best set of evidence about the toy. Therefore, the learner can safely assume that the teacher would provide information that is not only true of the world but also *sufficient* for the learner to draw the correct inference. As a result, the learner not only learns that the toy has an interesting function (i.e., pressing the lever generates a sound), but also infers that there are *no other functions* to learn; if there were more, the teacher would have demonstrated them, too (Figure 9.2, right). This example illustrates how a learner's expectations about the properties of pedagogically communicated evidence can support inferences that go beyond the teacher's demonstration. Critically, such inferences can be measured through children's spontaneous exploration of the toy; when children are given a chance to explore the toy after observing an experimenter pedagogically demonstrate one of its functions, they show limited exploration of the toy and instead spend the majority of their time playing with the demonstrated function. However, when the teacher's demonstration was non-pedagogical (e.g., apparently interrupted after showing one function, suggesting that she intended to show more, or when the teacher accidentally discovered the function, suggesting that she was ignorant of the toy's functions), children explore the toy more broadly, searching for potential additional functions.

C9.P18 One might wonder whether these results suggest a serious problem in how we educate our young: Does teaching always hinder children's learning by limiting spontaneous exploration and discovery? If the toy actually had multiple functions, learners who learned about one of its functions from a teacher in a pedagogical context would likely have failed to explore and discover the other functions. However, it is important to note that this study deliberately used a teacher who failed to provide all relevant information about the toy, as a way to show how such reduction in exploration can be a natural consequence of the inductive inferences that, ironically, make pedagogical learning so powerful and efficient in most typical pedagogical interactions. Indeed, if the toy actually had only one function, children would have learned this function accurately and efficiently while avoiding the costs of fruitless exploration. If the toy actually had other functions (as was the case in this study), a good teacher in a real pedagogical context would indeed have shown these functions or even have encouraged the learner to explore and learn about them. This suggests that pedagogical learning is beneficial insofar as the teacher is knowledgeable and helpful. One important challenge for real-world "teachers" (e.g., parents, caregivers, educators) is to understand how pedagogical contexts can influence learners' inferences, when its power can backfire, and how pedagogical instructions can be used in ways that maximize its benefit and minimize its unintended side effects.

C9.P19 These studies also raise questions about whether such inferences depend on formal schooling and repeated exposure to pedagogical contexts. Recent cross-cultural data suggest that these inferences are observed even in cultures where children rarely receive direct pedagogical instruction from adults (Shneidman et al., 2016). Toddlers in Yucatec Mayan communities show patterns of exploration in pedagogical versus non-pedagogical contexts similar to those of their same-aged peers in the United States, who presumably receive more frequent instruction from their caregivers. These results provide important empirical support for the idea that children's inferential abilities emerge from a basic understanding of others' knowledge, goals, and their actions, rather than from culturally specific teaching practices.

C9.P20 Together, these studies suggest that young human learners are not just passively absorbing information from others; they actively interpret others' actions and their outcomes based on a rich understanding of other minds. In particular, when someone deliberately selects evidence with a clear communicative goal (e.g., to teach), children draw powerful inferences that go beyond the evidence. Even though children's observations of others' behaviors are often noisy and sparse, these inferences support rapid, robust, yet accurate social learning even from small amounts of data.

C9.S4

2.2. Young Children as Active Evaluators of Others' Informativeness

C9.P21 In pedagogical contexts, strong inferences are licensed because the learners expect their source of information to be someone who is knowledgeable and helpful. However, the power of pedagogical learning can turn into a hazard when such expectations are violated. After all, some people may be less knowledgeable or less helpful than others, and some may even intend to deliberately mislead the learner. Therefore, in order to *learn from others*, it is critical to *learn about others' informativeness* to avoid learning from untrustworthy sources.

C9.P22 Let's go back to the example of a teacher who demonstrates one function of a toy. While this would be considered reasonable and fully informative if the toy had just one function, the same demonstration would be considered underinformative if the toy had three other functions that were left undemonstrated. In fact, by leaving out relevant information, this teacher would be committing a "*sin of omission*," misguiding the learner to falsely believe that the toy has just one function when it in fact has more. Compared to telling a lie or providing false information (Birch, Vauthier, & Bloom, 2008; Koenig & Harris, 2005; Sabbagh & Baldwin, 2001), this is a rather subtle form of misinformation. The teacher's demonstration would lead to inaccurate learning, even though her demonstration is not technically false! Critically, recognizing and evaluating sins of omission requires both prior knowledge about the world and an understanding of pedagogical sampling. Without prior knowledge about the toy's actual functions, one might not even recognize that the teacher omitted some functions of the toy.

Without an understanding of pedagogical sampling, even when an omission is detected, one might not evaluate it as a "sin" due to the lack of understanding that the teacher should have selected a better set of evidence for the learner (i.e., showing all of its functions).

C9.P23 Building on prior work on young children's understanding of pedagogical sampling (Bonawitz et al., 2011), one study tested early school-aged children's evaluations of sins of omission by comparing children's ratings of a fully informative versus an underinformative teacher (Gweon et al., 2014, see Figure 9.3). One group of six- to seven-year-olds were shown a toy with one function, while another group of children were shown a nearly identical toy with *four* functions (including the same function as the one-function toy). All children then rated a teacher who taught a naive learner how the toy works. Critically, the teacher always taught the function that was common to both toys. Even though the teacher's demonstration was true, useful, and identical across conditions, children rated the teacher as less helpful when the toy had other undemonstrated functions (suggesting that the teacher was under-informative, showing only some of its functions) than when the toy had only one function (suggesting the teacher was fully informative). Thus children's evaluations were not only based on whether the teacher's information was *accurate*, but also whether the information was *complete*.

C9.P24 Importantly, such evaluations are not just fleeting impressions; a follow-up study showed that six-year-old children learn about others' informativeness and modulate their subsequent learning (Gweon et al., 2014; see Figure 9.3, right). As in the first study, children saw a teacher who provided underinformative or fully informative demonstration about a toy. However, instead of rating the teacher, children then saw the teacher demonstrate a second toy that they had never seen before. Given a chance to play with the toy, children explored the toy more broadly (in search of potential additional functions) if the teacher had been underinformative in the past than when he had been fully informative. In other words, children who saw the underinformative teacher continued to question his credibility, and engaged in compensatory exploration to see if the toy had additional functions the teacher did not demonstrate. These results suggest that children can interpret the same pedagogical demonstration differently depending on the teacher's past informativeness, and modulate their subsequent inferences and exploration accordingly.

C9.P25 Pedagogical sampling assumptions may be a particularly strong instantiation of Grice's cooperative principles (Grice, 1975): Speakers should

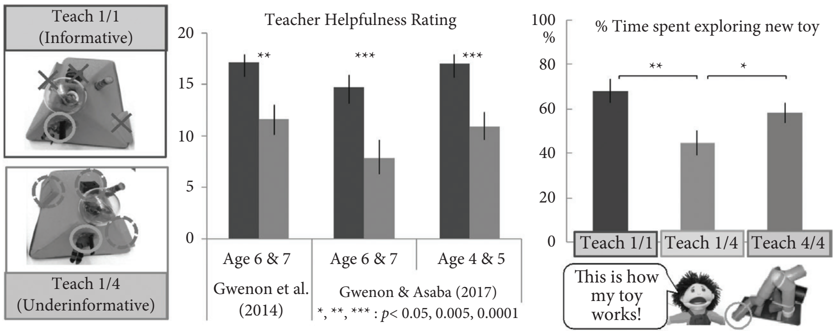

C9.F3 **Figure 9.3** Evaluation of under-informative teachers. Left: Children rated a teacher who taught the only function of a novel toy (Teach 1/1) higher than a teacher who taught the same function while omitting 3 others (Teach 1/4). Four- and five-year-olds show this distinction only when they've seen a clear example of an informative teacher (Gweon & Asaba, 2017). Right: Children's exploration of a new toy is modulated by the teacher's prior informativeness (Gweon et al., 2014).

be truthful (Maxim of Quality) and as informative as required (Maxim of Quantity). In linguistic communication, these cooperative expectations about the interlocutor can support various pragmatic inferences, allowing listeners to go beyond the literal meaning of speakers' utterances (Grice, 1975; Clark, 1996; Goodman & Frank, 2016). For instance, saying that "the boy ate *some* of the cookies" is underinformative when the boy, in fact, ate *all* of the cookies; although it is logically consistent with either possibility, it pragmatically misleads the listener to infer that the boy presumably ate some (and not all). A series of recent studies on pragmatic implicature have shown that even though preschool-aged children typically fail to evaluate such underinformative speakers, they do succeed given enough contextual support (Barner, Brooks, & Bale, 2011; Skordos & Papafragou, 2016). More specifically, children succeed when they understand that the speaker *could have used* the word "all" instead of "some." Children's surprising yet limited competence in pragmatic implicatures supports an interesting hypothesis: If children's evaluations of underinformative teachers are also based on a similar reasoning about what a teacher *could have shown* about a toy, we might see parallel limitations in younger children's evaluations of teachers who communicate via goal-directed actions (Baldwin, Loucks, & Sabbagh, 2008).

C9.P26 A recent study provides empirical support for this hypothesis (Gweon & Asaba, 2017). Using a similar setup as in Gweon et al. (2014), children were shown two teachers who both demonstrated one interesting function of a toy; one was fully informative (because the toy had a single function), and the other was underinformative (because the toy had three additional functions). The critical manipulation was the order in which the teachers were evaluated: Some children saw the fully informative teacher first, while others saw the underinformative teacher first. If seeing and evaluating the fully informative teacher provides a useful reference point for evaluating the underinformative teacher, children younger than six years of age might successfully penalize the underinformative teacher, but only when the fully informative teacher is presented first. The results were consistent with these predictions. While six-year-olds penalized the underinformative teacher regardless of the order in which the teachers were presented (replicating Gweon et al., 2014), preschoolers (four- and five-year-olds) did so only when they rated the fully informative teacher first. Importantly, seeing a teacher who was highly rated for a different reason (e.g., accurately naming an object) did not help; the effect was specific to seeing a fully informative teacher (e.g., providing all relevant functions of a toy). Furthermore, after seeing both teachers, even four-year-olds appropriately preferred the fully informative teacher, providing the earliest evidence for children's ability to detect "sins of omission."

C9.P27 Even though these studies do not include computational models that generate specific predictions for children's ratings, exploratory play, or teacher choice, the results are generally consistent with the computational account of pedagogical reasoning. The failure to provide all relevant information for a learner violates the assumptions of pedagogical sampling (i.e., a knowledgeable teacher selects data to increase the learner's belief in the correct hypothesis; see Shafto et al., 2012; 2014), and supports the inference that the teacher is either ignorant or unhelpful (or both). Note, however, that these studies do not distinguish between these more nuanced attributions; they only show that children readily detect and evaluate sins of omission in pedagogical contexts.

C9.P28 Yet, as adults, we understand that not all omissions are equally blameworthy. For instance, failing to show one of four functions is less egregious than failing to show three of four, and omitting a useless function is better than omitting an important one. If the teacher only shows one function because *he did not know* about the other three, one might be more inclined to pardon the ignorant teacher than a knowledgeable teacher who knowingly omitted these functions. Furthermore, if a teacher shows only one of four functions because the learner already knew the other three, omission could actually be deemed desirable for its efficiency. Thus, evaluation of omission depends on what and how much information was omitted, as well as the knowledge state of the learner and the teacher.

C9.P29 Recent computational work has extended prior models of pedagogical reasoning to incorporate such nuanced evaluations of teachers, and showed that adults' evaluations are consistent with such predictions (Bass et al., 2015; see also Shafto et al., 2012). These models incorporate knowledge states of both the learner and the teacher, and predict graded judgments of teacher quality based on the amount, value, and cost of information taught by the teacher. Consistent with this work, an ongoing series of studies suggest that young children's evaluations of teachers also reflect a rich, sophisticated understanding of what it means to be helpful and informative. For instance, preschool-aged children exonerate underinformative teaching if the teacher couldn't be more informative due to her ignorance (Bass, Bonawitz, & Gweon, 2017). Furthermore, children as young as five years of age understand that an exhaustive demonstration (e.g., showing every button on a multi-button toy even though only a few generate an effect) can be necessary or overinformative depending on the learner's prior knowledge (Gweon, Shafto, & Schulz, 2018). The fact that they penalize redundancy or overinformation (e.g., demonstrating functions that a learner already knows) suggests that children understand the trade-off between informativeness and efficiency. That is, a teacher who provides redundant information has failed to achieve efficiency at the expense of being maximally informative. These results suggest that the ability to reason about the costs and benefits of others' actions (Jara-Ettinger et al., 2016) may underlie these nuanced evaluations of teacher informativeness.

C9.P30 In sum, children don't simply avoid people who are wrong or incomplete, nor do they simply prefer those who provide more evidence. Instead, children reason about what the teacher knows and what the learner needs, and they evaluate whether the teacher selected information in ways that are beneficial for the learner. Collectively, this line of work shows how young children can protect themselves from the potential risks of social learning by learning *about others*; beyond detecting and recognizing misleading informants, they actively infer others' qualities as teachers based on how they select evidence and appropriately modulate their future learning and exploration.

C9.S5

2.3. Young Children as Active Communicators of Information

C9.P31 Children's ability to evaluate others' informativeness raises an interesting possibility: The same cognitive mechanisms that underlie their ability as learners might underlie their abilities as *informative teachers*. Just as children's understanding of pedagogical sampling supports their evaluations of other people's teaching, it may also allow children to conform to these expectations as teachers themselves. Some recent and ongoing work begins to test this hypothesis by placing children in the role of a teacher and studying their ability to communicate information to learners.

C9.P32 Recent work shows that around ages five and six, children readily infer the right amount of information to teach based on the learners' knowledge, goals, and competence, and tailor their demonstrations accordingly (Gweon, Shafto, and Schulz, 2018; Gweon & Schulz, 2019). In one study, given a chance to demonstrate a toy that had 20 identical buttons (only three of which played music, and the rest inert), children demonstrated different numbers of buttons depending on the prior knowledge of the learner (Gweon, Shafto, and Schulz, 2018). When the learner had already seen similar toys, and thus expected the toy to have just a few working buttons, children provided efficient demonstrations by showing just the three working buttons; however, when the learner was naive (thus expecting all buttons to work in the same way), children provided exhaustive demonstrations by pressing all 20 buttons which can help the learner acquire an accurate belief about how the toy works (i.e., only three buttons work). Another study suggests that children selectively provide costly causal demonstrations for a learner only when they were necessary given the learner's goals (e.g., whether the learner wants to learn about the toy vs. observe the toy's effect) or competence (e.g., whether the learner was introduced as an ordinary, naive learner vs. an exceptionally smart learner; Gweon & Schulz, 2019). Along with other recent studies on children's ability to teach (e.g., Ronfard & Corriveau, 2016; Clegg & Legare, 2016; Ronfard, Was, & Harris, 2016; see also Strauss, Calero, & Sigman, 2012 for a review), these results provide compelling evidence that children's teaching is informative, effective, and efficient; they provide as much information as required given the learner's mental states (e.g., goals, knowledge, competence) and avoid providing more than what is necessary.

C9.P33 As in the case of evaluation of teachers, children's ability to balance the trade-off between informativeness and efficiency suggests an early-developing foundation for utility-based social reasoning (Jara-Ettinger et al., 2016). However, the studies reviewed previously do not distinguish whether children taught efficiently simply to reduce their own costs of teaching, rather than having a genuine concern for reducing others' costs. One study shows that children can indeed represent others' cost of learning and use it to decide *what to teach, and what to let learners discover on their own* (Bridgers, Jara-Ettinger, & Gweon, 2016, in press). Although direct instruction allows learners to acquire useful knowledge without the cost of exploration, not everything can or needs to be taught. Therefore, it is important to prioritize teaching what is important yet costly to learn. This study shows that even young children can decide what to teach by reasoning about what others know or want, as well as what is easy or difficult for others to learn. Given a choice of one of two toys to teach to a naive learner (such that the learner has to learn about the other toy on her own), children chose to teach the toy that would (1) increase the learner's expected rewards (i.e., the toy that generates a more enjoyable effect), and (2) decrease the learner's expected costs (i.e., the toy that is harder to figure out by oneself). To provide insights into *how* children make such decisions, this study compared their choices against the predictions of different computational models. Children's responses were consistent with the model that makes utility-maximizing decisions for the learner; it considers the learner's expected utilities of learning about one toy through social learning (i.e., being taught) and learning about the other through self-guided exploration. In addition to finding early competence for making rational teaching decisions, this study addresses an important question about human culture that has been sidestepped in prior literature. To explain how humans have developed cumulative cultural knowledge, it is critical to understand *how humans decide what knowledge is worthy of sharing and teaching*. Otherwise, our cultural knowledge would include lots of useless information, as well as obsolete knowledge that is no longer useful. These results suggest that basic social cognition supports selective transmission of information that prioritizes what is useful for others.

C9.P34 Typically, children are characterized as learners who are on the receiving end of pedagogy. However, the remarkable success of human social learning depends just as much on the teachers' communicative abilities as it does on learners' inferential abilities. The studies reviewed here collectively suggest that the sophistication of adults' intuitions as teachers may have early-developing roots; young children not only learn from others, but also *help others* learn by reasoning about others' goals, competence, and knowledge.

Non-human animals (e.g., meerkats, Thornton & McAuliffe, 2006) may show teaching-like behaviors that have evolved to transmit particular kinds of information (see Kline, 2015, for a review). However, the ability to adjust and tailor teaching based on *an understanding of other minds* is a feature that is distinctive of human social learning.

C9.S6

3. Distinctively Human Social Learning: Looking Back and Looking Ahead

C9.P35 The three lines of work reviewed in this chapter highlight an important feature of human social learning. It is a complex yet well-coordinated process between two individuals who reason about each other's mind. Together they show that, even early in life, humans make rational inferences, evaluations, and pedagogical decisions as learners and as teachers. First, children's inferences reflect more than simple imitation; they actively use information from others to learn about the world, and they reason about what the teacher wants to communicate. Second, their evaluations of teachers are richer than attribution of knowledge for people who are nice, popular, or older (e.g., Lane, Wellman, & Gelman, 2013; Chudek et al., 2012; Wood, Kendal, & Flynn, 2012); they use the properties of information others provide to learn about their informativeness as teachers and use this understanding to guide their future learning. Third, their communicative behaviors as teachers suggest that children's teaching goes beyond providing what is unknown (e.g., Liszkowski, Carpenter, & Tomasello, 2008); children actively consider learners' mental states to teach in ways that increase the benefits of learning while reducing unnecessary costs.

C9.P36 The richness and the power of these abilities may be grounded in at least two ways of reasoning about other minds. First, they require an understanding of how observable evidence (e.g., a teacher's behavior) influences the unobservable states of the mind (e.g., the knowledge state of the learner), a capacity has been traditionally studied in theory of mind research (Gopnik & Astington, 1988; Wellman et al., 2001). Second, they also involve an understanding of how agents' actions are mediated by their goals (Gergely & Csibra, 2003) as well as the costs of achieving these goals (Jara-Ettinger et al., 2015; Liu & Spelke, 2016). The representations that support these understandings—goals, beliefs, desires, as well as effort, competence, and costs—are seamlessly integrated in our everyday social interactions, and manifest in ways we learn from others and help others learn. More specifically, children's inferences, evaluations, and communicative behaviors are constrained by the expectation that other agents will act to pursue rewarding goals while minimizing the costs for achieving them. The idea that humans assume that other agents are utility-maximizers (i.e., the Naïve Utility Calculus; see Jara-Ettinger et al., 2016) provides a useful framework for understanding children's learning and teaching in social contexts. Even early in life, children might engage in an intuitive cost-benefit analysis of information transfer, and make effective decisions both as learners and as teachers.

C9.P37 Pedagogy is a cooperative interaction in which the teacher and the learner act together to achieve a joint goal (Tomasello, 2009). However, it is a rather peculiar kind of cooperation; there is an apparent asymmetry in the division of costs and rewards between the teacher and the learner. The teacher incurs a significant cost (e.g., effort, time) to teach the learner with little direct benefit, while the learner reaps the rewards of the teacher's knowledge while bypassing the costs of learning from directly exploring the world. How can such asymmetric relationships can be sustained, and even become so prevalent, in human societies? One intriguing possibility is that even though teachers are typically considered to incur costs on behalf of the learner (Caro & Hauser, 1992), there may be additional sources of rewards that the teacher reaps from her act of teaching. This could come from the recursive reward a teacher gains through the learner's reward (e.g., the joy of teaching), or other kinds of social reward such as praise, prestige, or status. Furthermore, the asymmetry can come in different forms. For instance, we sometimes provide information not to teach something useful for others but to deliberately *show off* something positive about ourselves. In such cases, we are still incurring costs as a teacher (i.e., provider of information) but only to acquire the benefits for ourselves rather than for the learner. Thus the share of the rewards and costs between the teacher and the learner depends on the context. Future work might investigate how children reason about the cost-reward distributions between agents in different social contexts, and how children modulate their learning or teaching accordingly.

C9.P38 Indeed, there are many open questions about how young children navigate the complexities of learning in social contexts. Nevertheless, the studies reviewed here begin to paint a picture of young humans as active interpreters, evaluators, and communicators of information in social contexts. These inferential and communicative practices allow children to *learn to learn* about the world by making the best use of their rich social surroundings, and *learn* *to help others learn* by selecting the best set of evidence for others. These abilities may be rooted in basic social-cognitive mechanisms for reasoning about the inner qualities of others, such as their goals, beliefs, preferences, and competences. So what makes human social learning so distinctive? The studies in this chapter suggest that the answers may lie in something we often take for granted: Our everyday common-sense reasoning. In particular, the ability to understand other minds—the basic component of human social cognition—may give rise to the power and the richness of human social learning.

C9.S7 References

C9.P39 Aslin, R. N., Saffran, J. R., & Newport, E. L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science*, *9*(4), 321–324.

C9.P40 Baldwin, D., Loucks, J., & Sabbagh, M. (2008). Pragmatics of human action. In Shipley, T. F., & Zacks, J. M. (Eds.), *Understanding Events: From Perception to Action* (pp. 96–129). New York: Oxford University Press. doi:10. 1093/acprof:oso/9780195188370.003.0006.

C9.P41 Barner, D., Brooks, N., & Bale, A. (2011). Accessing the unsaid: The role of scalar alternatives in children's pragmatic inference. *Cognition*, *118*(1), 84–93.

C9.P42 Bass, I., Bonawitz, L. B., & Gweon, H. (2017). Didn't know, or didn't show? Preschoolers consider epistemic state and degree of omission when evaluating teachers. In Gunzelmann, G., Howes, A., Tenbrink, T., & Davelaar, E. (Eds.), *Proceedings of the 39th Annual Conference of the Cognitive Science Society* (pp. 105–110). Austin, TX: Cognitive Science Society.

C9.P43 Bass, I., Hawthorne, D., Goodman, N. D., & Gweon, H. (2015). Not by number alone: The effect of teacher's knowledge and its value in evaluating "sins of omission." In Noelle, D. C., Dale, R., Warlaumont, A. S., Yoshimi, J., Matlock, T., Jennings, C. D., and Maglio, P. P. (Eds.), *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (pp. 166–171). Austin, TX: Cognitive Science Society.

C9.P44 Birch, S. A. J., Vauthier, S. A., & Bloom, P. (2008). Three- and four-year-olds spontaneously use others' past performance to guide their learning. *Cognition*, *107*(3), 1018–1034. http://doi.org/10.1016/j.cognition.2007.12.008.

C9.P45 Bonawitz, E., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. Cognition, 120(3), 322–330. http://doi.org/10.1016/j.cognition.2010.10.001.

C9.P46 Bridgers, S., Jara-Ettinger, J., & Gweon, H. (2016). Children consider others' expected costs and rewards when deciding what to teach. In Papafragou, A., Grodner, D., Mirman, D., & Trueswell, J. C. (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 710–715). Austin, TX: Cognitive Science Society.

C9.P47 Bridgers, S., Jara-Ettinger, J., & Gweon, H. (in press). Young children consider the expected utility of others' learning to decide what to teach. *Nature Human Behaviour*.

C9.P48 Butler, L. P., & Markman, E. M. (2012). Preschoolers use intentional and pedagogical cues to guide inductive inferences and exploration. *Child Development*, *83*(4), 1416–1428. http://doi.org/10.1111/j.1467-8624.2012.01775.x.

C9.P49 Carey, S. (2011). *The Origin of Concepts*. New York: Oxford University Press.

C9.P50 Caro, T., & Hauser, M. (1992). Teaching in nonhuman animals. *Quarterly Review of Biology*, *67*(151), 74.

C9.P51 Chudek, M., Heller, S., Birch, S., & Behavior, J. H. E. A. H. (2012). Prestige-biased cultural learning: Bystander's differential attention to potential models influences children's learning. *Evolution and Human Behavior*, *33*, 46–56. http://doi.org/10.1016/j.evolhumbehav.2011.05.005.

C9.P52 Clark, H. (1996) *Using Language*. New York: Cambridge University Press.

C9.P53 Clegg, J. M., & Legare, C. H. (2016). Instrumental and conventional interpretations of behavior are associated with distinct outcomes in early childhood. *Child Development*, *87*(2), 527–542. http://doi.org/10.1111/cdev.12472.

C9.P54 Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, *13*(4), 148–153.

C9.P55 Csibra, G., & Gergely, G. (2011). Natural pedagogy as evolutionary adaptation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*(1567), 1149–1157. http://doi.org/10.1098/rstb.2010.0319.

C9.P56 Csibra, G., & Shamsudheen, R. (2015). Nonverbal generics: Human infants interpret objects as symbols of object kinds. *Annual Review of Psychology*, *66*(1), 689–710.

C9.P57 Dean, L. G., Kendal, R. L., Schapiro, S. J., Thierry, B., & Laland, K. N. (2012). Identification of the social and cognitive processes underlying human cumulative culture. *Science*, *335*(6072), 1114–1118. http://doi.org/10.1126/science.1213969.

C9.P58 Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naive theory of rational action. *Trends in Cognitive Sciences*, *7*(7), 287–292.

C9.P59 Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, *20*(11), 818–829. http://doi.org/10.1016/j.tics.2016.08.005.

C9.P60 Gopnik, A. (2012). Scientific thinking in young children: Theoretical advances, empirical research, and policy implications. *Science*, *337*(6102), 1623–1627. http://doi.org/10.1126/science.1223416.

C9.P61 Gopnik, A., & Astington, J. (1988). Children's understanding of representational change and its relation to the understanding of false belief and the appearance-reality distinction. *Child Development*, *59*(1), 26–37.

C9.P62 Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: Causal models, Bayesian learning mechanisms, and the theory theory. *Psychological Bulletin*, *138*(6), 1085–1108. http://doi.org/10.1037/a0028044.

C9.P63 Grice, H. P. (1975). Logic and conversation. In Cole, P., & Morgan, J. (Eds.), *Syntax and Semantics*: vol. 3, *Speech Acts* (pp. 41–58). New York: Academic Press.

C9.P64 Gweon, H., & Asaba, M. (2017). Order matters: Children's evaluation of underinformative teachers depends on context. *Child Development.* doi: 10.1111/cdev.12825.

C9.P65 Gweon, H., Pelton, H., Konopka, J. A., & Schulz, L. E. (2014). Sins of omission: Children selectively explore when teachers are under-informative. *Cognition*, *132*(3), 335–341. http://doi.org/10.1016/j.cognition.2014.04.013.

C9.P66 Gweon, H., & Schulz, L. E. (2011). 16-month-olds rationally infer causes of failed actions. *Science*, *332*(6037), 1524.

C9.P67 Gweon, H., & Schulz, L. E. (2019). From exploration to instruction: Children learn from exploration and tailor their demonstrations to observers' goals and competence. *Child Development*, *90*(1), e148–e164. doi:10.1111/cdev.13059.

C9.P68 Gweon, H., Shafto, P. & Schulz, L.E. (2018). Development of children's sensitivity to over-informativeness in learning and teaching. *Developmental Psychology*, *54*(11), 2113–2125. doi: http://dx.doi.org/10.1037/dev00005802113

C9.P69 Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences*, *107*(20), 9066–9071. http://doi.org/10.1073/pnas.1003095107.

C9.P70 Hewlett, B. S., Fouts, H. N., Boyette, A. H., & Hewlett, B. L. (2011). Social learning among Congo Basin hunter-gatherers. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*(1567), 1168–1178. http://doi.org/10.1016/j.socscimed.2009.05.016.

C9.P71 Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The Naïve Utility Calculus: Computational principles underlying commonsense psychology. *Trends in Cognitive Sciences*, *20*(8), 589–604. http://doi.org/10.1016/j.tics.2016.05.011.

C9.P72 Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, *76*(6), 1261–1277. http://doi.org/10.1111/j.1467-8624.2005.00849.x.

C9.P73 Kline, M. A. (2015). How to learn about teaching: An evolutionary framework for the study of teaching behavior in humans and other animals. *Behavioral and Brain Sciences*, *38*, e31. http://doi.org/10.1017/S0140525X14000090.

C9.P74 Kushnir, T., Xu, F., & Wellman, H. (2010). Young children use statistical sampling to infer the preferences of other people. Psychological Science, 21(8), 1134.

C9.P75 Lane, J. D., Wellman, H. M., & Gelman, S. A. (2013). Informants' traits weigh heavily in young children's trust in testimony and in their epistemic inferences. *Child Development*, *84*(4), 1253–1268. http://doi.org/10.1111/cdev.12029.

C9.P76 Legare, C. H., & Nielsen, M. (2015). Imitation and innovation: The dual engines of cultural learning. *Trends in Cognitive Sciences*, *19*(11), 688–699. http://doi.org/10.1016/j.tics.2015.08.005.

C9.P77 Liszkowski, U., Carpenter, M., & Tomasello, M. (2008). Twelve-month-olds communicate helpfully and appropriately for knowledgeable and ignorant partners. *Cognition*, *108*(3), 732–739. http://doi.org/10.1016/j.cognition.2008.06.013.

C9.P78 Liu, S., Ullman, T. D., Tenenbaum, J. B., & Spelke, E. S. (2017). Ten-month-old infants infer the value of goals from the costs of actions. *Science*, *358*(6366), 1038–1041.

C9.P79 Ma, L., & Xu, F. (2013). Preverbal infants infer intentional agents from the perception of regularity. *Developmental Psychology*, *49*(7), 1330–1337. http://doi.org/10.1037/a0029620.

C9.P80 MacDonald, K. (2007). Cross-cultural comparison of learning in human hunting. *Human Nature*, *18*(4), 386–402.

C9.P81 Meltzoff, A. N. (2007). "'Like me": A foundation for social cognition. *Developmental science*, *10*(1), 126–134.

C9.P82 Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., et al. (2010). Why copy others? Insights from the social learning strategies tournament. *Science*, *328*(5975), 208–213. http://doi.org/10.1126/science.1184719.

C9.P83 Repacholi, B., & Gopnik, A. (1997). Early reasoning about desires: Evidence from 14-and 18-month-olds. *Developmental Psychology*, *33*(1), 12–20.

C9.P84 Rogoff, B. (2003). *The Cultural Nature of Human Development*. Oxford: Oxford University Press.

C9.P85 Ronfard, S., & Corriveau, K. H. (2016). Teaching and preschoolers' ability to infer knowledge from mistakes. *Journal of Experimental Child Psychology*, *150*, 87–98. http://doi.org/10.1016/j.jecp.2016.05.006.

C9.P86 Ronfard, S., Was, A. M., & Harris, P. L. (2016). Children teach methods they could not discover for themselves. *Journal of Experimental Child Psychology*, *142*, 107–117. http://doi.org/10.1016/j.jecp.2015.09.032.

C9.P87 Schulz, L. (2012). The origins of inquiry: Inductive inference and exploration in early childhood. *Trends in Cognitive Sciences*, *16*(7), 382–389. http://doi.org/10.1016/j.tics.2012.06.004.

C9.P88 Sabbagh, M. A., & Baldwin, D. A. (2001). Learning words from knowledgeable versus ignorant speakers: Links between preschoolers' theory of mind and semantic development. *Child Development*, *72*(4), 1054–1070.

C9.P89 Shafto, P., Goodman, N. D., & Frank, M. C. (2012). Learning from others: The consequences of psychological reasoning for human learning. *Perspectives on Psychological Science*, *7*(4), 341–351. http://doi.org/10.1177/1745691612448481.

C9.P90 Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology*, *71*, 55–89. http://doi.org/10.1016/j.cogpsych.2013.12.004.

C9.P91 Shafto, P., Gweon, H., Fargen, C., & Schulz, L. (2012). Enough is enough: Inductive sufficiency guides learners' ratings of informant helpfulness. In Miyake, N., Peebles, D., & Cooper, R. P. (Eds.), *Proceedings of the 34th Annual Conference of the Cognitive Science Society* (pp. 977–982). Austin, TX: Cognitive Science Society.

C9.P92 Shneidman, L., Gweon, H., Schulz, L. E., & Woodward, A. L. (2016). Learning from others and spontaneous exploration: A cross-cultural investigation. *Child Development*, *87*(3), 723–735. http://doi.org/10.1111/cdev.12502.

C9.P93 Skordos, D., & Papafragou, A. (2016). Children's derivation of scalar implicatures: Alternatives and relevance. *Cognition*, *153*, 6–18. http://doi.org/10.1016/j.cognition.2016.04.006.

C9.P94 Strauss, S., Calero, C. I., & Sigman, M. (2014). Trends in neuroscience and education. *Trends in Neuroscience and Education*, *3*(2), 38–43. http://doi.org/10.1016/j.tine.2014.05.001.

C9.P95 Téglás, E., Vul, E., Girotto, V., Gonzalez, M., Tenenbaum, J. B., & Bonatti, L. L. (2011). Pure reasoning in 12-month-old infants as probabilistic inference. *Science*, *332*(6033), 1054–1059. http://doi.org/10.1126/science.1196404.

C9.P96 Thornton, A., & McAuliffe, K. (2006). Teaching in wild meerkats. *Science, 313*(5784), 227–229. http://doi.org/10.1126/science.1128727.

C9.P97 Tomasello, M. (2009). *Why We Cooperate*. Cambridge, MA: MIT Press.

C9.P98 Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development*, *72*(3), 655–684.

C9.P99 Wood, L. A., Kendal, R. L., & Flynn, E. G. (2012). Context-dependent model-based biases in cultural transmission: Children's imitation is affected by model age over model knowledge state. *Evolution and Human Behavior*, *33*(4), 387–394. http://doi.org/10.1016/j.evolhumbehav.2011.11.010.

C9.P100 Xu, F., & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences*, *105*(13), 5012–5015. http://doi.org/0704450105 [pii] 10.1073/pnas.0704450105.

C9.P101 Xu, F., & Kushnir, T. (2013). Infants are rational constructivist learners. *Current Directions in Psychological Science*, *22*(1), 28–32. http://doi.org/10.1177/0963721412469396.

C9.P102 Xu, F., & Tenenbaum, J. B. (2007). Sensitivity to sampling in Bayesian word learning. *Developmental Science*, *10*(3), 288–297. http://doi.org/10.1111/j.1467-7687.2007.00590.x.