Communication as Influence
Preschool-Aged Children Can Infer What Speakers Know Based on How They Influence Others

Honors Thesis
Presented to the faculty of the Department of Symbolic Systems
In partial fulfillment of the requirements for the degree of
Bachelor of Science with Honors

Robert Z. Sparks
Primary Advisor: Hyowon Gweon, Ph.D., Stanford Psychology
Second Reader: Aaron Chuey, Ph.D. Candidate, Stanford Psychology

Department of Symbolic Systems
Stanford University

June 2022

To the Directors of the Program on Symbolic Systems:
I certify that I have read the thesis of Robert Sparks in its final form for submission and have found it to be satisfactory for the degree of Bachelor of Science with Honors.
Signed Electronically.
June 3rd, 2022
Hyowon Gweon
Stanford Psychology

To the Directors of the Program on Symbolic Systems:
I certify that I have read the thesis of Robert Sparks in its final form for submission and have found it to be satisfactory for the degree of Bachelor of Science with Honors.
Signed Electronically.
June 3rd, 2022
Aaron Chuey
Stanford Psychology

Preschool-Aged Children Can Infer What Speakers Know Based on How They Influence Others

Robert Z. Sparks (bsparks@stanford.edu)
Department of Symbolic Systems, Stanford University

Abstract

How do we know what others know? Prior work has examined how children use evidence about isolated agents, like their perceptual access and actions, to infer what they know. However, humans are rarely fully isolated; instead, we are often surrounded by others whom we interact with, influence, and are influenced by. In these contexts, we can use a speaker's communication and the way it causes a listener to behave to infer what that speaker knows - even if we do not know the specific content of what was communicated. The present studies investigated how preschool-aged children use two pieces of evidence about listeners to reason about what speakers know: changes in the outcomes of a listener's actions following communication (Study 1) and changes in a listener's actions themselves following communication (Study 2). In both studies, children observed two scenarios where a listener failed to activate a toy before succeeding. In Study 1, children observed a speaker produce nonsense language towards a listener after they failed but before they succeeded to activate a toy, as well as another speaker who spoke to a listener prior to initial failure. In Study 2, children observed a speaker communicate with a listener before a distinct change in action, followed by success, as well as another speaker who communicated with a listener resulting in no distinct change in action, followed by success. When asked which speaker knows how to make the toy work, 5 year-olds chose the speaker who appeared to cause the listener to succeed (Study 1) or change their action (Study 2). These results suggest that preschool-aged children are sensitive to the way speakers influence others via communication and can use evidence of that influence to infer what speakers know. More broadly, these studies highlight children's ability to reason about the knowledge of one agent (a speaker) based primarily on evidence about another agent (a listener).

Introduction

Our capacity to reason about what others know is essential for navigating the social world we inhabit. Indeed, understanding what others know enables us to better predict others' actions (Jara-Ettinger, Gweon, Tenenbaum, Schulz, 2015), cooperate (Thomas, et. al., 2014; Grueneisen, et. al., 2015), compete (Schmidt, Hardecker, Tomasello, 2016), and especially interact pedagogically (Gweon, Shafto, Schulz, 2018). As learners, reasoning about what others know allows us to identify knowledgeable informants (Einav, Robinson, 2011) and evaluate others' testimony (Sperber, et. al., 2010). As teachers, we can tailor what we teach to a learner depending on their prior knowledge, providing optimal instruction (Bridgers, Jara-Ettinger, Gweon, 2020).

Impressively, we can even make epistemic inferences by observing pedagogical exchanges themselves. Consider two students, Kate and Betty, sitting next to each other in their science class. When the instructor asks Kate, "What are the only two non-silvery metals?", Kate pauses and does not provide an answer; it is apparent that she does not know. However, Betty then leans into Kate's ear and whispers something. Suddenly, Kate shouts, "Gold and Copper!". It seems obvious that, in addition to Kate ultimately knowing the answer, Betty probably knows the answer too. By observing a pedagogical exchange between the two, we were able to infer what both Kate, the learner, and Betty, the teacher, knows.

A growing body of literature has shown that our ability to reason about others' knowledge develops remarkably early. In one study, 6-month-olds attributed a preference for an object to an agent when an alternative object was visible, expecting the agent to reach for the preferred object when presented with both objects (Luo, Johnson, 2009). However, infants did not attribute a preference when the agent could not see the alternative object; when both objects became visible, infants expected the agent to reach for each object equally. This suggests that even infants as young as 6-months expect visible, but not non-visible objects, to influence an agent's goal-directed actions. Likewise, 13-month-olds expect an agent to successfully find an object when the agent had seen its location, but not when its location was obscured (Surian, Caldi, Sperber, 2007), suggesting that 13-month-olds can predict where agents will search for objects - and whether their searches will be successful - based on their perceptual access. Although these findings do not definitively entail knowledge representation, they demonstrate that infants understand how perceptual access mediates others' behaviors. As adults, we understand that agents acquire knowledge from what they perceive, and act based on what they know and desire.

As children continue to develop, they begin to generate more abstract representations of others' epistemic states. By around four years, children can infer that agents subject to deception have false beliefs and can discern between false beliefs and ignorance (Hogrefe, Wimmer, Perner, 1986). By around the same age, young children also form expectations about how agents revise their beliefs. For example, 5 year-olds can infer how an agent will revise their beliefs about the contents of a population of objects based on observing new perceptual evidence via a random or selective sample of the population (Magid, et. al., 2018).

Young children's understanding of the relationship between others' knowledge and how they act also develops substantially in the preschool years. For example, preschoolers infer that agents whose actions avoid costs and obtain greater reward are more knowledgeable than agents who incur costs and fewer rewards (Jara-Ettinger et. al., 2017). Young children can also make the reverse inference, predicting that knowledgeable agents will avoid costs and obtain higher reward compared to ignorant agents. Likewise, preschoolers also possess an understanding of knowledge that goes beyond mere notions of accuracy. 5 year-olds selectively attribute prior knowledge to agents who accurately predict the outcome of an event, but not to agents who merely describe the outcome of an event after it has occurred (Aboody, Huey, Jara-Ettinger, 2018). Preschoolers also attribute knowledge to agents as a means of explaining their actions; by 5 years children will attribute additional knowledge to ignorant agents whose prior experience cannot explain their actions nor their success (Aboody, et. al., 2019). This suggests that preschool-aged children expect the actions of both ignorant and knowledgeable agents to be guided by their prior knowledge more broadly. By the elementary school years, children can even use an agent's prior knowledge to make inferences about other pieces of knowledge they might possess, attributing greater domain knowledge to those who know how objects or animals work, but not greater knowledge about unrelated categories (Chuey, Lockhart, Sheskin, Keil, 2020).

Together, these studies demonstrate that by early childhood, children can infer what others know using their perceptual access, actions, and even prior knowledge. However, consider the science class example from earlier. We were able to infer what Betty knows despite possessing none of this information. Betty's perceptual access was unclear (she was not looking at a relevant book, for example), her prior knowledge was unknown, and while her actions (whispering in particular) might seem informative, nothing about whispering in isolation suggests that she knows anything about metals and their properties. Instead, the way Betty appeared to influence Kate's knowledge and her subsequent actions by communicating with her seemed to suggest that Betty herself was knowledgeable.

Some prior work has explored infant and toddlers' ability to reason about others' minds in communicative contexts. 12-month-olds expect a listener to successfully perform an action when a communicator who had successfully performed the action in the past spoke a novel word to the listener, but not when the communicator coughed towards the listener (Vouloumanos, Onishi, Pogue, 2012). Additionally, 6-month-olds expect a listener to grasp the preferred toy of a communicator only when the communicator produced speech, but not coughed, towards the listener (Vouloumanos, Martin, Onishi, 2014). These findings demonstrate that by as young as 6-months, even infants understand that speech, but not other vocalizations such as coughing, can influence a listener's actions based on the abilities or preferences of the speaker.

Other work has examined infants' ability to reason about how communication can correct false-beliefs. 18-month-olds expect a listener to revise their false-belief about the location of a ball that was moved from a box to a cup when told "the ball is in the cup!" by a speaker who viewed the ball's true location, but not when the speaker says "I like the cup" (Song, et. al., 2008). In a similar study, 18-month-old toddlers expected an agent to revise their false-belief when told "the ball is in the cup!", but not when told "the ball and the cup." Together, these studies show that toddlers understand how specific utterances can correct a listener's false beliefs depending on its semantic content. More specifically, toddlers expect explicitly relevant speech, but not irrelevant speech, to correct a listener's false belief.

Young children's understanding of the relationship between communication and knowledge is also evident in their own communicative behavior. Before they can speak, infants selectively point to an object when an adult did not see where it was dropped (Liszkowski, Carpenter, Tomasello, 2008) and spontaneously point to the new location of a moved object when an adult attempts to retrieve an object from its former location (Knudsen, Liszkowski, 2012). Once they can speak, toddlers spontaneously communicate the name and location of a hidden object to someone who did not witness the object being hidden (O'Neill, 1996).

These findings point to an early-emerging understanding that communication, particularly speech, can alter the epistemic states of listeners. However, in the science class example, we also inferred what Betty, the speaker, knew. How did we make this inference? Kate was initially ignorant but later became knowledgeable. In particular, Betty's whisper appeared to cause Kate to become knowledgeable, leading us to infer that Betty herself was knowledgeable. That is, we used Betty's influence over Kate, mediated by Betty's communication, to infer that Betty was knowledgeable. While whispering in isolation may not be enough to infer Betty's knowledge, her whisper's apparent influence over Kate's knowledge and actions was.

The bulk of work investigating children's epistemic inferences has focused on isolated agents who acquire knowledge from their environment. Even when cases of communication were considered, the emphasis was placed on the recipient of communication and its epistemic outcome (i.e., a listener coming to know something) rather than on the source of communication and the epistemic states that necessitated those outcomes (i.e., the speaker possessing the relevant knowledge that was communicated to the listener). Yet, as humans, we typically find ourselves in social environments where agents influence and are influenced by others. While understanding others in these environments certainly demands more complicated inferences about their minds, social environments also enable more complex mental state inferences by providing evidence that agents in isolation simply do not, such as a speaker's communicative influence over a listener. The current studies ask whether children can use a speaker's influence over a listener to infer what that speaker knows.

Across two studies, children viewed two videos in which a listener unsuccessfully attempts to activate a toy before succeeding. At some point in each video, a communicator communicates with the listener and then leaves. As with the case of Kate and Betty, the semantic content of the communication was unavailable to children (agents used a nonsense utterance identified as a novel language). In Study 1, we manipulated the temporal structure of the events such that in one video, the speaker appeared to cause the listener to succeed, while in the other video, the speaker appeared to cause the listener to fail. After watching both scenarios, we asked children which one of the speakers knew how the toy worked. We predicted that children between 3 to 5 years would prefer the speaker who appeared to cause the listener to succeed rather than fail.

Study 1

Methods

Participants Seventy-two children between ages 3 and 5 (24 of each age) were preregistered to participate in the study. 76 total children participated (Age: 3;0.0 - 5;11.9; Mage=4;7.6), but 4 children were excluded based on preregistered exclusion criteria (2 for experimenter error, 1 for parental interference, and 1 for technical difficulties). Children participated online via Zoom Meeting video conferencing software.

Materials The study was presented to children using Microsoft PowerPoint and the screen-sharing feature on Zoom Meeting. The presentation included still photos and two prerecorded video clips of two agents and a causal toy. All agents were monster puppets, introduced as "wubs." A green wub appeared in both videos as a listener, and the blue and orange wub each appeared in only one of the two videos as the communicator (color counterbalanced). The causal toy was constructed from foam boards and colored tape, and topped with two round buttons.

Procedure Children were first introduced to four different colored monster puppets called "wubs" (two green, one blue, one orange) and told the wubs did not speak English, but spoke a different language called 'Jabberwocky' that only the wubs know. Children were then introduced to a novel box-shaped causal toy and had their attention directed to the two buttons on the top of the toy. Children were told "do you see the two buttons on top? To make this toy go, you have to press both buttons on top at the same time. If you press just one of the buttons, the toy doesn't work!". The two green wubs were then shown on either side of the toy with question marks above their heads. The experimenter then said, "The green wubs have never seen the toy before, so they don't know how it works." The blue wub and orange wub then appeared on either side of the toy, and children were told ".. but luckily, one of these two wubs here knows how the toy works! But the other one doesn't know how it works. I'm going to show you two videos about these wubs, and your job is to figure out which of these wubs know how the toy works. It could be the blue wub, or it could be the orange wub."

Children were then shown two videos, each depicting a different condition: an Effective Communicator and an Ineffective Communicator condition (order counterbalanced). In the Effective Communicator condition, the green wub approached the toy, pressed each button individually, vocalized a grunt in frustration, and then pressed each button individually again. A second wub (the communicator, blue/orange counterbalanced) then approached the green wub, vocalized a short nonsense phrase, and left. The green wub then pressed both buttons at the same time, causing the toy to make a ringing sound. The green wub then raises both hands and vocalizes an expression of joy.

After watching the first video, the experimenter said, "so this is what happened to these wubs! Let's see what happens in the other video." and played the next condition. The Ineffective Communicator condition was identical to the Effective Communicator condition, but the communicator approached the green wub before the green wub pressed any buttons. The green wub approached the toy, the communicator immediately approached the green wub and vocalized a short nonsense phrase, and left. The green wub then pressed each button individually, vocalized a grunt in frustration, and then pressed each button individually again. The green wub then presses both buttons at the same time, causing the toy to make a ringing sound. The green wub then raises both hands and vocalizes an expression of joy.

After watching the second video, the experimenter said, "so that is what happened with these wubs! Now I have a question for you," and children were shown a slide with the toy and the blue and orange wubs on either side of the toy. Children were then asked, "remember the blue wub and the orange wub?" One of them knows how the toy works, but the other one doesn't know how the toy works. Which wub do you think knows how the toy works? The blue wub or the orange wub?"

Results and Discussion

In-line with our preregistered analysis plan, we conducted three binomial tests for each age group (3 year-olds, 4 year-olds, and 5 year-olds) testing the proportion of children who chose the effective communicator. 3 year-olds (12/24, Binomial Test: $p = 1.1612$, 95% CI = [.296, .704]) and 4 year-olds (13/24, Binomial Test: $p = .8388$, 95% CI = [.338, .756]) did not show a significant preference for either communicator. However, 5 year-olds (18/24, Binomial Test: $p = .023$, 95% CI = [.573, .927]) preferred the effective communicator significantly greater than chance.

These results provide evidence that by around 5 years, children can reason about a communicator's knowledge using the influence they exert over another agent. Although children had no explicit access to the content of the communicators' speech, and the actions of both the communicators and listeners were identical across the two videos, they were still able to infer which communicator was knowledgeable using the listener's behavior following a communicative act. In particular, children were able to make this inference based on whether

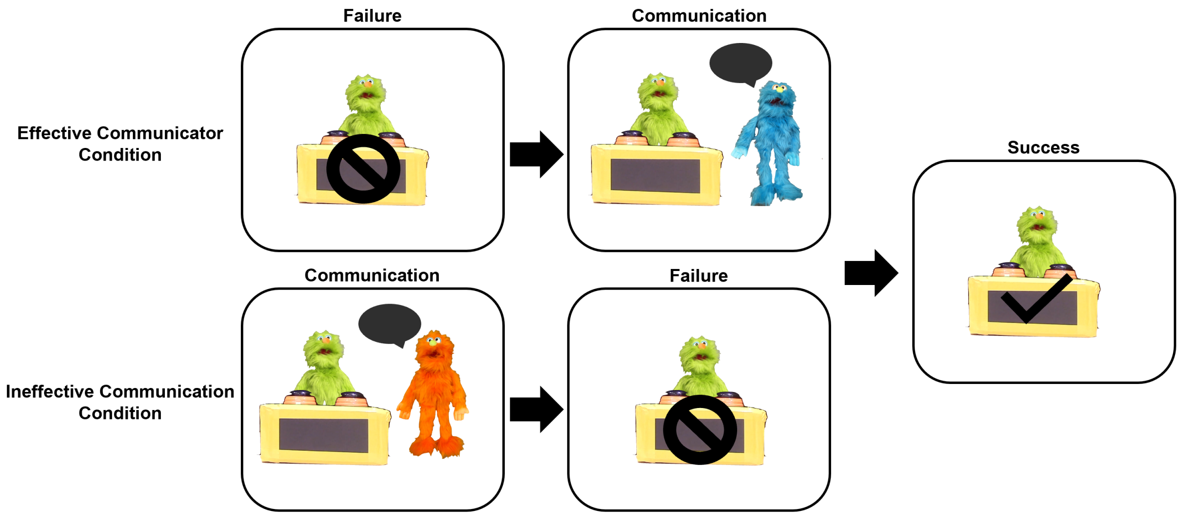

Figure 1: Study 1 Effective and Ineffective Communicator conditions. Across both conditions, the green wub fails to activate a two button toy twice before ultimately succeeding. However, the timing of the speaker's communication differs. In the Effective Communicator condition (top), the green wub fails to activate the toy twice before a communicator enters, speaks, and then leaves. Following communication, the green wub successfully activates the toy. In the Ineffective Communicator condition (bottom), a communicator enters, speaks, and then leaves before the green wub touches the toy. Following communication, the green fails to activate the toy twice before succeeding.

the communication appeared to bring about a successful or unsuccessful outcome. These results build on prior work that has shown that young children expect relevant communication to alter the epistemic states of listeners (Song, et. al., 2008; Jin, et. al., 2019) by showing that young children also attribute knowledge to the source of that epistemic change.

Although 5 year-olds demonstrated competence in this task, 3 and 4 year-olds both performed at chance. One explanation is that children this young simply cannot make this kind of epistemic inference - that while young children can make epistemic inferences in single-agent contexts or about listeners in communicative ones, reasoning about what a speaker knows following a brief communicative exchange may simply be too complex. Alternatively, young children may be capable of making these kinds of epistemic inferences, but the current task itself may have been too cognitively demanding. The current task has high memory demands with copious exposition; it also requires children to track very similar events across two videos while simultaneously attributing specific events to the appropriate agent at test. Children may also struggle to interpret the Ineffective Communicator condition. The ambiguity of the agent's nonsense utterance combined with the eventual success of the listener generates a wide domain of potential utterances that the communicator could have produced as well as potential interpretations of those utterances by the listener. This would suggest that post-communication success/failure does not offer a strong enough cue to young children to discern between a naive communicator and a knowledgeable one.

Study 2 attempts to address these latter two concerns by making adjustments to the study presentation and design in an attempt to reduce the cognitive load for children and by exploring another cue for communicative influence: changes in a listener's actions following communication.

Study 2

The prior study considered the result of communicative influence on the outcomes of a listener's actions (post communication success versus post communication failure), with agents otherwise performing the same actions across the two videos. In addition to outcomes, communicative influence can also alter the listener's actions themselves (from pressing one button to pressing both buttons). Can children also use a listener's costly change in action following communication to infer what a speaker knows?

Consider yourself and a friend, Cora, who are in a library together. Cora goes off on her own in search of a particular book. As you look off into the distance, you can see Cora scanning rows of books, pulling books out, and putting them back. Suddenly, a stranger approaches Cora, and you see them have a conversation. Cora then walks to the opposite end of the library and approaches a distant shelf. One might infer that the stranger knew where the book was located, but as was the case with Kate and Betty in their science class, we once again have very few direct cues that explicitly point to the stranger's knowledge. In fact, in this case, we do not even know if Cora managed to find the book she was looking for. However, Cora's sudden change in behavior in favor of a more costly action following her conversation with the stranger suggests that the stranger themselves might know

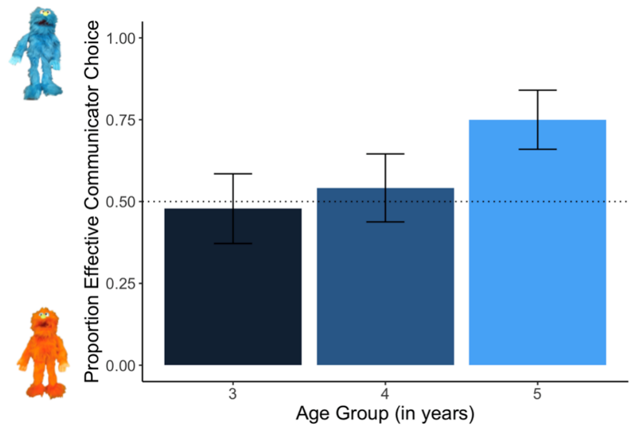

Figure 2: Results from Study 1, depicting children's preference for the effective communicator across different ages; error bars indicate standard error (.5 ± .1043, .543 ± .1039, .75 ± .0903), and the dotted line indicates chance performance (50%).

where the book is located, or at least that they have a strong belief about it.

Prior work has shown that preschool-aged children will ask an agent to complete a more difficult task with low probability of random success rather than an easy task with guaranteed success in order to assess the agent's knowledge (Aboody, Denison, Jara-Ettinger, 2021), suggesting that children can jointly use the properties of a task to assess an agent's epistemic states. Further, preschool-aged children attribute more knowledge to an agent who expresses satisfaction and commitment to a choice compared to an agent who expresses dissatisfaction and revises a choice (Jara-Ettinger, et. al., 2017). Together, these findings suggest that children recognize spontaneous, costly changes in action as being potentially indicative of a change in epistemic state, and that costly changes in action are made in service of goal-directed behaviors guided by those epistemic state changes. Therefore, children might also be able to attribute such changes in a listener's actions following communication to a speaker's influence.

The present study asks whether young children can use observable changes in a listener's actions to infer what a speaker knows. Similar to Study 1, Children viewed two scenarios where a listener unsuccessfully attempted to activate a toy before ultimately succeeding. In one scenario, following the speaker's communication, the listener performs an unlikely, costly action before activating the toy. In the other scenario, the listener shows no observable change in action following the speaker's communication, before ultimately activating the toy. After the children watched both scenarios, we asked the children which speaker knew how to make the toy ring (activate the toy). We predicted that by around 5, children would prefer the speaker whose communication led to a costly change in action.

Methods

Participants Twenty-Four children between the ages of 4 and 5 participated in a pilot study. Children participated online via Zoom Meeting video conferencing software. 24 children participated (Age: 4;1.0 - 5;11.9; Mage=5;0.83) and one child was tested but excluded due to exclusion criteria related to technical difficulties.

Materials The study was presented to children using Microsoft PowerPoint and the screen-sharing feature on Zoom Meeting. The presentation included still photos, animations and two pre-recorded video clips of two agents and a causal toy. All agents were monster puppets, introduced as "wubs" wearing different colored vests: a green wub with a red vest, a green wub with a gray vest, a blue wub with an orange vest, and an orange wub with a blue vest. A green wub appeared in both videos as a listener, and the blue and orange wub each appeared in only one of the two videos as the communicator (color counterbalanced). The causal toy was constructed from foam boards and colored tape, and topped with 8 round buttons (6 green buttons, 1 white button, and 1 pink button). Both the white button and pink button were equidistant from the center of the toy, separated by 2 green buttons, and both with 2 green buttons to the other side.

Procedure Children were first introduced to the four different colored wubs (two green, one blue, one orange) and were told that the wubs did not speak English, but spoke a different language called 'Jabberwocky' that only the wubs know. Children were then introduced to a novel box-shaped causal toy and told, "Have you seen a toy like this before? This toy rings when you press the correct button". A ringing noise was then produced to show the child what sound the toy makes. The experimenter then directed the attention of the child to the pink button on the toy and said, "Do you see this pink button? To make this toy go, you have to press the pink button! If you press any of the other buttons, the toy doesn't do anything." Children were then asked, "Can you tell me how to make this toy ring?".

The two green wubs then appeared, without the toy, with question marks above their heads. The experimenter then

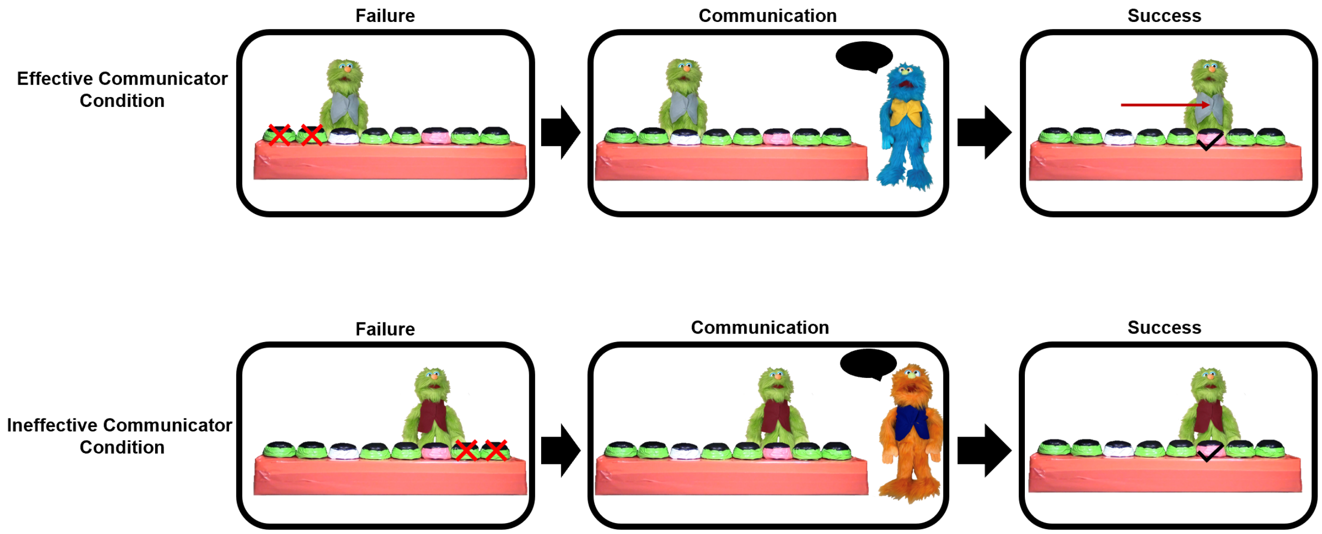

Figure 3: Study 2 Effective and Ineffective Communicator conditions. In both conditions, a green wub fails to activate an eight button toy twice before ultimately succeeding on their third attempt. However, the green wub's actions following communication differ by condition. In the Effective Communicator condition (top), the green wub starts on the left-side of the toy and fails to activate the toy twice. The green wub then moves in front of the white button (inert) and raises its hand as if about to press it. A communicator then interrupts the green wub, speaks a novel language, and then leaves. Following communication, the green wub moves down the length of the toy and successfully activates the toy by pressing the pink button. In the Ineffective Communicator condition (bottom), the green wub starts on the right-side of the toy and fails to activate the toy twice. The green wub then moves in front of the pink button (efficacious) and raises its hand as if about to press it. A communicator then interrupts the green wub, speaks, and then leaves. Following communication, the green wub does not move and immediately activates the toy by pressing the pink button.

said, "The green wubs have never seen the toy before, so they don't know how to make the toy ring, but they really want to make the toy ring!" The blue wub and orange wub then appeared on either side of the green wubs, and children were told "But look! See this blue wub and orange wub? One of these two wubs has played with the toy before and knows that the pink button makes the toy ring, but I'm not sure which wub knows. It could be the blue wub or it could be the orange wub!" Children were then told, "But the videos you're going to watch have some clues. I really need your help, can you watch the videos very carefully and help me figure out whether the blue wub or orange wub really knows that you have to press the pink button to make the toy ring?"

The first video was then introduced to the child. Children were told, "Ok, so this first video is going to be about the blue wub and this green wub" and children were then shown an animation with the orange wub and other green wub exiting, leaving only the blue and green wub. Children were then shown two videos, each depicting a different condition: one video showing an Effective Communicator condition, and the other showing an Ineffective Communicator condition (order counterbalanced). Still photos of the agents in each video were displayed above the video to help remind children whom the video was about.

In the Effective Communicator condition, the green wub (listener) approached the toy, looked down at the toy and vocalized an expression of curiosity, and then moved to the side of the toy farthest from the pink button. The green wub then looked down, pressed the first green button, and vocalized a grunt in frustration. The green wub then moved closer to the next green button, looked down, pressed the next green button, and vocalized a grunt in frustration. The green wub then moved closer to the white button, looked down, and raised its hand as if to press the white button. A second wub (communicator, blue/orange counterbalanced) then approached the green wub. The green wub lowered its hand, the communicator vocalized a short nonsense phrase, and the communicator left. The green wub then moved down the toy and positioned itself in front of the pink button. The green wub looked down and pressed the pink button, causing the toy to make a ringing sound. The green wub then vocalized an expression of joy.

After watching the first video, the experimenter said, "So that was a video about this green wub and the blue wub, now we're going to watch a video about a different green wub and the orange wub." Children were then shown an animation where green wub and blue wub exit, and the other green wub and orange wub enter the slide. The experimenter then asked ,"Ready to watch?" and played the next video.

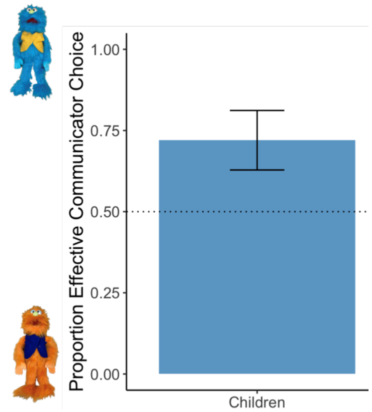

Figure 4: Children's preference for the effective communicator in Study 2. The error bar indicates standard error (.708 ± .0948), and the dotted line indicates chance performance (50%).

The Ineffective Communicator condition was identical to the Effective Communicator condition, but the listener begins by moving towards the side of the toy closest to the pink button. The green wub (listener) approached the toy, looked down at the toy and vocalized an expression of curiosity, and then moved to the side of the toy closest to the pink button. The green wub then looked down, pressed the first green button, and vocalized a grunt in frustration. The green wub then moved closer to the next green button, looked down, pressed the next green button, and vocalized a grunt in frustration. The green wub then moved closer to the pink button, looked down, and raised its hand as if to press the pink button. A communicator then approached the green wub, the green wub lowered its hand, the communicator vocalized a short nonsense phrase, and then the communicator left. The green wub then paused, looked back down at the pink button, and pressed the pink button, causing the toy to make a ringing sound. The green wub then vocalized an expression of joy.

After watching the second video, the experimenter said, "so that is what happened with these wubs! Now I have a question for you," and children were shown a slide with only the blue wub and orange wub. Children were then told, "Remember, only one of these wubs here knew that the pink button makes the toy ring - now we need to figure out whether it was the blue wub or orange wub that really knew how to make the toy ring!" Children were then asked, "Which wub knew that the pink button makes the toy ring, the blue wub or the orange wub?"

Results

A large proportion of children (17/24) chose the effective communicator when asked which wub knew that the pink button makes the ring (2-Tailed Binomial Test: $p = .064$, 95% CI = [.522, .894]). This provides some preliminary evidence that by around 5, children appear to use the change in a listener's actions following communication to infer what a speaker knows. Although, more data is necessary to draw a reliable conclusion.

General Discussion

In Study 1, children viewed two videos involving a speaker and a listener with the actions of both sets of speakers and listeners being identical. In both videos, the speaker merely entered the scene, spoke a nonsense language, and left. Listeners in both videos failed to activate a toy twice before ultimately succeeding in activating the toy. However, in one video the speaker communicated to the listener directly before the listener succeeded, while in the other video, the speaker appeared before the listener initially failed. Remarkably, children as young as 5 appear to selectively attribute knowledge to the speaker from the former video, suggesting they can infer the knowledge of a speaker based on whether communication caused a listener to produce a successful or unsuccessful outcome.

In Study 2, there were no temporal differences in the videos - the speakers in both videos appeared directly before the listener succeeded. However, the listener's actions following communication differed across both videos. In one video, the listener performed a costly action following communication, moving away from the button they were about to press and down the length of the toy in order to successfully activate it. In the other video, the listener exhibited no observable change in action following communication, pressing the same button they were about to press before communication in order to successfully activate the toy. Preliminary results show that by around 5, children reliably attribute knowledge about activating the toy to the speaker in the former video, suggesting that around this age, children can infer the knowledge of a speaker based on whether communication caused a listener to spontaneously produce a more costly action. Results from Study 2 are further supported by a more recent pilot with 10 5 year-olds, in which a large proportion (8 of the 10) of participants chose the effective communicator when asked which wub knew that the pink button makes the ring. Continuing data collection may offer stronger evidence of a preference for the effective communicator amongst 5 year-olds.

Together, these results provide initial evidence that children as young as 5 years can reason about a communicator's knowledge using the influence they appear to have over another agent's actions and their outcomes. While some work has investigated how children reason about the knowledge of other agents in communicative contexts, prior studies have focused on listeners and how their behavior and knowledge are affected by communication. In contrast, the present studies break new ground by examining how children reason about the knowledge of communicators. This is particularly impressive given that children had no relevant information about the speaker's perceptual access, their toy-related actions, their prior knowledge, or even the semantic content of their speech. Instead, children were able to use a listener's actions and their outcomes to infer what a speaker, a completely different agent, knows. More broadly, these findings suggest that children can use evidence about an agent (a listener) to make inferences about another agent (a communicator).

Young children's ability to reason about the knowledge of other agents via their communicative influence suggests a number of interesting implications. Building on work that supports children's ability to identify and assess knowledgeable informants (Einav, Robinson, 2011; Bass, Bonawitz, Gweon, 2017), the current studies suggest that young children might use a speaker's influence over others to reason about the speaker's competence, particularly in pedagogical contexts (Bass, et. al., 2022; Sierksma, Shutts, 2021). In turn, this can help children identify effective teachers and helpers, as well as seek help themselves. Similarly, this may also help children identify ineffective teachers and helpers and guide children in accessing the need for pedagogical intervention. Future work could examine how children integrate epistemic inferences about communicators and their information seeking behaviors more broadly.

These findings also inform our understanding of young children's causal reasoning. Children's ability to attribute knowledge to influential communicators rests on their more basic capacity to identify causal agents within a social system (i.e., agents that cause other agents to engage in desirable actions or achieve desirable outcomes). This expands on prior work showing that children can identify causal relationships within mechanical-object systems (Lucas, et. al., 2014; Goddu, Gopnik, 2020) and that children infer mental state and physical state attributions as causes of an agent's behavior (Seiver, et. al., 2013). Identifying causal relationships within a social system entails understanding that people can exert causal influence over others. The transmission of knowledge functions as a mechanism by which communication alters others' epistemic states and actions; the extent to which young children understand this causal relationship warrants future study.

A number of limitations also limit the scope of our current and potential conclusions. 3 and 4 year-olds were at chance in Study 1, suggesting that younger children may not be capable of making this kind of inference about a speaker's knowledge. Although we suspect that children around this age are capable of reasoning about a speaker's knowledge, children's general difficulty with engaging with and understanding epistemic language and reasoning paired with the complexity of the communicative interaction may make explicitly reasoning about a speaker's knowledge too difficult for younger children. Given this possibility, further exploration into whether children are sensitive to the influence that a speaker has over a listener as evidence for producing preference for an agent may implicate younger children's understanding of communicative influence.

Alternatively, the current methods may underestimate children's capacity to reason about communicators' knowledge. The complexity of the current methods may simply be too demanding to elicit the necessary epistemic inference from younger children. Relatively verbose exposition, the presence of up to four agents, and the need to remember the content of two very similar videos might contribute to children's failure on the task. Both studies also only offer children a binary choice of which communicator is knowledgeable, rather than a richer gradient of knowledge that may more accurately reflect children's (and adults') own intuitions. Future work should attempt to weaken these constraints to both better demonstrate children's underlying competence, and measure richer inferences that are more akin to children's real world experience. Interestingly, a number of findings with infants and toddlers (e.g.,Tauzin, Gergely, 2018) suggests that an understanding of knowledge and communication may develop quite early. Given the right methodology, it is possible that competence in the present inferences can be probed with much younger children.

Conclusion

A growing body of work has explored young children's ability to reason about the knowledge of others, focusing primarily on children's ability to reason about the knowledge of isolated agents using their perceptual access, actions, and prior knowledge. Yet, very little work has investigated how children reason about others' knowledge in social contexts where agents' minds influence and are influenced by others. The current studies investigated preschool-aged children's ability to infer the knowledge of communicators via two kinds of evidence: changes in the outcome of a listener's action following communication (Study 1) and changes in the listener's actions themselves following communication (Study 2). Five year-olds consistently attributed knowledge to communicators who appeared to influence a listener to successfully activate a toy. Results from these studies suggest that by preschool years, children can attribute knowledge to communicators as sources of epistemic change using their causal influence over listeners, even in the absence of traditionally studied evidence such as perceptual access, speech content, actions, and prior knowledge.

Acknowledgments

I would like to express my deepest appreciation to Hyowon Gweon for allowing me the opportunity to work with her lab and for supporting me professionally and academically as my Major Advisor, my Thesis Advisor, and as my PI over the past two years. She is deserving of all the recognition and celebration of her generosity, patience, kindness, and enduring commitment to supporting the mission and members of her lab.

I would also like to extend my deepest gratitude to Aaron Chuey for allowing me to contribute to his projects and for being the most amazing mentor that anyone could ever ask for. His genuine and enthusiastic support for my professional and personal well-being and the extensive lengths to which he has been willing to collaborate with me has meant the world over these past 2 years.

To my Mother, my Grandparents, my Uncle Kevin, my Aunt Cathy, my Uncle Paul, my Aunt Cindy, and to Amy Larimer & Drew Kracfik, I would like to express my most sincere love and appreciation for being pillars in my life and for their unconditional belief in and celebration of me.

Finally, I'd like to thank all of my loved ones who have supported me through my time at Stanford. The outpouring of love, support, and the unconditional nurturing of my soul and well-being during my time at Stanford have sustained me and inspired my continuous growth and pursuit of joy.

References

Aboody, R., Denison, S., & Jara-Ettinger, J. (2021). Children consider the probability of random success when evaluating knowledge. *Proceedings of the 43rd Annual Conference of the Cognitive Science Society*.

Aboody, R., Flowers, M., Zhou, C., & Jara-Ettinger, J. (2019). Ignorance = doing what is reasonable: Children expect ignorant agents to act based on prior knowledge. *Proceedings of the 41st Annual Conference of the Cognitive Science Society*.

Aboody, R., Huey, H., & Jara-Ettinger, J. (2018). Success does not imply knowledge: Preschoolers believe that accurate predictions reveal prior knowledge, but accurate observations do not. *Proceedings of the 40th Annual Conference of the Cognitive Science Society*.

Bass, I., Bonawaitz, E., Hawthorne-Madell, D., Vong, W. K., Goodman, N. D., & Gweon, H. (2022). The effects of information utility and teachers' knowledge on evaluations of under-informative pedagogy across development. *Cognition*, *222*.

Bass, I., Bonawitz, E., & Gweon, H. (2017). Didn't know, or didn't show? preschoolers consider epistemic state and degree of omission when evaluating teachers. *Cognitive Science*.

Bridgers, S., Jara-Ettinger, J., & Gweon, H. (2020). Young children consider the expected utility of others' learning to decide what to teach. *Nature Human Behavior*, *4*.

Chuey, A., Lockhart, K., Sheskin, M., & Keil, F. (2020). Children and adults selectively generalize mechanistic knowledge. *Cognition*, *199*.

Einav, S., & Robinson, E. J. (2011). When being right is not enough: Four-year-olds distinguish knowledgeable informants from merely accurate informants. *Psychological Science*, *22*.

Goddu, M. K., & Gopnik, A. (2020). Learning what to change: Young children use "difference-making" to identify causally relevant variables. *Developmental Psychology*, *56*.

Grueneisen, S., Wyman, E., & Tomasello, M. (2015). "i know you don't know i know. . . " children use second-order false-belief reasoning for peer coordination. *Child Development*, *86*.

Gweon, H., Shafto, P., & Schulz, L. E. (2018). Development of children's sensitivity to over-informativeness in learning and teaching. *Developmental Psychology*, *54*.

Hogrefe, G. J., Wimmer, H., & Perner, J. (1986). Ignorance versus false belief: A developmental lag in attribution of epistemic states. *Child Developement*, *57*.

Jara-Ettinger, J., Floyd, S., Tenenbaum, J. B., & Schulz, L. E. (2017). Children understand that agents maximize expected utilities. *Journal of Experimental Psychology*, *146*.

Jara-Ettinger, J., Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2015). Children's understanding of the costs and rewards underlying rational action. *Cognition*, *140*.

Jin, K. S., Song, M., Kim, Y. J., Lee, H., Lee, Y., Cha, M., & Song, H. J. (2019). Fourteen- to eighteen-month-old infants use explicit linguistic information to update an agent's false belief. *Frontiers in Psychology*, *10*.

Kim, E. Y., & Song, H. J. (2015). Six-month-olds actively predict others' goal-directed actions. *Cognitive Development*, *33*.

Knudsen, B., & Liszkowski, U. (2012). Eighteen-and 24-month-old infants correct others in anticipation of action mistakes. *Developmental Science*, *15*.

Liszkowski, U., Carpenter, M., & Tomasello, M. (2008). Twelve-month-olds communicate helpfully and appropriately for knowledgeable and ignorant partners. *Cognition*, *108*.

Lucas, C. G., Bridgers, S., Griffiths, T. L., & Gopnik, A. (2014). When children are better (or at least more open-minded) learners than adults: developmental differences in learning the forms of causal relationships. *Cognition*, *131*.

Luo, Y., & Johnson, S. C. (2009). Recognizing the role of perception in action at 6 months. *Developmental science*, *12*.

Magid, R. W., Yan, P., Siegel, M. H., Tenenbaum, J. B., & Schulz, L. E. (2018). Changing minds: Children's inferences about third party belief revision. *Developmental Science*, *21*.

O'Neill, D. K. (1996). Two-year-old children's sensitivity to a parent's knowledge state when making requests. *Child Development*, *67*.

Schmidt, M. F. H., Hardecker, S., & Tomasello, M. (2016). Preschoolers understand the normativity of cooperatively structured competition. *Journal of Experimental Child Psychology*, *143*.

Seiver, E., Gopnik, A., & Goodman, N. D. (2013). Did she jump because she was the big sister or because the trampoline was safe? causal inference and the development of social attribution. *Child Development*, *84*.

Song, H. J., Onishi, K. H., Baillargeon, R., & Fisher, C. (2008). Can an agent's false belief be corrected by an appropriate communication? psychological reasoning in 18-month-old infants. *Cognition*, *109*.

Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G., & Wilson, D. (2010). Epistemic vigilance. *Mind Language*, *25*.

Surian, L., Caldi, S., & Sperber, D. (2007). Attribution of beliefs by 13-month-old infants. *Psychological Science*, *18*.

Tauzin, T., & Gergely, G. (2018). Communicative mind-reading in preverbal infants. *Scientific Reports*, *8*.

Thomas, K. A., DeScioli, P., Haque, O. S., & Pinker, S. (2014). The psychology of coordination and common knowledge. *Journal of Personality and Social Psychology*, *107*.

Vouloumanos, A., Martin, A., & Onishi, K. H. (2014). Do 6- month-olds understand that speech can communicate? *Developmental Science*, *17*.

Vouloumanos, A., Onishi, K. H., & Pogue, A. (2012). Twelve-month-old infants recognize that speech can communicate unobservable intentions. *Proceedings of the National Academy of Sciences*, *109*.