

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Ignorance = doing what is reasonable: Children expect ignorant agents to act based on prior knowledge

Permalink

<https://escholarship.org/uc/item/1ct9s9vc>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 41(0)

Authors

Aboody, Rosie

Flowers, Madison

Zhou, Caiqin

et al.

Publication Date

2019

Peer reviewed

Ignorance = doing what is reasonable: Children expect ignorant agents to act based on prior knowledge

Rosie Aboody¹ (rosie.aboody@yale.edu), Madison Flowers¹ (madison.flowers@yale.edu), Caiqin Zhou² (czhou@wellesley.edu), & Julian Jara-Ettinger¹ (julian.jara-ettinger@yale.edu)

¹ Department of Psychology, Yale University. New Haven, CT 06520 USA.

² Department of Psychology, Wellesley College. Wellesley, MA 02481 USA.

Abstract

When deciding how to act in new situations, we expect agents to draw on relevant prior experiences. This expectation underlies many of our mental-state inferences, allowing us to infer agents' prior knowledge from their current actions. Do children share this expectation, and use it to infer others' epistemic states? In Experiment 1, we find that five- and six-year-olds (but not four-year-olds) attribute additional knowledge to agents whose prior experiences cannot explain their success. In Experiment 2, we find that six-year-olds (but not younger children) also attribute greater knowledge to agents whose prior experience cannot explain their failure. We show that by age five or six, children expect ignorant agents' beliefs (and therefore their actions) to be guided by their prior knowledge. This work adds to a growing body of research suggesting that, while infants can represent mental states, the ability to infer mental states continues to develop throughout early childhood.

Keywords: Ignorance; Knowledge; Social Cognition; Theory of Mind

Introduction

To discuss someone's ambitions, frustrations, or disappointments is to talk about a mind that works much like our own, except that we cannot see it or know what it knows. Yet, we make surprisingly accurate inferences about what others think or want, just by watching how they act. For example, if your friend gives you her keys but later rummages in her bag upon reaching the car, you might infer that she forgot you have them. If she doesn't slow for a pedestrian at a crosswalk, you'd probably assume she didn't see them. And if she suddenly takes a detour, you might suspect she knows something you don't (perhaps the usual route is under construction).

The ability to infer other people's thoughts and desires from their behavior involves building a working model of how their mental states relate to their actions. The foundations of this capacity, called a Theory of Mind (Dennett, 1987; Gopnik & Wellman, 1992), are in place and at work early in infancy (Woodward, 1998; Liu, Ullman, Tenenbaum, & Spelke, 2017) but continue to mature throughout early childhood (Wellman, Cross, & Watson, 2001), and well into adolescence (Richardson, Lisandrelli, Riobueno-Naylor & Saxe, 2017).

Within Theory of Mind, our ability to reason about other people's beliefs—what they know, what they don't, and what they think they know—is particularly slow to develop. While infants can represent other people's beliefs (Onishi &

Baillargeon, 2005), knowledge (Surian, Caldi & Sperber, 2007), and ignorance (O'Neill, 1996), children do not use these representations explicitly until several years later (Bartsch & Wellman, 1995; Wellman, et al., 2001).

As adults, we understand that other people's past experiences shape their current beliefs, and that these beliefs guide their actions. If, for example, your friend starts their car by inserting and turning a key, you can reasonably predict they will try the same the first time they drive yours. And you'd expect this even if you know your car works differently (for example, starting when a button is pushed in proximity to the key fob).

This expectation not only allows us to predict how others will act: it also allows us to infer what they know by observing how they act. In the example above, if your friend defied your expectations by immediately locating the button that starts your car, you might wonder if they had some prior experience you didn't know about (perhaps they've driven other cars like yours before). Such reasoning may seem intuitive, but how exactly do we predict what actions agents are likely to take in new situations? Prior research suggests that adults solve this problem by integrating over agents' uncertainty (Baker, Jara-Ettinger, Saxe, & Tenenbaum, 2017). For instance, when we reason about an agent who does not know whether a car starts via a key or a button, we consider what they would do in each situation, and we expect them to choose a plan weighted by their confidence.

While effective, these types of inferences are computationally complex. They require considering multiple possible worlds (at least implicitly), and deciding what an agent would do in each. Perhaps unsurprisingly, children's expectations for how ignorant agents are likely to act appear to rely on simpler strategies. Children sometimes equate being ignorant with getting things wrong (Ruffman, 1996; Saxe, 2005); although, in other contexts, their intuitions reverse (Friedman & Petrashek, 2009; German & Leslie, 2001).

While expecting ignorant agents to fail may support accurate inferences and useful predictions, such strategies are limited. Even ignorant agents can make reasonable guesses based on past experience. For instance, even if you've only used PC's, you probably have some idea of what you'd try if you had to turn on a Mac. And ignorant agents can always get lucky, succeeding by chance.

Do children understand how previous experiences affect agents' future actions? And do they leverage this expectation to infer what an agent knows based on what she does? In the current work, we investigate these questions with four- to six-year-olds. The ability to explicitly and flexibly represent beliefs emerges in the mid-preschool years (e.g., Rubio-Fernández, 2019; Wellman et al., 2001). Therefore, if children have expectations about the relation between ignorance and action, we might expect them to emerge in this age range.

In two experiments, participants watched two puppets learn how to activate a novel toy. Each puppet later attempted to activate a different (but outwardly identical) toy. One agent's actions were consistent with their prior experience, while the other agent's actions were inconsistent with their prior experience. In Experiment 1, both agents succeeded in activating a toy. If children expect agents to act based on their prior knowledge, they should judge that the inconsistent agent (whose actions cannot be explained by their experience with the initial toy) must have had additional knowledge. We find that five- and six-year-olds (but not four-year-olds) attribute additional prior knowledge to this agent.

To control for the possibility that children attribute knowledge to agents who teach them something new, in Experiment 2, children learned how a toy worked, and then watched two agents fail to activate this toy. Children again judged that the inconsistent agent (whose action couldn't be explained by his experience with the initial toy) had *greater* additional knowledge. These results suggest that by age five, children expect ignorant agents to act according to their prior knowledge, and further, that children leverage this expectation to infer what others know from what they do. All experiments' procedures, predictions, exclusion criteria, and analyses were pre-registered.

Experiment 1

In Experiment 1, children watched two puppets learn how to activate a novel toy. Next, each puppet was given the chance to activate a different toy (always outwardly identical to the original). One puppet stated that his chosen toy worked the same as the original, and pressed the same button he had seen activate the original toy. The other puppet stated that his chosen toy worked differently to the original, and pressed a different button. Both puppets succeeded in activating their chosen toy. Children were then asked which of the two agents already knew how the toys worked.

If children expect ignorant agents to behave in accordance with their prior beliefs, then they should judge that the agent who acted inconsistently with their prior experience is more likely to be knowledgeable. But if children attribute epistemic states by relying on a rule of thumb (e.g., expecting ignorant agents to be wrong), or have no representation of what it means to be ignorant, then children should have no preference for either agent.

Method

Participants 72 four-, five- and six-year-olds (mean age: 5.46 years, range: 4.05 – 6.99 years; $n = 24$ participants per age group) were recruited at a local children's museum. 22 participants were excluded from the analyses and replaced because: they did not pass the pre-registered inclusion questions ($n = 9$), due to experimenter error ($n = 5$), interruptions from other children ($n = 3$), because the participant did not answer the test question within 30s ($n = 2$), distraction ($n = 1$), interference with the procedure ($n = 1$), or due to developmental delays ($n = 1$).

Stimuli Stimuli consisted of two male puppets, and three novel toys. These toys were externally identical machines, each covered in black construction paper and measuring approximately 5 x 3 x 2.75 in. Toys had three buttons on top: a red button in the middle, and two black buttons flanking the red one (see Figure 1).

Although they all looked the same, the toys worked in different ways. The first toy (called the "training" toy) activated and played music only when the central red button was pressed. Of the remaining toys, the "consistent" one worked the same way. However, the "inconsistent" toy worked differently: only pressing the black button to the participant's far left made it activate. For clarity, we refer to this button as the "correct" black button, and the other as the "incorrect" black button (since it did not activate the toy).

Procedure First, participants were familiarized with the training toy (which turned on when the central red button was pressed). Participants learned that the red button made the toy go, but that the black buttons did nothing. They were then given a chance to press all of the buttons themselves. Next, participants were introduced to two puppets. The experimenter explained that she was going to show the puppets how the toy worked, and told the puppets that while the red button made the toy go, the black buttons did not do anything. Upon the experimenter's request, the puppets pressed the red button together.

Next, the remaining toys were placed on the table (one on either side of the training toy). The experimenter explained that one of the puppets had snuck out from under the table and played with all the toys, and discovered which buttons made the toys play music. The other puppet had stayed underneath the table, and hadn't seen anything. The child's task was to help figure out which puppet had snuck out and played with all the toys.

Each puppet was questioned individually, while the other agent was placed under the table. During his turn, each puppet was asked: "Can you show us how to make one of these toys go?" To make the relation between agents' actions and their experience with the initial toy more explicit, each agent explained himself as he acted. One puppet chose the consistent toy, saying, "Hmm. Well, the red button made this [original] toy go, so the red button makes this toy go too," pointing to the two relevant buttons as he spoke. Finally he pressed the red button, successfully

activating the toy. The other puppet chose the inconsistent toy, saying, “Hmm. Well, the red button made this [original] toy go, but this black button makes this toy go,” pointing to the two relevant buttons as he spoke. Finally he pressed the correct black button, successfully activating the toy.

After each puppet demonstrated one of the toys, the experimenter asked the test question: “[Child name], remember how I told you at the beginning of the game that only one of my friends snuck out from underneath the table, and played with all the toys? Can you tell me, which one of my friends snuck out and played with all the toys?” Participants were then asked to explain their answer. The memory check questions (pre-registered as inclusion questions) were asked last, with subjects asked to match each puppet to the toy he had demonstrated: “[Child name], can you remind me, which friend showed us how to make this toy go [both puppets point to a toy]? And which friend showed us how to make this toy go [both puppets point to the other toy]?”

Puppets always demonstrated the toy they were standing closest to. This was to avoid pragmatic concerns that could arise if puppets undertook a cost to demonstrate a particular toy. Therefore, the puppet on the experimenter’s left hand demonstrated the leftmost toy, and vice versa. The identity of the puppet whose turn was first, and the toy this agent acted on was always counterbalanced. Additionally, the side each puppet was presented on (left/right) was randomized.




Figure 1: Procedure of both experiments. In Experiment 1 both puppets succeeded in activating the toy. In Experiment 2, both failed. Crucially, one agent’s actions were always consistent with his prior experience (pressing the red

button); the other agent’s were not (he pressed one of the black buttons).

Results and Discussion

Two coders who were not involved in data collection determined exclusions. The first coder determined whether the experiment had been run correctly, blind to children’s final answers. The second coder coded only children’s answers, unaware of each puppet’s role (that is, whether he demonstrated the consistent or inconsistent toy). 22 participants were excluded and replaced (see Participants).

Overall, of 58.3% of children judged that the agent who pressed the black button (and acted inconsistently with his prior experience) was more likely to have had additional knowledge. This proportion is not reliably different from chance (42 of 72; 95% CI: 47.2 – 69.4). However, a logistic regression predicting performance based on age revealed a significant age difference ($\beta = 0.87, p = .006$). While only 37.5% of four-year-olds judged that the agent who activated the inconsistent toy had prior knowledge (9 of 24; 95% CI: 16.67 – 58.33), 66.6% of five-year-olds (16 of 24; 95% CI: 50 – 87.5) and 70.8% of six-year-olds (17 of 24; 95% CI: 54.17 – 87.5) selected this agent. And consistent with five- and six-year-olds’ success, a logistic regression predicting performance based on age also predicts that children will be more likely to answer the test question correctly (as opposed to incorrectly) by 5.04 years of age.

These results suggest that children do not simply expect ignorant agents to act successfully or unsuccessfully. Rather, by age five, children seem to expect ignorant agents to act reasonably, applying their prior knowledge in novel situations. This is consistent with prior findings that children do not think ignorance means having a false belief (Friedman & Petrashek, 2009; Jara-Ettinger, Floyd, Tenenbaum, & Schulz, 2017). If children assumed that ignorant agents should fail due to a false belief, then participants should have judged that both agents were equally knowledgeable (since both were successful). Our results suggest that by age five, children make principled belief inferences from agents’ behavior. Specifically, children expect both knowledgeable and ignorant agents to act consistently with their prior knowledge, and they use these expectations to infer what other people know.

Note, however, that children were only ever taught how the training toy worked. If children (reasonably) assumed all the toys worked in the same way, they may have been surprised to see a puppet activate the inconsistent toy. Perhaps children attributed greater knowledge to this agent not because his actions were inconsistent with his prior knowledge, but because the actions (and their outcome) were inconsistent with children’s own beliefs. In other words, children might simply attribute knowledge to agents who teach them something new, or show them something unexpected. We test this possibility in Experiment 2.

Experiment 2

Participants in Experiment 1 learned only how the first (training) toy worked. If participants attributed greater knowledge to the inconsistent actor because he taught them something new or unexpected, teaching children how all the toys work should cause performance to fall to chance because, now, neither agent can provide any novel information.

To address this, Experiment 2 differs in three substantial ways. First, we taught participants how all the toys worked. To reduce concerns about memory load, we used only two machines in this task: the training toy, and the inconsistent toy. Second, when trying to activate the novel toy, both puppets failed. One puppet pressed the red button (consistent with his prior experience), and one pressed the incorrect black button (inconsistent with his prior experience). Finally, we emphasized throughout that one of the puppets knew more, but not all, about the toy, making it plausible that both puppets could fail. Together, these changes allow us to test whether children attribute greater prior knowledge to agents whose actions are not explained by their prior experience, even when the agent fails to achieve their goal.

Method

Participants 72 four-, five- and six-year-olds (mean age: 5.56 years, range: 3.99 – 6.92 years; $n = 24$ participants per age group) were recruited at a local children's museum. 26 participants were excluded from analyses and replaced because: they did not pass the pre-registered inclusion questions ($n = 13$), due to experimenter error ($n = 5$), interruptions or interference with the procedure ($n = 3$), because the participant did not answer the test question within 30s ($n = 3$), because the participant had already participated in the past ($n = 1$), or due to developmental delays ($n = 1$).

Stimuli Materials were identical to those of Experiment 1, except that now only two machines were used: the training toy, and the inconsistent toy.

Procedure Experiment 2 began identically to Experiment 1. Participants and then puppets were familiarized with the training toy. Next, after placing the puppets underneath the table, the experimenter produced the additional (inconsistent) toy. In contrast to Experiment 1, the experimenter told participants that this toy was “a little bit different.” She explained that the red button did not activate this toy, and that only one of the black buttons (the correct black button) made the toy play music. She demonstrated all of the buttons, and then allowed the participant to press each button. Thus, participants were explicitly taught how the toys worked, and experienced for themselves that the toys worked differently.

Next, both puppets returned. The experimenter explained that one of the puppets had seen the toy before, and knew a little bit about it. And she explained that the other puppet

had never seen the toy before. The experimenter noted that one of the puppets knew more about the toy, but she didn't know which one. The participant's task was to help the experimenter identify which puppet knew more about the toy.

Each puppet was asked to make the toy go in turn. During each puppet's turn, the other agent was placed underneath the table. One puppet's actions were consistent with his prior knowledge, saying, “Hmm. Well, the red button made this [original] toy go, so the red button makes this toy go too,” pointing to the two relevant buttons as he spoke. He pressed the red button. The button did not activate the toy, and the puppet exclaimed “oh!” in surprise when nothing happened. The other puppet's actions were inconsistent with his prior knowledge, saying, “Hmm. Well, the red button made this [original] toy go, but this black button makes this toy go,” pointing to the two relevant buttons as he spoke. He pressed the incorrect black button. The button also did not activate the toy, and the puppet exclaimed “oh!” in surprise when nothing happened.

After each puppet pressed a button, the experimenter asked the test question: “[Child name], remember how I told you that one of my friends knows more about this toy? Can you tell me, which friend knows more?” Participants were asked to explain their answer. The inclusion questions were asked last, with children asked to match each puppet to the button he had pressed on the novel (inconsistent) toy: “[Child name], can you remind me, which one of my friends pressed this button [both puppets point to one button]? And which one of my friends pressed this button [both puppets point to the other button]?”

The identity of the puppet whose turn was first, and the button this agent pressed was always counterbalanced. Additionally, the side each puppet was presented on (left/right) was randomized.

Results and Discussion

Results were coded as in Experiment 1, with 26 participants excluded and replaced (see Participants). Overall, 61.1% of participants attributed knowledge to the puppet who pressed the black button, a proportion reliably higher than chance (44 of 72; 95% CI: 50 - 72.2). A logistic regression predicting performance based on age did not reveal a significant age difference ($\beta = 0.42$, $p = .14$). But while participants in all age groups preferred to attribute knowledge to the agent whose actions were inconsistent with his prior experience, only six-year-olds' preferences were robust. While 70.8% of six-year-olds judged that the agent who pressed the black button was more knowledgeable (17 of 24; 95% CI: 54.17 - 87.5), only 54% of four-year-olds (13 of 24; 95% CI: 33.33 - 75) and 58% of five-year-olds (14 of 24; 95% CI: 37.5 - 79.17) also made this judgment. In sum, although no age difference was obtained, only six-year-olds reliably judged that the agent whose failure was inconsistent with his prior experience had greater knowledge.

These findings suggest that children do not simply attribute knowledge to agents who show them something new. If they did, they should have performed at chance, as neither puppet taught children anything new. Instead, our results suggest that, by age six, children not only expect ignorant agents to act based on their prior knowledge, but also understand that knowledge runs along a continuum: agents can know more or less about any given topic. Thus, by age six, children attribute more knowledge to agents whose prior experience cannot explain their actions, even when these actions fail to fulfill their goal.




Figure 2: Results from both experiments. The error bars are bootstrapped 95% confidence intervals, and the dotted line indicates chance performance (50%). In Experiment 1, five- and six-year-olds judged that an agent whose success could not be explained by his prior experience had additional knowledge. In Experiment 2, six-year-olds judged that an agent whose failure could not be explained by his prior experience had additional (albeit incomplete) knowledge.

General Discussion

To successfully interact with others, we must understand what they know and believe, what they feel, and what they want. Children understand the link between mind and behavior early in life, inferring goals (Csibra, Gergely, Bíró, Koós & Brockbank, 1999; Jara-Ettinger, Gweon, Schulz & Tenebaum, 2016), beliefs (Onishi & Baillargeon, 2005; Rubio-Fernandez & Geurts, 2013) and desires (Doan, Denison, Lucas & Gopnik, 2015; Repacholi & Gopnik, 1997) from others' actions. Yet, while much work has shown that even young children have expectation about how knowledgeable agents should behave (Surian, Caldi & Sperber, 2007), less work has investigated whether children understand how ignorant agents might apply their prior knowledge to new situations.

Here we found that preschoolers expect ignorant agents to act based on their prior knowledge. When agents' past experience cannot explain their actions, children infer that

these agents must have additional knowledge. In Experiment 1, five- and six-year-olds (but not four-year-olds) judged that an agent whose observable past experience could not explain his successful actions must've had additional knowledge. In Experiment 2, four- to six-year-olds (but only six-year-olds reliably) judged that an agent whose observable past experience could not explain his failure must've had some (incomplete) additional knowledge.

Our results show that, by age five, children expect past experiences to shape agents' beliefs and guide their actions in new situations. These results are consistent with related work, which suggests that children do not reliably link ignorance to specific outcomes (Friedman & Petrashek, 2009; German & Leslie, 2001; Ruffman, 1996).

These findings also suggest several broader implications. First, while we often talk about "knowing" or "not knowing," knowledge is not binary. People are rarely completely ignorant or completely knowledgeable. More frequently, knowledge lies along a continuum. In Experiment 2, six-year-olds succeeded in identifying which of two agents knew more, even when both agents were wrong. If children believe that agents can only be fully knowledgeable or fully ignorant, they may not have attributed even partial knowledge in this case (perhaps judging that any agent who is wrong is equally ignorant). The results of this experiment suggest that, by age six, children represent knowledge and ignorance as two poles of the epistemic continuum, leveraging their expectations about how prior experience should affect agents' actions to infer the extent of their knowledge.

Second, these findings provide insight into the development of children's epistemic inferences. While prior work has thoroughly investigated young children's ability to represent others' beliefs (e.g., Onishi & Baillargeon, 2005; Wellman et al., 2001), less research has investigated how children infer belief from action. In our tasks, children had to infer agents' beliefs from their actions. This required understanding that each agent pressed the button they believed would make the toy go, and considering what role their past experiences played in shaping these beliefs. Past work suggests that children infer knowledge from action via a naïve theory of knowledge: a set of expectations about how ignorant and/or knowledgeable agents should act (Aboody, Huey, & Jara-Ettinger, 2018). Our results are consistent with this account, demonstrating that across varied contexts, children can infer what others know or believe by observing their actions.

Our results also open avenues for future work. First, Experiment 2 shows that children do not simply attribute knowledge to agents who show them something new or surprising. However, other simple rules may explain participants' performance. For example, children may expect ignorant agents to act the same way they've acted in the past, without representing their knowledge or beliefs. In our studies, specifically, children may have solved the task by matching agents' current actions to their prior acts,

licensing knowledge any time these acts were inconsistent. Future work can address this possibility by providing agents with knowledge, but not experience (e.g., by telling the puppets in Experiment 1 how the toy works but not allowing them to try it for themselves).

A second possibility is that children expect ignorant agents to try whatever is most reasonable, not in the context of agents' knowledge, but in the context of what children themselves think is reasonable. For example, children in our task could have assumed that the red button was the most obvious thing to try (regardless of agents' past experiences), and attributed prior knowledge to any agent who rejected this obvious solution. While it is unclear whether children in fact find the red button to be the obvious solution in this task, future work can address this possibility by reversing Experiment 1, and introducing children to a training toy that works the same as the inconsistent toy. If children now attribute greater knowledge to the agent who presses the (more visually salient) red button, this would show that children do not just think that ignorance means trying the most perceptually obvious answer.

Third, in both experiments, puppets' actions differed, but so did their explanations of their actions. Namely, one agent said: "Hmm. Well, the red button made this [original] toy go, **so** the red button makes this toy go **too**," and the other said, "Hmm. Well, the red button made this [original] toy go, **but** this black button makes this toy go." Although only two words differed between explanations, it is possible that this could explain children's epistemic attributions in our task. Note, however, that this would be consistent with our account, showing that children attribute knowledge to those who explicitly reject past experience. In addition, if the linguistic cue guides children's inferences, this would be interesting in its own right—the difference between "so too" and "but" is subtle, and to our knowledge, little work has investigated how such words affect children's belief inferences. To identify whether these explanations were critical to children's inferences, future work will leave them out. If children make the same judgments, this would provide evidence that performance in this task did not hinge upon puppets' explanations.

Fourth, in Experiment 2, it is possible that children did not think both puppets were equally wrong. Conceptually, the puppet who pressed the black button may have been closer to being right (since he knew that one of the black buttons made the toy go). It is possible that children didn't consider whether agents' prior knowledge explained their actions, and instead simply attributed greater knowledge to the agent who was closer to being correct. While possible, this account does not explain children's success in Experiment 1. Furthermore, it is unclear how to operationalize what it means to be "closer" to being right in Experiment 2: while one agent was conceptually closer (pressing a black button), the other was physically closer (pressing the red button, which was right next to the correct black button). It is unclear how the magnitude of agents' errors may have guided children's inferences in the current

task, but future work should investigate how this factor affects children's epistemic judgments.

Last, across both experiments, children's preferences strengthened with age (significantly in Experiment 1, and non-significantly in Experiment 2). Four-year-olds' failures in both experiments are consistent with prior work, which suggests that the ability to infer knowledge from behavior continues to develop between the ages of four and five (Aboody, Huey & Jara-Ettinger, 2018). But while five-year-olds succeeded in Experiment 1, they were not reliably above chance in Experiment 2. Why might this be?

One possibility is that identifying a completely knowledgeable agent (Experiment 1) is easier than judging which agent has greater (but still incomplete) knowledge (Experiment 2). Furthermore, given that children may equate accuracy with knowledge (Brosseau-Liard & Birch, 2010; Ronfard & Corriveau, 2016), it might be harder for them to attribute knowledge in the face of a failure.

It is also possible that five-year-olds do attribute knowledge based on a rule (for example, attributing knowledge to agents who act in a surprising way). This could explain their weaker performance in Experiment 2, although it is unclear why four-year-olds would not have followed the same rule (which would have led to success in Experiment 1). It is possible that four-year-olds have no rule for inferring belief from knowledge, five-year-olds depend on a rule (e.g., knowledge = rejecting the obvious), and six-year-olds have a deeper understanding of how prior knowledge shapes beliefs. Finally, it is always possible that task demands affected children's performance, although this would fail to explain the difference in five-year-olds' performance across the two studies. Future work will address these possibilities to further clarify how children's epistemic intuitions emerge and develop.

In sum, across two experiments, we find evidence that young children have expectations for how prior knowledge is likely to shape people's beliefs and guide their behavior. We find that children use these expectations to infer what others know (or don't know) from their actions and that, by age five, children do not expect ignorant agents to act as blank slates; rather, they expect ignorant agents to leverage relevant prior knowledge when planning their actions. Altogether, our findings suggest that even young children may understand how ignorance begets belief and action.

Acknowledgments

We thank the Boston Children's Museum, the Peabody Natural History Museum, and the families who participated in this research. We thank Sarah Wong and Ivana Bozic for help with coding, and Lindsay Stoner for help with coding and data collection. This material is based upon work supported by the Center for Brains, Minds, and Machines (CBMM), funded by NSF-STC award CCF-1231216.

References

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*.
- Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, *1*(4), 0064.
- Bartsch, K., & Wellman, H. M. (1995). *Children talk about the mind*. Oxford university press.
- Brosseau-Liard, P. E., & Birch, S. A. (2010). 'I bet you know more and are nicer too!': what children infer from others' accuracy. *Developmental Science*, *13*(5), 772-778.
- Csibra, G., Gergely, G., Bíró, S., Koós, O., & Brockbank, M. (1999). Goal attribution without agency cues: the perception of 'pure reason' in infancy. *Cognition*, *72*(3), 237-267.
- Dennett, D. C. (1987). *The intentional stance*. Cambridge, MA: MIT Press
- Doan, T., Denison, S., Lucas, C., & Gopnik, A. (2015). Learning to reason about desires: An infant training study. *Proceedings of the 37th Annual Conference of the Cognitive Science Society*.
- Friedman, O., & Petrashek, A. R. (2009). Children do not follow the rule "ignorance means getting it wrong". *Journal of Experimental Child Psychology*, *102*(1), 114-121.
- German, T. P., & Leslie, A. M. (2001). Children's inferences from 'knowing' to 'pretending' and 'believing'. *British Journal of Developmental Psychology*, *19*(1), 59-83.
- Gopnik, A., & Wellman, H. M. (1992). Why the child's theory of mind really is a theory. *Mind & Language*, *7*(1 - 2), 145-171.
- 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.
- Jara-Ettinger, J., Floyd, S., Tenenbaum, J. B., & Schulz, L. E. (2017). Children understand that agents maximize expected utilities. *Journal of Experimental Psychology: General*, *146*(11), 1574.
- 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.
- O'Neill, D. K. (1996). Two-year-old children's sensitivity to a parent's knowledge state when making requests. *Child Development*, *67*(2), 659-677.
- Onishi, K. H., & Baillargeon, R. (2005). Do 15-month-old infants understand false beliefs? *Science*, *308*(5719), 255-258.
- Repacholi, B. M., & Gopnik, A. (1997). Early reasoning about desires: evidence from 14-and 18-month-olds. *Developmental Psychology*, *33*(1), 12.
- Ronfard, S., & Corriveau, K. H. (2016). Teaching and preschoolers' ability to infer knowledge from mistakes. *Journal of Experimental Child Psychology*, *150*, 87-98.
- Richardson, H., Lisandrelli, G., Riobueno-Naylor, A., & Saxe, R. (2018). Development of the social brain from age three to twelve years. *Nature Communications*, *9*(1), 1027.
- Rubio-Fernández, P. (2019). Memory and inferential processes in false-belief tasks: An investigation of the unexpected-contents paradigm. *Journal of Experimental Child Psychology*, *177*, 297-312.
- Rubio-Fernández, P., & Geurts, B. (2013). How to pass the false-belief task before your fourth birthday. *Psychological Science*, *24*(1), 27-33.
- Ruffman, T. (1996). Do children understand the mind by means of simulation or a theory? Evidence from their understanding of inference. *Mind & Language*, *11*(4), 388-414.
- Saxe, R. (2005). Against simulation: the argument from error. *Trends in Cognitive Sciences*, *9*(4), 174-179.
- Surian, L., Caldi, S., & Sperber, D. (2007). Attribution of beliefs by 13-month-old infants. *Psychological Science*, *18*(7), 580-586.
- 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.
- Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, *69*(1), 1-34.