Cognition 228 (2022) 105212

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Preschoolers decide who is knowledgeable, who to inform, and who to trust via a causal understanding of how knowledge relates to action

Rosie Aboody [a,*], Holly Huey [b], Julian Jara-Ettinger [a]

[a] Yale University, United States of America
[b] University of California, San Diego, United States of America

ARTICLE INFO

Keywords:
Social cognition
Social learning
Theory of Mind
Testimony
Knowledge

ABSTRACT

Preschoolers are discerning learners, preferring to trust people who are accurate, reliable, and appropriately-informed. Do these preferences reflect mental-state reasoning, where children infer what others know from their behavior, or do they reflect a reliance on simple cues? In Experiment 1 we show that four- and five-year-olds can infer knowledge from others' behavior when superficial cues and actions are matched across agents. Experiments 2a and 2b further suggest that children track how agents acquired their knowledge, and may use this to determine what different agents will (and will not) know. Finally, Experiment 3 shows that children independently make these knowledge inferences when deciding whom to trust. Our findings suggest that by age four, children have expectations about how knowledge relates to action, use these expectations to infer what others know from what they do, and rely on these inferences when deciding whom to trust.

1. Introduction

Most of what we know, we learn from others. We learn what foods are safe to eat without poisoning ourselves, compute derivatives without rediscovering calculus, and find out what's happening in Washington without witnessing it firsthand. Yet, while social learning removes the cost and risk of exploration, it introduces a new challenge: How do we know whether others' testimony is true? Agents can be misinformed, unreliable, or malicious, and it is critical to distinguish those who are knowledgeable and trustworthy from those who are not (Lackey, 2010; Sperber et al., 2010).

The ability to identify useful informants is rooted early in development. By preschool, children selectively learn from (and engage with) those who are accurate (preferring agents who label objects correctly; Birch, Vauthier, & Bloom, 2008; Jaswal & Neely, 2006; Koenig, Clément, & Harris, 2004; Koenig & Harris, 2005; Koenig & Woodward, 2010; Scofield & Behrend, 2008; Pasquini, Corriveau, Koenig, & Harris, 2007), reliable (preferring agents who behave consistently over time; Chow, Poulin-Dubois, & Lewis, 2008; Zmyj, Buttelmann, Carpenter, & Daum, 2010), and adequately-informed (preferring agents who had access to relevant information; Koenig, 2012; Robinson, Butterfill, & Nurmsoo, 2011; Brosseau-Liard & Birch, 2011; Nurmsoo & Robinson, 2009a; see Sobel & Kushnir, 2013 for review).

What are the underlying representations that support this behavior? One possibility is that children choose who to trust based on their understanding of other people's minds—a *Theory of Mind* (Gopnik & Meltzoff, 1997; Wellman, 2014). This is consistent with evidence that children trust accurate speakers but not accurate inanimate objects (e.g., an audio speaker; Koenig & Echols, 2003), prefer to learn from agents who can answer questions unaided (rather than agents who ask for help, and simply repeat the accurate answer they receive; Einav & Robinson, 2011), recover trust in previously inaccurate agents if they now have the right knowledge (Brosseau-Liard & Birch, 2011; Nurmsoo & Robinson, 2009a), dismiss previously accurate agents when their statements conflict with children's own experience (Clément, Koenig, & Harris, 2004; Jaswal, McKercher, & VanderBorght, 2008), and reverse their baseline trust in consensus (Corriveau, Fusaro, & Harris, 2009; Fusaro & Harris, 2008) when a dissenter is actually better-informed than the majority group (Aboody, Yousif, Sheskin and Keil, 2022; Einav, 2014). Children at this age also understand what kinds of knowledge other people can have (Lockhart, Goddu, Smith, & Keil, 2016; Pillow, 1993) and how it ought to be verified (Butler, Schmidt, Tavassolie, & Gibbs, 2018; Butler, Gibbs, & Tavassolie, 2020). They can represent complex epistemic states like expertise (Kushnir, Vredenburgh, & Schneider, 2013; Lutz & Keil, 2002), knowledge and ignorance (Koenig et al., 2015; Kushnir & Koenig, 2017; O'Neill, 1996; Ronfard & Corriveau, 2016; Sabbagh & Baldwin,

* Corresponding author at: Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, United States of America
E-mail address: rosie.aboody@yale.edu (R. Aboody).

https://doi.org/10.1016/j.cognition.2022.105212
Received 11 October 2021; Received in revised form 12 May 2022; Accepted 23 June 2022
Available online 28 July 2022

2001; Surian, Caldi, & Sperber, 2007), and use these representations to guide learning (Mangardich & Sabbagh, 2018; Mills, Legare, Bills, & Mejias, 2010; Mills, Legare, Grant, & Landrum, 2011; Sabbagh & Shafman, 2009). Moreover, children infer knowledge following the same principles that they use to infer desires, through an expectation that agents act to maximize their expected utilities (selecting action plans that produce the greatest expected rewards while incurring the lowest possible costs; Aboody, Zhou, & Jara-Ettinger, 2021; Jara-Ettinger, Floyd, Tenenbaum, & Schulz, 2017).

However, children's preference for accurate, reliable, and well-informed agents may not be grounded in mental state reasoning. Preschooler Theory of Mind is notably brittle (Wellman, Cross, & Watson, 2001) and may continue developing well past the early preschool years (Richardson, Lisandrelli, Riobueno-Naylor, & Saxe, 2018). For instance, even older preschoolers (who generally pass explicit false belief tasks) have poor intuitions about the expected behavior of ignorant agents (Chen, Su, & Wang, 2015; Aboody, Zhou, Flowers, & Jara-Ettinger, 2019; Friedman & Petrashek, 2009; German & Leslie, 2001). Moreover, preschoolers are not always sensitive to the causes behind other people's errors, failing to distinguish inaccuracy arising from a lack of perceptual access (a blindfolded agent mislabeling an item they cannot see) from inaccuracy that implies a lack of knowledge (an agent mislabeling an item they can clearly see; Nurmsoo & Robinson, 2009b).

These studies raise an open question: How complex is children's understanding of the relation between knowledge and accuracy? And to what extent is such an understanding in use when preschoolers decide whom to learn from? If preschoolers lack a full causal model of the relation between epistemic states and actions (e.g., struggling to understand how ignorant or partially-informed agents are likely to act; Aboody et al., 2019; Chen et al., 2015), then approximating mental-state inferences through simpler rules might be easier and generally effective. Indeed, children often rely on superficial cues that do not necessarily reveal agents' knowledge, preferentially learning from those who are familiar (Corriveau & Harris, 2009; Lucas et al., 2017), dominant (Bernard et al., 2016), confident (Kominsky, Langthorne, & Keil, 2016; Tenney, Small, Kondrad, Jaswal, & Spellman, 2011), as well as those who are in-group members (Elashi & Mills, 2014; Hetherington, Hendrickson, & Koenig, 2014), and who share their accent (Kinzler, Corriveau, & Harris, 2011; although accuracy trumps accent; Corriveau, Kinzler, & Harris, 2013).

The goal of our paper is to shed light on the representations underlying children's choices for whom to learn from, focusing on three questions. First, to what extent do children select informants based on superficial behavioral cues? Do children circumvent mental-state inferences by associating different epistemic states with different observable cues (e.g., possibly equating knowledge with accuracy, and ignorance with error; Ruffman, 1996; although see Friedman & Petrashek, 2009)? Or can children infer knowledge through their Theory of Mind (e.g., considering which knowledge states best explain an observed set of actions)? Second, what are children's attributions of knowledge like? Are they broad attributions (e.g., general beliefs like "Max is knowledgeable," as suggested by halo-like effects in children; Brosseau-Liard & Birch, 2010) or representations that track the content of what others know (e.g., concrete beliefs like "Max is knowledgeable about what's inside this box")? Finally, do children independently choose to derive such inferences when they would be helpful (e.g., when deciding whom to trust)? Or do they do so only when explicitly prompted to make a judgment about others' epistemic states?

Here we present three studies that shed light on these questions, using simple social events where two agents take identical actions and produce identical outcomes, differing only in the order in which they produce each action. This enabled us to design events where agents are indistinguishable in terms of superficial cues, but whose behavior reveals different epistemic states when analyzed through a Theory of Mind. We focus on four- and five-year-olds because younger children struggle to explicitly reason about other people's beliefs and knowledge (Robinson & Whitcombe, 2003; Wellman et al., 2001).

2. Approach to analyses

Following current recommendations for statistical best practices, we take an estimation approach to data analysis (as opposed to relying on null-hypothesis significance testing; Cohen, 1994; Cumming, 2014). We estimate effect sizes by bootstrapping our data and obtaining 95% confidence intervals. For individual effects, we take confidence intervals that do not cross chance as evidence of a reliable effect. For differences across conditions, we take confidence intervals that do not cross 0 (the point of no difference) as evidence of a reliable effect. Finally, to test for developmental effects while controlling for question type and other task features, we use Bayesian mixed-effects logistic regressions (along with the 95% equal-tailed credible interval from the posterior distribution over each coefficient). We take intervals that do not cross 0 (the point of no difference) as evidence of a reliable effect.

3. General methods

The procedure, analysis plan, sample size, and exclusion criteria for all experiments were pre-registered, unless otherwise indicated. All pre-registrations, scripts, data, and analysis files are available in the OSF project page: https://osf.io/mhcvt/ (Aboody, Huey & Jara-Ettinger, 2022). The pre-registered sample size for all experiments was determined through Monte Carlo power analyses. A post-hoc analysis confirmed that our sample was appropriately powered, with power > 0.8 for every test question in every experiment (see Supplemental Materials for details).

In Experiments 1 and 2a, participants were asked two test questions, order counterbalanced. As we collected data for Experiments 1 and 2a (run approximately concurrently), we realized that our studies could be underpowered for detecting order effects in our test questions, particularly as a function of age (which we did not originally anticipate, but later realized we would need to test for, to ensure that an order effect could not explain any potential developmental patterns we might obtain). Thus, before finishing data collection, we chose to replicate each experiment. For both experiments, performance did not differ between samples (see Supplemental Materials). We report analyses on each sample separately, and also aggregate this data to compute effect sizes more accurately—but do not compute *p*-values.

4. Experiment 1

Experiment 1 tests whether preschoolers can infer which of two agents is knowledgeable when agents are matched for low-level superficial cues. Participants were introduced to two puppets, one that accurately predicted what was under two cups before revealing their contents; and one that accurately stated what was under the cups after revealing their contents. If children rely on simple behavioral cues like accuracy or checking, they should believe that both agents are equally knowledgeable. But if children consider the epistemic states that best explain each agent's actions, they should judge that the predictor is more likely to have been knowledgeable.

4.1. Method

4.1.1. Participants

64 four- and five-year-olds (mean age: 5.02 years, range: 4.07–5.99 years; $n = 32$ participants per age group) were recruited at children's museums ($n = 57$), preschools ($n = 4$), or in-lab ($n = 3$). 32 participated in the pre-registered original experiment and an additional 32 participants in a direct replication. The direct replication was not pre-registered but was identical in all aspects to the original pre-registered experiment. 23 additional participants were recruited but not included in the study (21 at museums, and 2 at preschools; see Results).

4.1.2. Stimuli

Stimuli consisted of two male puppets, three paper cups (red, blue and yellow), and three small animal figurines (a fox, a hippo, and a deer).

4.1.3. Procedure

The participant and the experimenter sat on opposite sides of a table. The experimenter first showed participants three cups (sitting inverted on the table), and lifted each cup to reveal an animal figurine hidden underneath. The experimenter then introduced two puppets, Max and Sam, explaining that, "Right before you came in here, one of our friends snuck out from under the table and peeked underneath all the cups! And one friend stayed under the table, and he never saw anything." The experimenter then explained that she'd "ask our friends some questions to find out who peeked underneath all the cups."

Participants were allowed to choose which puppet they wanted to hear from first, and which cup they wanted to ask the puppet about. This was to avoid any pragmatic concerns that could arise if participants interpreted the experimenter's choice of puppet or cup as meaningful (e.g., perhaps the experimenter suspects the puppet she questions first?). Puppets' roles (predictor/observer) were assigned after participants chose which puppet to ask first, so the role assigned to the first puppet was counterbalanced across participants.

When the predicting agent was asked what was under a cup, he correctly stated the animal name, lifted the cup to reveal its contents, and then looked at the animal (e.g., saying 'There's a hippo under this cup,' revealing a toy hippo, and looking down at it; Fig. 1). By contrast, when the observing agent was asked what was under a cup, he first lifted the cup to reveal its contents, looked down, and then correctly stated the animal name (e.g. revealing a toy hippo, looking down at it, and then saying 'There's a hippo under this cup'). Thus, both puppets performed identical actions, but in the opposite order: the predictor first said what was under the cup and then looked; the observer first looked under the cup, and then said what was there (see Fig. 1). After each puppet stated the contents of the first cup, participants chose a second cup to ask the puppets about and each puppet repeated the same actions (predicting or observing the cup's contents), in the same order as in the first trial. Each puppet was always alone when asked what was under each cup.

Participants then answered two test questions, order counterbalanced. In the *prior knowledge* test question, participants were asked, "which one of our friends peeked?" And in the *generalization* test question, participants were asked, "which one of our friends knows what's under this cup?" (referring to the last remaining cup, which neither of the puppets had interacted with); see Supplemental Materials for scripts and additional details.

4.2. Results

For the 94% of participants whose sessions were video or audio taped (n = 82/87), two coders who were not involved in data collection determined exclusions according to pre-registered criteria. The first coder was blind to participants' final answers, checking for any experimenter errors, family interference, and ensuring that the participant was attentive. The second coder, blind to condition, checked whether participants answered the test questions, and whether the experimenter or family members behaved in any way that could affect participants' choices. For participants who were not video or audio taped (6% of participants; $n = 5/87$), the experimenter took notes on any deviations from the procedure, and the first author determined exclusions by comparing these notes to the pre-registered inclusion criteria. Twenty-three participants were recruited but not included in the study, because the participant interfered with the study (by revealing the contents of the cups to the puppets; $n = 11$), due to experimenter error ($n = 6$), because the participant did not complete the study ($n = 4$), or because they did not speak English fluently ($n = 2$).

Of the final 64 participants included in the study, 68.8% judged that the predicting agent had peeked under the cups ($n = 44$ of 64; 95% CI: 57.8–79.7) and 75% judged that the same agent also knew what was inside the last cup ($n = 48$ of 64; 95% CI: 65.6–85.9). See Fig. 2. There was no reliable difference in the proportion of participants who selected the predicting agent in response to each test question ($\Delta = 6.25\%$; 95% difference CI: −9.4–21.9). To test for effects of age and question type, we conducted a Bayesian mixed-effects logistic regression predicting performance as a result of age, question type, and their interaction (with random intercepts by participant and question order, as well as a random slope of question order per test question and centered participant age; regression not pre-registered). This regression revealed no effects of age, question type, or their interaction (all $|\beta|$'s between [0.23, 0.72]; all 95% CI's crossing 0, lower bound ≤ -0.74, upper bound $\geq$ 2.14).

All results are qualitatively the same when performance is analyzed separately for our original study and our replication, with participants in both samples reliably selecting the predicting agent when asked who peeked (respectively: $n = 21$ of 32, 95% CI: 50–81.3; $n = 23$ of 32, 95% CI: 56.3–87.5), and when asked who would know what was under the final cup (respectively: $n = 27$ of 32, 95% CI: 71.9–96.9; n = 21 of 32, 95% CI: 50–81.3). Consistent with this, there was no reliable difference in the proportion of participants who selected the predicting agent in response to each test question in either the original sample ($\Delta = 18.75\%$; 95% difference CI: −3.1–40.6) or the replication sample ($\Delta = -6.25\%$; 95% difference CI: −28.1–15.6).

In the original sample, a Bayesian mixed-effects logistic regression

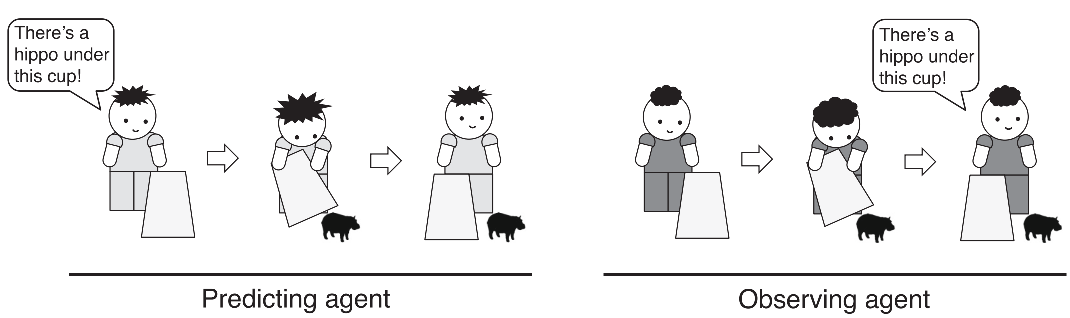

Fig. 1. Agents' behavior in Experiments 1–3. The predicting agent first stated the contents of the cup, and then revealed them. The observing agent first revealed the contents of the same cup, and then described them. In Experiment 1, participants were asked which puppet had peeked under the cups before the task began, and which puppet knew what was under the final cup, which neither puppet had interacted with. In Experiments 2a and 2b, the animal under the final cup was replaced without the puppets' knowledge. In Experiment 2a participants were asked the same questions as in Experiment 1. In Experiment 2b participants were again asked which puppet had peeked, and they were given an opportunity to tell one of the puppets the name of the animal currently under the final cup. Finally, in Experiment 3, participants watched the puppets disagree about the contents of the final cup, and were asked to endorse one agent's testimony.

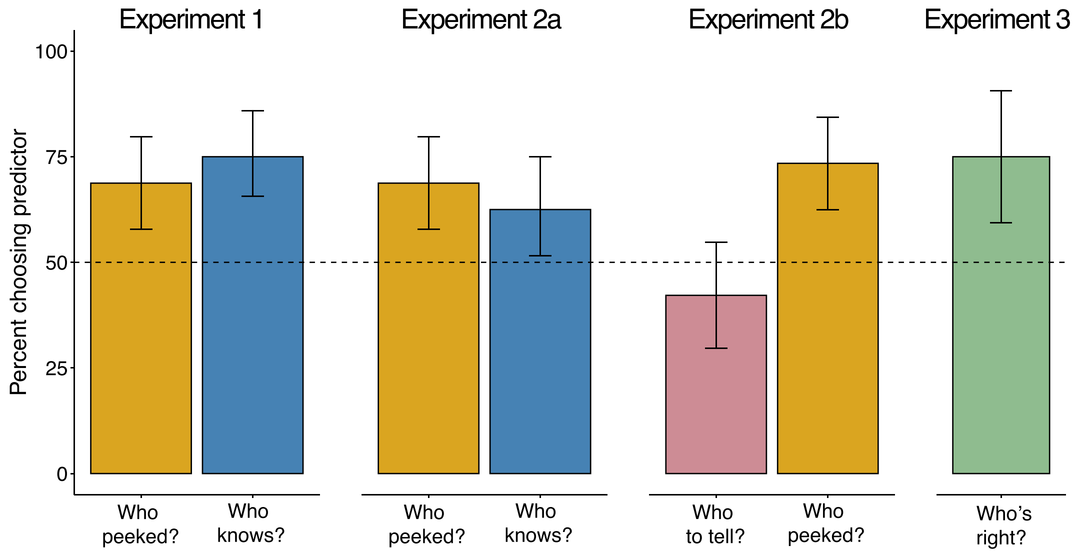

Fig. 2. Results from all three experiments. The dotted line indicates chance performance, and the error bars are bootstrapped 95% confidence intervals. In Experiment 1, participants judged that the predicting agent had peeked underneath the cups, and knew what was under the remaining cup. In Experiment 2a, participants judged that the predicting agent had peeked under the cups, but also judged that he knew what was under the remaining cup (whose contents had been switched out). In Experiment 2b, participants again judged that the predicting agent had peeked, but had no preference when given a chance to tell an agent the name of the new animal under the last cup. In Experiment 3, participants preferred to endorse the predicting agent's testimony when the two agents disagreed about the contents of the last cup, suggesting that they inferred who was knowledgeable without any explicit prompts.

(identical to the regression run on the full data) revealed no effects of age, question type, or their interaction (all $|\beta|$'s between [1.13, 1.84]; all 95% CI's crossing 0, lower bound ≤ -3.07, upper bound ≥ 1.81; regression not pre-registered). In the replication sample the same regression revealed no effect of age or question type ($|\beta|$'s between [0.47, 0.77]; all 95% CI's crossing 0, lower bound ≤ -2.86, upper bound ≥ 2.64), but did reveal an interaction between age and test question, with participants performing better on the “who peeked” test question with age ($\beta = 2.21$, 95% CI: 0.07–4.79; regression not pre-registered). Note, however, that we obtained no such effect when aggregating our data, or in a meta-analysis over all experiments (see Supplemental Materials for details of the meta-analysis, supplementary analyses, and regression tables), suggesting that this effect may reflect a false positive due to the small sample size.

5. Experiment 2a

Experiment 1 shows that children can infer which of two agents is knowledgeable when superficial cues are matched, distinguishing an agent who makes accurate predictions from one who merely makes accurate observations. In Experiment 2 we explore the representational content of these epistemic attributions. Do children's inferences result in coarse, broad knowledge attributions (for instance, that one agent is generally knowledgeable in the relevant domain, or even just more knowledgeable overall)? Or do children use an agent's behavior not only to infer if they are knowledgeable, but to determine the exact contents and limits of what they know (tracking what a knowledgeable agent knows, but also what they might not)?

5.1. Method

5.1.1. Participants

64 four- and five-year-olds (mean age: 4.97 years, range: 4.0–5.99 years; $n = 32$ participants per age group; 32 in the pre-registered original experiment, and 32 in a direct replication) were recruited at children's museums ($n = 52$), preschools ($n = 8$), or in-lab ($n = 4$). The direct replication was not pre-registered; it is identical in all aspects to the original pre-registered experiment. 21 additional participants were recruited but not included in the study (14 from children's museums, 2 from preschools, and 5 from festivals in the New Haven area; see Results).

5.1.2. Stimuli

Materials were identical to those of Experiment 1, with the addition of a small box (6 × 7 × 7 in.) containing six animal figurines: a cat, a duck, a penguin, a parakeet, a rabbit, and an ostrich.

5.1.3. Procedure

The procedure was identical to that of Experiment 1 with one exception. After the two puppets predicted or observed the contents of the first two cups, the puppets left and the experimenter said, “But you know what? We haven't asked our friends about this cup yet,” and pointed to the remaining cup. The experimenter brought out a box and continued, “And I thought we could play a trick. I have this box of animals—can you choose one?” After the participant chose a new animal figurine, the experimenter said, “Ok! So let's put [original animal] back in this box, and let's put [new animal] underneath. So now, [new animal] is here instead!” The puppets were then brought back, and participants answered the same test questions (prior knowledge and generalization; order counterbalanced); see Supplemental Materials for scripts and additional details.

5.2. Results

Results were coded as in Experiment 1, according to identical pre-registered criteria. 89.4% of participants were video or audio taped ($n = 76/85$). 21 participants were recruited but not included in the study because the participant interfered with the study (by revealing the contents of the cups to the puppets; $n = 12$), due to experimenter error ($n = 3$), because the participant was distracted ($n = 2$), non-neurotypical ($n = 2$), did not answer a test question ($n = 1$), or due to sibling interference ($n = 1$).

If children attribute epistemic states based on a causal understanding of how knowledge relates to action, they should continue to judge that the predictor peeked. However, they should no longer judge that he would also know what was underneath the last cup—because neither agent had seen the switch—performing at chance on the generalization question. Consistent with our first prediction, 68.8% of participants judged that the predicting agent had peeked under the cups ($n = 44$ of 64; 95% CI: 57.8–79.7), replicating our finding in Experiment 1. Contrary to our predictions, however, 62.5% of participants judged that the predicting agent also knew what was under the last cup, a proportion reliably higher than chance ($n = 40$ of 64; 95% CI: 51.6–75). There was no reliable difference in the proportion of participants selecting the predicting agent in response to each test question ($\Delta = -6.25\%$; 95% difference CI: −21.9–10.9). Moreover, the proportion of children who judged that the predicting agent knew what was under the last cup was not reliably lower than the proportion we found in Experiment 1 (75% of participants in Experiment 1 compared to 62.5%; 95% CI: −28.1–3.1). To test for effects of age and question type, we conducted a Bayesian mixed-effects logistic regression predicting performance as a result of age, question type, and their interaction (with random intercepts by participant and question order, as well as a random slope of question order per test question and centered participant age; regression not pre-registered). This regression revealed no effects of age, question type, or their interaction (all $|\beta|$'s between [0.30, 0.34]; all 95% CI's crossing 0, lower bound ≤ -1.03, upper bound ≥ 1.73).

These results are qualitatively similar when performance is analyzed separately in the original experiment vs. the replication, with participants in both samples reliably selecting the predicting agent when asked who peeked (respectively: $n = 23$ of 32, 95% CI: 56.3–87.5; $n = 21$ of 32, 95% CI: 50–81.3). However, while participants in the original sample reliably judged the predicting agent would know what was under the final cup ($n = 21$ of 32, 95% CI: 50–81.3), participants in the replication sample showed no reliable preference for either agent (with $n = 19$ of 32 judging the predicting agent was knowledgeable, 95% CI: 43.8–75). However, participants' judgments across samples did not differ for either test question (*prior knowledge*: $\Delta = -6.25\%$; 95% difference CI: −28.1–15.6; *generalization*: $\Delta = -6.25\%$; 95% difference CI: −31.3–15.6), and within each sample, there was no reliable difference in the proportion of participants who selected the predicting agent in response to each test question (original sample: $\Delta = -6.25\%$; 95% difference CI: −28.1–15.6; replication sample: $\Delta = -6.25\%$; 95% difference CI: −31.3–15.6). Additionally, our combined sample suggests that participants overall did expect the predictor to be knowledgeable. Taken together, these results suggest that the effect was simply too low to be detected with a smaller sample size. Furthermore, treating age as a continuous independent variable, a Bayesian mixed-effects logistic regression (identical to the regression run on the full data) revealed no effects of age, question type, or their interaction in either the original sample (all $|\beta|$'s between [0.16, 0.88]; all 95% CI's crossing 0, lower bound ≤ -1.22, upper bound ≥ 2.35) or the replication sample (all $|\beta|$'s between [0.22, 2.08]; all 95% CI's crossing 0, lower bound ≤ -2.11, upper bound ≥ 2.50; analyses not pre-registered). See Supplemental Materials for supplementary analyses and regression tables.

5.2.1. Exploratory response-time analysis

Contrary to what we expected, participants judged that the predictor still knew what was under the final cup when the animal figurine had just been replaced, rather than performing at chance. One possible explanation is that children made a broad knowledge attribution to the agent who made an accurate prediction, without tracking the boundaries of what the agent is likely to know. Note, however, that this test question had no correct answer, since neither puppet knew what was under the final cup. While we expected that, in the absence of a correct answer, participants would respond randomly, it is also possible that participants selected the predictor simply due to a lack of a better option (after all, the predicting agent knew what was under the final cup at one point in time, while the observing agent never did).

To explore this possibility, we tested whether children's response times varied as a function of test question. A coder blind to our hypotheses recorded participants' response times to the test questions in Experiments 1 and 2a (for all participants where audio or video was available; 95.3% and 89.1% of participants in Experiments 1 and 2a, respectively), measuring the time between the end of the test question and the onset of the participant's answer.

This exploratory analysis revealed no significant difference in children's response times in Experiment 1 as a function of test question (*prior knowledge* question, $M = 2.56$ s; *generalization* question, $M = 3.20$ s; $t(60) = -0.96$, $p = 0.34$, $d = 0.12$ by paired *t*-test). By contrast, children's response times in Experiment 2a were significantly different, with participants taking an average of 2.18 additional seconds to answer the *generalization* question relative to the *prior knowledge* question (*prior knowledge* question, $M = 1.73$ s, *generalization* question, $M = 3.91$ s; $t(56) = -3.13$, $p = 0.0028$, $d = 0.42$ by paired t-test). See Fig. 3. These results suggest that participants were uncertain about how to answer the *generalization* question in Experiment 2a, opening the possibility that children indeed recognized that neither agent knew what was under the last cup, but defaulted to the predicting agent due to the forced-choice nature of our task. Experiment 2b further tests this possibility.

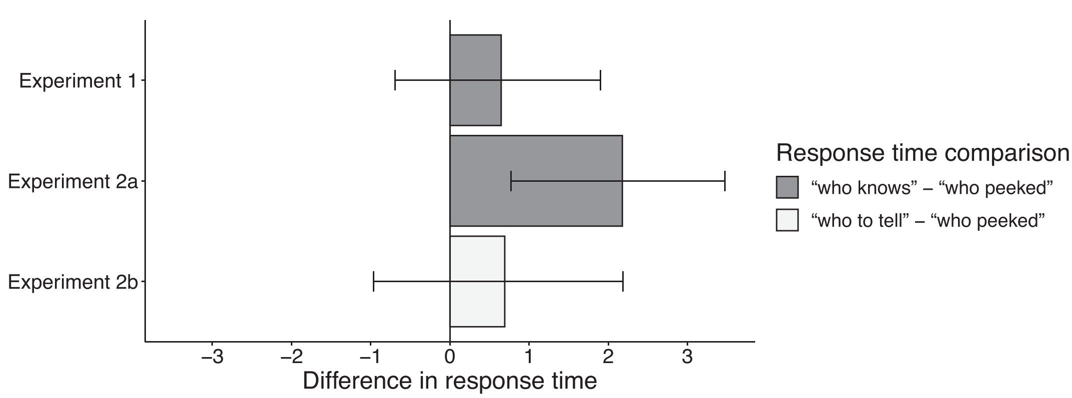

Fig. 3. The difference in response time in seconds for Experiments 1, 2a, and 2b. Error bars are paired 95% bootstrapped confidence intervals over the difference in response time. The fill of each bar indicates the test questions asked, and how the response time difference was computed. In Experiment 1, there is no reliable difference in response time across the two test questions ("who knows" – "who peeked"). In Experiment 2a, there is a reliable difference, with the same participants taking reliably longer to answer the "who knows" generalization test question than they took to answer the "who peeked" prior knowledge test question. Finally, in Experiment 2b, there is no reliable difference in response time across the two test questions ("who to tell" – "who peeked").

6. Experiment 2b

Participants' choices in Experiment 2a suggest that they struggled to track what different agents know, over-attributing knowledge to the agent that made accurate predictions. However, this experiment required participants to decide which of two ignorant agents was knowledgeable (leading to our original prediction of chance performance). In the absence of a better option, children may have simply chosen the agent who had previously been more knowledgeable. Indeed, participants took significantly longer to answer this question, suggesting that children were unsure about what to say.

To further test the possibility that children struggle to track what different agents know, Experiment 2b replicated Experiment 2a; but rather than asking participants to judge who knew what was under the final cup (as neither actually knew what was there), participants were given the chance to select one agent to inform.

If participants believe that the predictor knows what's under the final cup even after its contents are switched (as suggested by their answers in Experiment 2a), they should prefer to inform the observer (i.e., the only agent whom they believe to be ignorant). However, if participants understand that both agents are ignorant once the animal figurine is switched (as suggested by their response time in Experiment 2a), they should be equally likely to inform either puppet, responding at chance.

6.1. Method

Data collection for this study was interrupted due to the COVID-19 pandemic and this task was therefore completed online. To ensure that the switch to online testing did not affect our results, we first conducted a validation pilot, which revealed no differences between in-person and online samples (see Supplemental Materials). The revised plan in response to the pandemic was pre-registered, and both the original and revised pre-registration are available in our OSF repository.

6.1.1. Participants

64 four- and five-year-olds (mean age: 5.06 years, range: 4.0–5.91 years; $n = 32$ participants per age group) participated. The first 15 were recruited in-person, at children's museums ($n = 14$) and preschools ($n = 1$); the final 49 were recruited online. 8 additional participants were recruited but not included in the study (2 from preschools, 1 from a children's museum, and 5 recruited online; see Results).

6.1.2. Stimuli

For the in-person procedure, materials were identical to those of Experiment 2a. For online participants, the experimenter narrated while playing a pre-recorded puppet show (embedded in a PowerPoint presentation). The script and pre-recorded show were kept as similar as possible to the in-person version. The PowerPoint presentation is available in the project OSF page.

6.1.3. Procedure

The original in-person experiment began and proceeded in the same way as Experiment 2a up to the test question phase of the task. At this point, instead of asking participants who would know what was underneath the final cup, the experimenter offered participants the chance to tell one of the puppets what was there, saying, “So, we haven't asked our friends about this cup yet [pointing to cup]. But guess what! You can tell one of our friends the name of the animal underneath this cup. Which one of our friends do you want to tell?” Participants in this task were always asked who they wanted to inform before being asked which puppet had peeked. This pre-registered decision helped us avoid any potential influence of the peeking question on children's decision of who to inform (this was a conservative decision: we did not observe any order effects in our previous experiments, with the exception that in Experiment 1, participants were more likely to correctly answer the “who peeked” prior knowledge question when it was asked first).

The online task was designed to be as close as possible to the in-person design, with three pre-registered procedural adjustments. First, participants were no longer allowed to choose the puppet they wanted to question first, and the cups they wanted to ask about, as this was impractical for online testing given that the puppet show was prerecorded. Each puppet's role as predictor or observer was still counterbalanced across participants. Second, because participants could not reach in and choose an animal to place under the final cup, the actor in the video chose one randomly. To achieve this, the actor in the video placed the box of animals on the table, and then upended it, spreading out all of the animals on the table. The experimenter narrated, “I have this box of animals. Look! Do you see the animals? Look!” After participants had a chance to see the animals, the experimenter continued, “Let's put them back in the box. And we'll close our eyes and choose one.” The actor replaced the animals in the box, shook it around for several seconds, reached in, sampled an animal, and then held it out to the camera to show that she had pulled out a cat. The experimenter narrated, “Look! We got a kitty cat!” As the actor replaced the contents of the blue cup, the experimenter said, “Ok! So let's put the hippo back in this box, and let's put the kitty cat underneath. So now the kitty cat is here instead!”

Finally, each puppet was associated with a color (blue and green, indicated by their shirt) so that children could select a puppet by choosing a color (in line with recent work showing color-based response prompts to be an effective method for online testing of four- and five-year-olds; Aboody, Yousif, Sheskin & Keil, 2022; Kominsky et al., 2021; Sheskin & Keil, 2018). In addition, the experimenter emphasized their shirt color when the puppets were introduced, saying, “This is Sam, wearing a green shirt. And this is Max, wearing a blue shirt.” Similarly, when asking the test questions, the experimenter emphasized the options by referring to each agent's shirt color. For the *tell* test question, the experimenter said, “So, we haven't asked our friends about this cup yet. But guess what! You can tell one of our friends the name of the animal underneath this cup. You can tell our friend in the blue shirt, or you can tell our friend in the green shirt. Which one of our friends do you want to tell? The one in the blue shirt, or the one in the green shirt?” And for the *prior knowledge* test question, the experimenter said, “And can you tell me: which one of our friends peeked? The one in the blue shirt, or the one in the green shirt?”

6.2. Results

Results were coded in the same way as Experiments 1 and 2a, according to identical pre-registered exclusion criteria (but with the addition that coders would also check for any internet connectivity issues). 88.9% of participants were video or audio taped ($n = 64/72$). 8 participants were recruited but not included in the study due to experimenter error ($n = 4$), because the participant interfered with the study (by revealing the contents of the cups to the puppets; $n = 1$ in-person participant), because they did not complete the experiment ($n = 1$), because they did not answer a test question within 30s ($n = 1$), or because they had already participated in the past ($n = 1$).

When asked whom they wanted to inform, 42.2% of participants selected the predicting puppet, and 57.8% selected the observing puppet. These proportions are not reliably different from chance ($n = 27$ of 64 chose to inform the predicting puppet; 95% CI: 29.7–54.7). This chance performance, however, was not due to a failure to track or understand the experiment: 73.4% of participants judged that the predicting puppet had peeked ($n = 47$ of 64; 95% CI: 62.5–84.4), replicating our findings from Experiments 1 and 2a. These results suggest that children's inference that the predictor had peeked and was therefore knowledgeable also allowed them to track the limits of what the predictor would know.

Additionally, an exploratory response time analysis revealed no difference in response time between the two test questions (*tell* question, $M = 2.57$ s; *prior knowledge* question, $M = 1.88$ s; $t(56) = 0.87$, $p = 0.39$,

$d = 0.12$ by paired t-test). This difference is consistent with the possibility that, by replacing the unanswerable generalization question from Experiment 2a with a conceptually similar prompt that had no wrong answer, participants' uncertainty was reduced (see Fig. 3).

To test for effects of age and question type, we conducted a Bayesian mixed-effects logistic regression predicting performance as a result of age, question type, and their interaction (with random intercepts by participant; regression not pre-registered). This regression revealed a main effect for question type, with participants more likely to select the predictor when asked who peeked than when deciding whom to inform ($\beta = 1.5$, 95% CI: 0.72–2.30). This is consistent with our main results, showing that children were reliably more likely to select the predictor when asked who peeked, relative to when they were asked who they'd like to inform. In addition, the regression revealed no main effect of age, and no interaction between age and question type (all $|\beta|$'s between [0.71, 1.33]; all 95% CI's crossing 0, lower bound $\leq$ –0.13, upper bound $\geq$ 0.28), suggesting that performance did not differ with age.

Finally, we ran an additional (not pre-registered) analysis to explore the effect of online vs in-person testing. A Bayesian mixed-effects logistic regression predicting children's preference as a function of testing format (in-person vs. online), test question (who peeked vs. whom to inform), and their interaction (with random intercepts by participant), found no main effect of testing format (in-person vs. online; $\beta = 1.06$, 95% CI: -0.33–2.57). Consistent with our main results, the regression revealed a main effect of test question (with participants more likely to select the predictor when asked who peeked vs. when asked who they wanted to inform; $\beta = 3.39$, 95% CI: 1.46–5.47). Finally, the regression also revealed an interaction between testing format and test question, with participants tested online more likely to judge that the observer peeked ($\beta = -2.39$, 95% CI: -4.83 to -0.27), and more likely to want to inform the predictor ($\beta = 2.40$, 95% CI: 0.25–4.90; see Supplemental Materials for regression tables).

Although in-person participants performed better on the "who peeked" test question ($n = 13$ of 15, or 86.7%), most online participants also selected the predictor in response to this test question ($n = 34$ of 49, or 69.4%), with the only difference being in the strength of the effect size (rather than a qualitatively different pattern of responses). Similarly, although online participants were more likely to inform the predictor ($n = 23$ of 49, or 46.9%; as compared to in-person participants, $n = 4$ of 15, or 26.7%), note that only 15 of our 64 final participants were run in-person; thus, it is unclear whether these differences would be reliable given a larger in-person sample. Even if participants were truly better able to track the limits of the puppets' knowledge in our online task (although again, it is unclear whether this would be the case given a larger in-person sample), this would still serve as evidence that young children are capable of doing so.

7. Discussion, Experiments 1-2b

Experiment 1 showed that children's inferences about knowledge cannot be reduced to a simple sensitivity to superficial cues. Instead, children appeared to infer knowledge by considering what mental states best explained each agent's behavior. Next, Experiment 2a raised the possibility that children assumed the predictor was generally knowledgeable, rather than tracking what this agent was likely to know and not know. This was evidenced by the fact that children judged the predicting agent would know what was under the final cup, even after its contents had been replaced out of his view. This is consistent with the possibility that participants did not track the contents of each agent's knowledge representations, simply representing the predictor as the more knowledgeable of the two—and if so, children may have relied on a similar inference to answer the generalization question in Experiment 1.

However, it is also possible that participants were uncertain how to answer this question only in Experiment 2a, because both agents in this task were actually ignorant (whereas in Experiment 1, the predicting agent presumably still knew what was under the final cup, as its contents were not changed). This may have caused children in Experiment 2a to settle for the predictor despite realizing he was ignorant (after all, neither agent knew what was under the cup, but the predictor at least knew what was there at one point). Consistent with this possibility, an exploratory analysis revealed that participants were significantly slower at answering the generalization question than the prior knowledge question; in contrast, participants in Experiment 1 showed no difference in response time between the two questions. This raised the possibility that participants were unsure how to answer the generalization question.

Experiment 2b sought to explore this possibility. The procedure of Experiment 2b was identical to Experiment 2a, with the difference that, after switching the contents of the cup, children were given the opportunity to tell one of the puppets what was underneath (rather than asking them who was knowledgeable). If children truly believed that the predictor knew what was under the cup (as their responses in Experiment 2a suggest), they should prefer to inform the observing agent. But if participants believed that both agents were ignorant (as their response times in Experiment 2a suggest), they should show no preference over whom to inform. Our results are consistent with the latter: participants identified that the predictor had peeked under the cups, but they were equally likely to inform either agent. Additionally, as in Experiment 1, participants showed no difference in response time between the two test questions. These results are consistent with the idea that children tracked the contents of the predicting agent's knowledge, but were uncertain how to answer the generalization test question in Experiment 2a (perhaps because they realized both agents were ignorant).

At the same time, children's responses in Experiment 2b do not provide conclusive evidence that children track the limits of what others know, as there are alternative explanations for these results. For instance, participants may have been split on whom to inform because some children were motivated to punish the predicting agent (who had snuck out to peek under the cups, and thus could have been interpreted as naughty[1]), and some motivated to correct his false belief, or to affiliate with a more knowledgeable agent. We return to this point in the General Discussion.

8. Experiment 3

Experiments 1-2b suggest that children can infer knowledge from agents' actions in the absence of superficial cues, and that their epistemic attributions may track what agents do or do not know. However, participants in the prior experiments were asked to make epistemic judgments explicitly. In Experiment 3, we test whether these capacities underlie preschoolers' decisions about whom to trust when two agents provide conflicting testimony.

8.1. Method

8.1.1. Participants

32 four- and five-year-olds (mean age: 5.02 years, range: 4.14–5.99 years; $n = 16$ participants per age group) were recruited at a children's museum. Because we asked only one test question, it was not necessary to increase our sample to account for potential test question order effects. Seven additional participants were recruited but not included in the study (these participants were also recruited at children's museum; see Results).

8.1.2. Stimuli

Materials were identical to those of Experiment 1.

[1] We thank an anonymous reviewer for raising this possibility.

8.1.3. Procedure

The task began in the same way as Experiment 1, with the difference that participants were not shown the contents of the cups at the beginning of the task. After the two puppets had interacted with the first two cups, the experimenter pointed to the third cup, and said, "Well, we didn't ask our friends about this cup yet. So let's ask both of our friends about what's under this cup." One puppet said, "There's a bear under this cup," and one said "There's a squirrel under this cup" (randomizing which puppet spoke first, and the animal they claimed was under the cup). Finally, participants were asked "Can you tell me: what animal is under this cup?" After the test question, participants were asked the same memory check questions from Experiment 1 (order randomized).

8.2. Results

Results were coded in the same way as Experiments 1-2b according to identical pre-registered criteria. 97% of participants were video or audio taped ($n = 38/39$). Seven participants were recruited but not included in the study because the participant did not answer the test question (despite prompting; $n = 5$), due to experimenter error ($n = 1$), or because the participant lifted a cup and revealed its contents during the puppet show ($n = 1$). 75% of participants endorsed the predicting agent's testimony (24 of 32; 95% CI: 59.4–90.6), suggesting that when deciding whom to trust, children independently used the agent's behavior to infer knowledge—even when they were not explicitly prompted to do so. See Fig. 2. A Bayesian mixed-effects logistic regression revealed no effect of age ($\beta = 0.54$, 95% CI: −0.87–1.98; regression not pre-registered; see Supplemental Materials).

9. General discussion

The capacity to determine and track what people know is critical for social life, from everyday communication (Bohn & Köymen, 2018) to social learning (Harris, 2012). Here we explored the representations and inferences that preschoolers use to navigate the social world, focusing on three questions. First, do preschoolers attribute knowledge through a sensitivity to simple cues (such as assuming that accuracy inevitably implies knowledge) or via mental-state reasoning, where children consider what epistemic states best explain someone's observed behavior? Second, when children attribute knowledge, do these representations include expectations of what an agent may or may not know? Or are they coarser attributions (e.g., inferences that an agent is more knowledgeable without tracking the limits of their knowledge)? Finally, do children make such inferences when deciding whom to trust—or only when explicitly prompted to make epistemic judgments? To test these questions, we presented children with simple events where two agents took identical actions to produce identical outcomes. By varying the order of agents' actions, we were able to create situations where children could distinguish the agents' epistemic states only if they reasoned about which mental states best explained each agent's behavior.

Experiment 1 showed that four- and five-year-olds distinguish between accurate predictions and accurate observations, and use this distinction to infer agents' causal history (who peeked?) and epistemic states (who knows what's under a new cup?). Experiments 2a and 2b provided some initial evidence suggesting that children's beliefs about how agents gained their knowledge may allow them to determine the limits of what an agent knows. Finally, Experiment 3 showed that children make these inferences independently when deciding whom to trust—inferring that the predictor was knowledgeable from his actions even without being asked to make an epistemic inference, and using this judgment when deciding whom to believe.

Our results are consistent with related work showing that, by age four, children recognize that agents who repeat facts that they heard from others, or agents that simply ask accurate questions, are not necessarily knowledgeable (Einav & Robinson, 2011; Luchkina, Sobel, & Morgan, 2018; Luchkina, Morgan, Williams, & Sobel, 2020). This work shows that children disregard accuracy when it doesn't emerge from an independently-produced statement (i.e., stated declaratively rather than as a question, and produced without help from others). Our work goes beyond previous research by testing whether such behavior reflects a causal understanding of how knowledge relates to accuracy, or a simpler belief that independently-produced accurate statements imply knowledge. Our work contributes to this literature by providing evidence that children's inferences are sensitive to subtle differences in behavior that can reveal agents' knowledge when analyzed through a causal mental model of others' epistemic states.

At the same time, our results do not imply that children never rely on heuristics when reasoning about knowledge. Related work shows that children select informants based on a variety of cues, such as agents' accent, familiarity, and dominance (see Introduction); young children will even favor familiar agents over accurate ones (Corriveau & Harris, 2009). This work opens three possibilities. A first possibility is that young children are simultaneously motivated to learn from knowledgeable agents and to affiliate with in-group members (Dunham, Baron, & Banaji, 2008). If so, children's apparent reliance on heuristics might reflect intergroup preferences rather than epistemic reasoning. A second possibility is that children do believe that seemingly irrelevant features, such as familiarity or accent, reveal whom to trust (perhaps because familiar people have had a long track record of generally being accurate, or because people with similar accents might have information that is relevant to their group; Begus, Gliga, & Southgate, 2016). Finally, a third possibility is that children begin to navigate the epistemic world by relying on simple heuristics. Such heuristics may be replaced or complemented by richer mental-state reasoning as children's Theory of Mind develops. From this standpoint, four- and five-year-olds may be at an intermediate stage of development where they need not necessarily rely on simple heuristics (vs. mental-state reasoning), but still rely on some superficial cues when these are available. It is also possible that, through mental state reasoning, children develop more effective shortcuts or rules to rely on (for instance, codifying intuitions grounded in Theory of Mind, or identifying the most efficient heuristics for different tasks; e.g., Horn, Ruggeri, & Pachur, 2016)—but still relying on their causal model of other minds to attribute knowledge in situations where simple rules cannot do. These are questions we hope to explore in future work.

Our work also opens a new question. What are the underlying computations that led children to infer that the predictor was knowledgeable? While this is an open question, our results can be readily explained by expanding current theories of mental-state inference. Research suggests that children and adults infer mental states through an assumption that agents maximize their subjective utilities—the difference between the cost they incur and the reward they obtain (Aboody et al., 2021; Gergely & Csibra, 2003; Jara-Ettinger, Gweon, Schulz, & Tenenbaum, 2016; Jern, Lucas, & Kemp, 2017). Our findings are consistent with this framework: When the goal is to provide accurate information, a knowledgeable agent can maximize their utilities by providing the correct answer (and thus incurring no unnecessary costs). An ignorant agent, by contrast, ought to incur a cost to obtain the information needed to fulfill their goal (provided that the cost of getting the information does not outweigh the reward associated with getting things right). Note, however, that in our task, the predictor lifted the cup after they stated what was inside it. Under this framework, this additional cost can be interpreted as evidence that the predictor's goal was not only to provide accurate information, but also to prove that the information was accurate. While more research is needed to test this prediction in children, in an additional study we have shown that adults indeed interpret the observer's cup-lifting as checking (i.e., providing information for the self) and the predictor's cup-lifting as showing (i.e., providing information to others; see Supplemental Materials).

Finally, our work has two chief limitations. First, our experiments used a two-alternative forced choice paradigm, which limits our ability to interpret children's choices. This limitation is particularly important for Experiment 2a, where children's pattern of responses was ambiguous. Although our exploratory reaction time analysis (Section 5.2.1) and Experiment 2b suggest that children in this experiment may have tracked the limits of agents' knowledge, our work would have benefitted from more nuanced measures, such as explicit measures of confidence (see Lapidow, Killeen, & Walker, 2022) or open-ended explanation prompts (although note that young preschoolers often struggle to produce codable explanations in experimental contexts; Legare & Lombrozo, 2014; Walker, Lombrozo, Legare, & Gopnik, 2014; Walker, Bonawitz, & Lombrozo, 2017). Future work should investigate whether more graded measures can capture nuances in children's understanding of how knowledge relates to action.

Our second related limitation is that our work did not conclusively show whether children can track the limits of what agents know. While Experiment 2a suggested that children may not track the contents of others' knowledge representations, Experiment 2b suggested that they might. Coding response times across Experiments 1-2b provided some converging evidence that participants in Experiment 2a may have been uncertain how to respond when asked which of two ignorant agents would know what was under the final cup: participants in Experiment 2a took significantly longer to answer this test question (as opposed to the prior knowledge test question), whereas participants' response times did not differ between test questions in Experiments 1 and 2b. Importantly, these analyses were exploratory, and a difference in response time does not definitively indicate uncertainty or confusion. While it is unclear why participants in Experiment 2b did not prefer to inform the observing agent if children truly believe accurate predictions imply knowledge (without tracking precisely what others know) it is possible that these participants were divided for other reasons. For instance, some participants may have preferred to avoid the predicting agent, who could be considered "naughty" for peeking, whereas others may have preferred to correct the predicting agent's false belief, or even to affiliate with this knowledgeable agent. Note that these explanations are not necessarily inconsistent with the idea that children track the contents and limits of epistemic states (in fact, a motivation to correct a false belief depends on this capacity). Nevertheless, we cannot rule out alternate interpretations of these results. Future work should test whether children track the precise contents of others' epistemic representations, or whether they make broader epistemic attributions (e.g., assuming that accurate predictions imply greater knowledge in a domain).

9.1. Conclusion

At an age where we gain knowledge primarily by learning from others, our results show that children do not select informants through coarse rules or heuristics. Instead, children navigate the epistemic world by observing others' behavior, and reasoning about the events and mental states likely to have caused that behavior. Our work highlights children's early inferential capacities, and advances our understanding of how these inferences work within children's developing Theory of Mind.

Credit author statement

Rosie Aboody: Conceptualization, Methodology, Investigation, Formal Analyses, Writing - original draft, Writing - review & editing, Visualization. **Holly Huey:** Methodology, Investigation, Writing - original draft, Writing - review & editing. **Julian Jara-Ettinger:** Conceptualization, Methodology, Funding acquisition, Supervision, Writing - original draft, Writing - review & editing.

Acknowledgments

We thank the Boston Children's Museum, the Peabody Natural History Museum, and the families who participated in this research. We thank Colin Jacobs for assistance with data collection. We thank Madison Flowers, Caiqin Zhou, Colin Jacobs, Gwyneth Heuser, Eleanor Iskander, Lauren Barragan, and Jenna Landy for help with coding. This work was supported by the Simons Center for the Social Brain, and NSF awardBCS-2045778 awarded to JJE. This material is based upon work supported by the Center for Brains, Minds, and Machines (CBMM), funded by NSF-STC award CCF-1231216.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.cognition.2022.105212.

References

Aboody, R., Huey, H., & Jara-Ettinger, J. (2022, May 2). *Preschoolers decide who is knowledgeable, who to inform, and who to trust via a causal understanding of how knowledge relates to action.* 10.17605/OSF.IO/MHCVT.
Aboody, R., Yousif, S. R., Sheskin, M., & Keil, F. (2022). Says who? Children consider informants' sources when deciding whom to believe. *Journal of Experimental Psychology: General*.
Aboody, R., Zhou, C., Flowers, M., & Jara-Ettinger, J. (2019). Ignorance = doing what is reasonable: Children expect ignorant agents to act based on prior knowledge. In A. K. Goel, C. M. Seifert, & C. Freksa (Eds.), *Proceedings of the 41st annual conference of the cognitive science society* (pp. 1297–1303). Montreal, CA: Cognitive Science Society.
Aboody, R., Zhou, C., & Jara-Ettinger, J. (2021). In pursuit of knowledge: Preschoolers expect agents to weigh information gain and information cost when deciding whether to explore. *Child Development, 92*(5), 1919–1931.
Begus, K., Gliga, T., & Southgate, V. (2016). Infants' preferences for native speakers are associated with an expectation of information. *Proceedings of the National Academy of Sciences, 113*(44), 12397–12402.
Bernard, S., Castelain, T., Mercier, H., Kaufmann, L., Van der Henst, J. B., & Clément, F. (2016). The boss is always right: Preschoolers endorse the testimony of a dominant over that of a subordinate. *Journal of Experimental Child Psychology, 152*, 307–317.
Birch, S. A., 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.
Bohn, M., & Köymen, B. (2018). Common ground and development. *Child Development Perspectives, 12*(2), 104–108.
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.
Brosseau-Liard, P. E., & Birch, S. A. (2011). Epistemic states and traits: Preschoolers appreciate the differential informativeness of situation-specific and person-specific cues to knowledge. *Child Development, 82*(6), 1788–1796.
Butler, L. P., Gibbs, H. M., & Tavassolie, N. S. (2020). Children's developing understanding that even reliable sources need to verify their claims. *Cognitive Development, 54*, Article 100871.
Butler, L. P., Schmidt, M. F., Tavassolie, N. S., & Gibbs, H. M. (2018). Children's evaluation of verified and unverified claims. *Journal of Experimental Child Psychology, 176*, 73–83.
Chen, Y., Su, Y., & Wang, Y. (2015). Young children use the "ignorance= getting it wrong" rule when predicting behavior. *Cognitive Development, 35*, 79–91.
Chow, V., Poulin-Dubois, D., & Lewis, J. (2008). To see or not to see: Infants prefer to follow the gaze of a reliable looker. *Developmental Science, 11*(5), 761–770.
Clément, F., Koenig, M., & Harris, P. (2004). The ontogenesis of trust. *Mind & Language, 19*(4), 360–379.
Cohen, J. (1994). The earth is round (p<. 05). *American Psychologist, 49*(12), 997.
Corriveau, K., & Harris, P. L. (2009). Choosing your informant: Weighing familiarity and recent accuracy. *Developmental Science, 12*(3), 426–437.
Corriveau, K. H., Fusaro, M., & Harris, P. L. (2009). Going with the flow: Preschoolers prefer nondissenters as informants. *Psychological Science, 20*(3), 372–377.
Corriveau, K. H., Kinzler, K. D., & Harris, P. L. (2013). Accuracy trumps accent in children's endorsement of object labels. *Developmental Psychology, 49*(3), 470.
Cumming, G. (2014). The new statistics: Why and how. *Psychological Science, 25*(1), 7–29.
Dunham, Y., Baron, A. S., & Banaji, M. R. (2008). The development of implicit intergroup cognition. *Trends in Cognitive Sciences, 12*(7), 248–253.
Einav, S. (2014). Does the majority always know best? Young children's flexible trust in majority opinion. *PloS one, 9*(8), Article e104585.
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*(10), 1250–1253.
Elashi, F. B., & Mills, C. M. (2014). Do children trust based on group membership or prior accuracy? The role of novel group membership in children's trust decisions. *Journal of Experimental Child Psychology, 128*, 88–104.
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.
Fusaro, M., & Harris, P. L. (2008). Children assess informant reliability using bystanders' non-verbal cues. *Developmental Science, 11*(5), 771–777.
Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naıve theory of rational action. *Trends in Cognitive Sciences, 7*(7), 287–292.

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., & Meltzoff, A. N. (1997). *Words, thoughts, and theories*. MIT Press.
Harris, P. L. (2012). *Trusting what you're told*. Harvard University Press.
Hetherington, C., Hendrickson, C., & Koenig, M. (2014). Reducing an in-group bias in preschool children: The impact of moral behavior. *Developmental Science, 17*(6), 1042–1049.
Horn, S. S., Ruggeri, A., & Pachur, T. (2016). The development of adaptive decision making: Recognition-based inference in children and adolescents. *Developmental Psychology, 52*(9), 1470.
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.
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.
Jaswal, V. K., McKercher, D. A., & VanderBorght, M. (2008). Limitations on reliability: Regularity rules in the English plural and past tense. *Child Development, 79*(3), 750–760.
Jaswal, V. K., & Neely, L. A. (2006). Adults don't always know best: Preschoolers use past reliability over age when learning new words. *Psychological Science, 17*(9), 757–758.
Jern, A., Lucas, C. G., & Kemp, C. (2017). People learn other people's preferences through inverse decision-making. *Cognition, 168*, 46–64.
Kinzler, K. D., Corriveau, K. H., & Harris, P. L. (2011). Children's selective trust in native-accented speakers. *Developmental Science, 14*(1), 106–111.
Koenig, M. A. (2012). Beyond semantic accuracy: Preschoolers evaluate a speaker's reasons. *Child Development, 83*(3), 1051–1063.
Koenig, M. A., Clément, F., & Harris, P. L. (2004). Trust in testimony: Children's use of true and false statements. *Psychological Science, 15*(10), 694–698.
Koenig, M. A., Cole, C. A., Meyer, M., Ridge, K. E., Kushnir, T., & Gelman, S. A. (2015). Reasoning about knowledge: Children's evaluations of generality and verifiability. *Cognitive Psychology, 83*, 22–39.
Koenig, M. A., & Echols, C. H. (2003). Infants' understanding of false labeling events: The referential roles of words and the speakers who use them. *Cognition, 87*(3), 179–208.
Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development, 76*(6), 1261–1277.
Koenig, M. A., & Woodward, A. L. (2010). 24-month-olds' sensitivity to the prior inaccuracy of the source: Possible mechanisms. *Developmental Psychology, 46*(4), 815.
Kominsky, J. F., Gerstenberg, T., Pelz, M., Sheskin, M., Singmann, H., Schulz, L., & Keil, F. C. (2021). The trajectory of counterfactual simulation in development. *Developmental Psychology, 57*(2), 253.
Kominsky, J. F., Langthorne, P., & Keil, F. C. (2016). The better part of not knowing: Virtuous ignorance. *Developmental Psychology, 52*(1), 31.
Kushnir, T., & Koenig, M. A. (2017). What I don't know won't hurt you: The relation between professed ignorance and later knowledge claims. *Developmental Psychology, 53*(5), 826.
Kushnir, T., Vredenburgh, C., & Schneider, L. A. (2013). "Who can help me fix this toy?" the distinction between causal knowledge and word knowledge guides preschoolers' selective requests for information. *Developmental Psychology, 49*(3), 446.
Lackey, J. (2010). *Learning from words: Testimony as a source of knowledge*. OUP Oxford.
Lapidow, E., Killeen, I., & Walker, C. M. (2022). Learning to recognize uncertainty vs. recognizing uncertainty to learn: Confidence judgments and exploration decisions in preschoolers. *Developmental Science, 25*(2), Article e13178.
Legare, C. H., & Lombrozo, T. (2014). Selective effects of explanation on learning during early childhood. *Journal of Experimental Child Psychology, 126*, 198–212.
Lockhart, K. L., Goddu, M. K., Smith, E. D., & Keil, F. C. (2016). What could you really learn on your own?: Understanding the epistemic limitations of knowledge acquisition. *Child Development, 87*(2), 477–493.
Lucas, A. J., Burdett, E. R., Burgess, V., Wood, L. A., McGuigan, N., Harris, P. L., & Whiten, A. (2017). The development of selective copying: Children's learning from an expert versus their mother. *Child Development, 88*(6), 2026–2042.
Luchkina, E., Morgan, J. L., Williams, D. J., & Sobel, D. M. (2020). Questions can answer questions about mechanisms of preschoolers' selective word learning. *Child Development, 91*(5), e1119–e1133.
Luchkina, E., Sobel, D. M., & Morgan, J. L. (2018). Eighteen-month-olds selectively generalize words from accurate speakers to novel contexts. *Developmental Science, 21*(6), Article e12663.
Lutz, D. J., & Keil, F. C. (2002). Early understanding of the division of cognitive labor. *Child Development, 73*(4), 1073–1084.
Mangardich, H., & Sabbagh, M. A. (2018). Children remember words from ignorant speakers but do not attach meaning: Evidence from event-related potentials. *Developmental Science, 21*(2), Article e12544.
Mills, C. M., Legare, C. H., Bills, M., & Mejias, C. (2010). Preschoolers use questions as a tool to acquire knowledge from different sources. *Journal of Cognition and Development, 11*(4), 533–560.
Mills, C. M., Legare, C. H., Grant, M. G., & Landrum, A. R. (2011). Determining who to question, what to ask, and how much information to ask for: The development of inquiry in young children. *Journal of Experimental Child Psychology, 110*(4), 539–560.
Nurmsoo, E., & Robinson, E. J. (2009a). Children's trust in previously inaccurate informants who were well or poorly informed: When past errors can be excused. *Child Development, 80*, 23–27.
Nurmsoo, E., & Robinson, E. J. (2009b). Identifying unreliable informants: Do children excuse past inaccuracy? *Developmental Science, 12*, 41–47.
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.
Pasquini, E. S., Corriveau, K. H., Koenig, M., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology, 43*(5), 1216.
Pillow, B. H. (1993). Preschool children's understanding of the relationship between modality of perceptual access and knowledge of perceptual properties. *British Journal of Developmental Psychology, 11*(4), 371–389.
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.
Robinson, E. J., Butterfill, S. A., & Nurmsoo, E. (2011). Gaining knowledge via other minds: Children's flexible trust in others as sources of information. *British Journal of Developmental Psychology, 29*(4), 961–980.
Robinson, E. J., & Whitcombe, E. L. (2003). Children's suggestibility in relation to their understanding about sources of knowledge. *Child Development, 74*(1), 48–62.
Ronfard, S., & Corriveau, K. H. (2016). Teaching and preschoolers' ability to infer knowledge from mistakes. *Journal of Experimental Child Psychology, 150*, 87–98.
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.
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.
Sabbagh, M. A., & Shafman, D. (2009). How children block learning from ignorant speakers. *Cognition, 112*(3), 415–422.
Scofield, J., & Behrend, D. A. (2008). Learning words from reliable and unreliable speakers. *Cognitive Development, 23*(2), 278–290.
Sheskin, M., & Keil, F. (2018, December 30). *TheChildLab.com a video chat platform for developmental research*. https://doi.org/10.31234/osf.io/rn7w5
Sobel, D. M., & Kushnir, T. (2013). Knowledge matters: How children evaluate the reliability of testimony as a process of rational inference. *Psychological Review, 120*(4), 779.
Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G., & Wilson, D. (2010). Epistemic vigilance. *Mind & Language, 25*(4), 359–393.
Surian, L., Caldi, S., & Sperber, D. (2007). Attribution of beliefs by 13-month-old infants. *Psychological Science, 18*(7), 580–586.
Tenney, E. R., Small, J. E., Kondrad, R. L., Jaswal, V. K., & Spellman, B. A. (2011). Accuracy, confidence, and calibration: How young children and adults assess credibility. *Developmental Psychology, 47*(4), 1065.
Walker, C. M., Bonawitz, E., & Lombrozo, T. (2017). Effects of explaining on children's preference for simpler hypotheses. *Psychonomic Bulletin & Review, 24*(5), 1538–1547.
Walker, C. M., Lombrozo, T., Legare, C. H., & Gopnik, A. (2014). Explaining prompts children to privilege inductively rich properties. *Cognition, 133*(2), 343–357.
Wellman, H. M. (2014). *Making minds: How theory of mind develops*. Oxford University Press.
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.
Zmyj, N., Buttelmann, D., Carpenter, M., & Daum, M. M. (2010). The reliability of a model influences 14-month-olds' imitation. *Journal of Experimental Child Psychology, 106*(4), 208–220.