

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Young children can identify knowledgeable speakers from their causal influence over listeners

Permalink

<https://escholarship.org/uc/item/9qh630dn>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors

Chuey, Aaron
Sparks, Robert
Gweon, Hyowon

Publication Date

2023

Peer reviewed

Young children can identify knowledgeable speakers from their causal influence over listeners

Aaron Chuey, Robert Sparks, & Hyowon Gweon

chuey@stanford.edu, bsparks@stanford.edu, gweon@stanford.edu

Department of Psychology, Stanford University

Abstract

Prior work demonstrates an early-emerging understanding of how speakers can alter listeners' minds and actions. Yet, an abstract understanding of communication entails more than *forward* inferences about its influence on the listener; it also supports *inverse* inferences about the speaker based on its causal influence over the listener. Can children reason about the minds of speakers based on their causal influence over listeners? Across three studies, children viewed two communicative exchanges where a listener attempted to activate a toy; we manipulated when speakers communicated (Exp.1), how listeners' subsequent actions changed (Exp.2), and whether speakers spoke or sneezed (Exp.3). By 5 years of age, children inferred the speaker who appeared to cause the listener to succeed was more knowledgeable, but only when they produced speech. These results suggest children can reason causally about the sources of communication, identifying knowledgeable speakers based on their influence over a listener's actions and their outcomes.

Keywords: Communication, Cognitive development, Social cognition, Theory of Mind

Introduction

Communication allows one person to influence another person's mind and behaviors. Based on how two people interact and communicate, we can make rich inferences about their minds and future actions. While these inferences are relatively easy when we know who said what to whom, in some cases, we can make reasonable guesses even without actually knowing the content of the communicative exchange.

Imagine you ask your friend Gus "what is the capital of Iceland?" but he fails to provide an answer, returning a blank stare. However, after your other friend Barry whispers something in his ear, Gus immediately says "Reykjavik", accurately answering the question. Observing this instance allows you to understand not only that Gus knows the capital of Iceland, but also that Barry probably knows, too. In fact, Barry was likely the source of Gus' knowledge!

Note that you had few direct cues to Barry's knowledge; you couldn't hear what Barry actually told Gus, Barry didn't explicitly claim he knew the answer, and in principle he could have whispered anything to Gus. Nonetheless, the combination of Gus' change in behavior (i.e., from silence to providing the right answer) and Barry's whisper that preceded the outcome was enough for you to guess what Barry knows. In other words, you inferred that Barry *caused* the change in Gus' behavior by imparting his knowledge to Gus.

This example captures a particularly interesting and challenging case of epistemic inference: using an agent's communicative influence on others to infer what that agent knows. This is possible because you have an abstract understanding of communication that supports not just forward inferences (how speakers influence listeners) but also inverse inferences; you can recover the speaker's mental states and even the content of the conversation based on changes in the listener's behaviors.

Although a speaker's knowledge can easily be gleaned from what they say or do, we often encounter situations where the content of communication is inaccessible due to perceptual (e.g., noise, whispering) or cognitive (e.g., speaking a different language) limitations. Inferring what others know in the absence of any spoken content, however, is far from trivial; in order to use a speaker's influence over the listener's actions to recover the speaker's knowledge, one must understand how information flows from one person to another and gives rise to the recipient's behaviors. When do children acquire such an understanding, and how does it support rich epistemic inferences?

An understanding of how communication transmits information between agents develops throughout early childhood. For instance, infants expect a listener to reach for a speaker's preferred object (Vouloumanos et al., 2014) or mimic their actions (Vouloumanos et al., 2012) when they speak, but not cough. They can even identify informative exchanges based on the pattern of auditory tone sequences (e.g., variable, unpredictable, patterns are more suggestive of communication, (see Tauzin & Gergely, 2018)). As their language comprehension develops, children begin to form more concrete expectations about the relationship between communication and knowledge. Toddlers expect listeners to update their beliefs based on the contents of a speaker's utterances (Jin et al., 2019; Song et al., 2008), and by around 4-years, children understand that environmental factors, such as auditory noise, can disrupt communication and its ability to alter a listener's beliefs (Chuey et al., 2022). These abilities develop in parallel to a broader understanding of how information influences agents' minds and behaviors, enabling them to predict how others will act (Kim & Song, 2015; Knudsen & Liszkowski, 2012; Luo & Baillargeon, 2007; Luo & Johnson, 2009; Surian et al., 2007), react (Moll et al., 2006; Tomasello & Haberl, 2003; Scott, 2017; Wu et al., 2018), and form beliefs (Aboody

et al., 2022; Baron-Cohen et al., 1985; Hogrefe et al., 1986; Magid et al., 2018) based on what they perceive.

These findings demonstrate an early understanding of how and when communication affects recipients of information (i.e., listeners) based on the presence, pattern, and content of communicative signals. However, open questions remain regarding how children reason about the sources of information (i.e., speakers), especially when the content of their speech is unavailable. In the Reykjavik example, you did not hear what Barry whispered to Gus. Nonetheless, you gathered that Barry's communication was effective (i.e., caused Gus to give the correct answer) and inferred Barry's knowledge. Can we use changes in listeners' actions and outcomes to infer what speakers know in the absence of explicit evidence? Below we consider at least two ways in which effective communication could cause changes in listeners' actions and outcomes.

First, effective communication is more likely to cause a listener to perform a successful action rather than a failed action. Thus, when a listener initially fails and ultimately succeeds, a speaker who communicated before the successful action would appear more knowledgeable than a speaker who communicates before the failure. Given that attribution of causality is modulated by the outcome as well as the temporal gap between causes and effects (e.g., Leslie & Keeble, 1987; Michotte, 2017), manipulating *when* a speaker communicates relative to the outcome of a listener's actions (i.e., success) might also modulate our inferences about the perceived influence of the speaker over the listener, and subsequently the degree to which people attribute knowledge to the speaker.

Second, even when communication precedes a positive outcome, we might have different intuitions about the speaker's causal role depending on the changes in the listener's actions alone. In our example, Barry's whisper preceded a behavior-producing the correct answer "Reykjavik"—which is not something he would have produced otherwise. If, however, the whisper preceded a behavior that was going to happen anyway (i.e., Gus was actually just about to say "Reykjavik" when Barry whispered in his ear), it would be less clear whether Barry actually knew the capital of Iceland; perhaps he did, but it is also possible that what he whispered to Gus was something completely irrelevant to the question. Just as we engage in counterfactual reasoning to make causal attributions in physical scenarios (Gerstenberg et al., 2021), such cases may also involve reasoning about whether a speaker's communication caused the listener to succeed by considering how a listener's actions would have unfolded without the speaker's communication.

The current studies ask whether preschool-aged children can identify and use speakers' causal influence over the actions and outcomes of listeners to infer what speakers know. Prior work suggests that even young children can reason about causality based on temporal contiguity (e.g., Leslie & Keeble, 1987), and by four years of age, engage in preliminary forms of counterfactual or hypothetical reasoning

(Kominsky et al., 2021; Nyhout & Ganea, 2019). Thus, it is possible that preschool-aged children can reason about speakers' epistemic states by considering their causal influence on listeners. However, children also struggle with genuine counterfactual reasoning (Gerstenberg, 2022) and tracking the knowledge states of multiple agents (Hogrefe et al., 1986). Given these findings, we recruited our participants from a relatively wide age range (3-, 4-, and 5-year-olds) to identify potential developmental change.

Across three studies, children viewed two scenarios where an agent (listener) unsuccessfully attempted to activate a toy before ultimately succeeding. At some point during each scenario, another agent (speaker) entered the scene and said something to the listener in an unknown language, rendering the semantic content of the speech unavailable to participants. We manipulated the timing of the speaker's actions (Studies 1 & 3) or the degree of change in the listener's actions (Studies 2 & 3) such that in one scenario, the speaker appeared to have caused the listener to successfully activate the toy (Effective Speaker) but not in the other (Ineffective Speaker). Afterwards, we asked children which speaker possessed prior knowledge about the toy. To examine whether children's inferences were specific to communicative acts, Study 3 manipulated whether both speakers spoke (communicative) or sneezed (non-communicative).

All experimental stimuli and videos are available at https://osf.io/derxp/?view_only=d3ad5730e321405da0e5347dfb35a3f0.

Study 1

The first study examined whether children can infer what a speaker knows based on how their communication affects the the outcome of a listener's actions by controlling for the listener's actions themselves and manipulating when the speaker communicates. We predicted that when speech followed a listener's failure and preceded their success, children would judge the speaker as more knowledgeable compared to when speech preceded a listener's failure and eventual success. We preregistered our hypotheses and analysis plan at https://aspredicted.org/W23_JD3.

Methods

Participants Seventy-two children (Age: 3.0 - 5.9; Mage=4.6), 24 of each year, participated online via Zoom. An additional 4 children were tested but excluded based on preregistered exclusion criteria (2 for experimenter error, 1 for parental interference, and 1 for technical difficulties).

Materials Children watched a slide show using MS PowerPoint and the screen-sharing feature on Zoom. The slide show presented still photos and prerecorded video clips of two characters (monster puppets referred to as "wubs") and a causal toy (a colored box with two buttons, see 1. Two green wubs appeared in both scenarios as a listener, while a blue wub and an orange wub each appeared in only one of the two scenarios as a speaker (color counterbalanced).




Figure 1: Study 1 schematic of the two scenarios (left) and results (right). In both scenarios, a listener failed to activate a toy twice before succeeding on their third attempt. The key difference was whether communication occurred before the final success (Effective) or before the initial failures (Ineffective). The graph shows the proportion of children in each age group who attributed knowledge to the Effective Speaker; error bars indicate 95% bootstrapped confidence intervals.

Procedure First, the experimenter introduced children to four different colored wubs (two green, one blue, one orange) who spoke a different language, Jabberwocky, that only wubs knew. Next, the experimenter described a 2-button causal toy: “To make this toy go, you have to press both buttons on top at the same time. If you press just one of these buttons, the toy doesn’t work!” Children then saw two green wubs appear, followed by question marks above their heads; the experimenter told them: “The green wubs have never seen the toy before, so they don’t know how it works”. Next, the experimenter introduced a blue wub and an orange wub, telling children that one of them played with the toy before and knows how the toy works while the other one does not know. Children were told they would see two videos about the wubs and that their job was to figure out which wub (blue or orange) knows how the toy works.

Children then viewed two videos (order counterbalanced). In the Effective Speaker scenario, the green wub (listener) moved towards the toy, pressed each button individually, let out a grunt of frustration, and repeated these actions a second time. A second wub (speaker) then entered, uttered a short nonsense phrase, and left. Next, the listener pressed both buttons at the same time, causing the toy to ring; it then raised its hands and produced a jubilant vocalization.

The Ineffective Speaker scenario was identical to the Effective Speaker scenario, except for when the speaker appeared. After the listener positioned itself over the toy, but before it pressed a button, the speaker entered the scenario, moved towards the listener, produced a nonsense utterance, and left. The listener then proceeded with the same actions as the Effective Speaker scenario, failing to activate the toy twice before ultimately succeeding.

After the second video, children viewed a slide with the toy flanked by the blue and orange wubs. The experimenter then

reminded children that only one of the two wubs knew how the toy works, and asked them which wub knew.

Results

We first conducted a logistic regression predicting speaker choice by age (years). We found a marginal effect of age on toy choice ($\beta = .53, p = .08$) suggesting that the tendency to choose the speaker in the Effective Speaker scenario increased with age. While 3-year-olds (12/24) and 4-year-olds (13/24) did not show a preference for either speaker (Binomial tests, two-tailed; $ps > .4$), a significant majority of 5-year-olds (18/24) preferentially chose the Effective Speaker ($p = .023$). See Figure 1.

Study 2

The results of Study 1 suggest that, by 5 years of age, children can infer what a speaker knows by using the timing of the speaker’s communication, which gave the impression that the speaker influenced the listener’s actions (pressing a single button vs. both buttons) as well as their consequences (failure vs. success). Can children also identify a speaker’s influence by considering the listener’s actions alone? Study 2 investigates this question by controlling for the outcome (i.e., whether the listener’s actions following communication resulted in success) while manipulating how the listener’s course of actions changed following communication. We selectively recruited 5-year-olds given that younger children did not show above-chance performance in Study 1, and Study 2 involves an even subtler manipulation.

Methods

Participants Twenty-six children (Age: 5.1 - 5.9; Mage=5.5), participated online via Zoom. An additional 2




Figure 2: Study 2 schematic (left) and results (right). In both scenarios, a listener sequentially pressed two inert buttons, and then following communication, pressed the effective button on their third attempt. The key difference was whether the listener pressed the effective button by skipping a few buttons (Effective) or continuing in sequence (Ineffective). The graph shows the proportion of children who attributed knowledge to the Effective Speaker; error bars indicate 95% bootstrapped confidence intervals.

children were tested but excluded (1 for experimenter error, and 1 for technical difficulties).

Materials The materials and procedure were similar to Study 1, except the wubs wore different colored clothes to further help children distinguish them, and we used a different toy (a long colored box with eight buttons, see 2¹).

Procedure The procedure was similar to Study 1 except for a few differences. After being introduced to the wubs, children were introduced to an 8-button toy; the toy had a pink button that made the toy ring and seven inert buttons, one white and six green. Children then saw two green wubs appear, followed by question marks above their heads; the experimenter told them “The green wubs have never seen the toy before, so they don’t know how to make it ring, but they really want to make it ring”. Next, the experimenter introduced a blue wub and an orange wub, telling children that one of them had played with the toy and knew the pink button made the toy ring. Children were told they would see two videos about the wubs and that their job was to figure out which wub (blue or orange) knew how to make the toy ring.

Children then viewed two videos (order counterbalanced). In the Effective Speaker scenario, the green wub (listener) moved towards the left side of the toy, pressed the leftmost button, and let out a grunt of frustration. They then pressed the next button and reacted similarly. Before they were about to press the next (white, inert) button, a second wub (speaker) entered the scenario from the right, spoke, and left. The listener then moved over three buttons to the right and pressed the pink button, causing the toy to produce a ringing sound;

the listener raised its hands, producing a jubilant vocalization.

The Ineffective Speaker scenario was similar to the Effective Speaker scenario, but the listener started from the right side of the toy instead, pressing the two rightmost buttons which failed to activate the toy. As they were about to press the next (pink, effective) button, the speaker entered, produced a nonsense utterance, and left. The listener then pressed the pink button, activating the toy.

After playing the two videos, the experimenter transitioned to a slide with just the blue and orange wubs. The experimenter reminded children that only one of the wubs on screen knew that the pink button made the toy ring, before asking them which of the two wubs knew.

Results

As predicted, a significant majority of 5-year-olds (19/26) chose the speaker in the Effective Speaker scenario as the one who knew how to make the toy ring (Binominal Test, $p = .029$). See Figure 2.

Study 3

Studies 1 and 2 demonstrate that 5-year-olds can attribute knowledge to speakers based on when their communication occurred relative to the outcomes of or changes in a listener’s actions. However, it is unclear whether these inferences are unique to communicative contexts: do children treat communication (speech) as a necessary ingredient for influencing listeners? Thus, Study 3 sought to verify that these inferences were limited to communicative signals (i.e., speech) rather than any vocalizations (e.g., sneeze, cough). We preregistered our hypotheses and analysis plan at https://aspredicted.org/9TD_XK6.

¹Some children viewed videos that contained a toy with green, orange, and purple buttons instead. There was no difference in children’s performance based on which toy they saw, $p > .5$, so we collapse across them for our analysis.




Figure 3: Study 3 schematic (left) and results (right). In both scenarios, a listener pressed two inert buttons before pressing the effective button on their third attempt. The scenarios differed in two ways: whether the speaker appeared before (Ineffective) or after (Effective) the listener’s failures, and whether the listener skipped buttons (Effective) or continued in sequence (Ineffective) to press the pink button. The graph shows the proportion of children who attributed knowledge to the Effective Speaker by age (years) and condition (Communicative or Non-communicative); error bars indicate 95% bootstrapped confidence intervals.

Methods

Participants We recruited 144 children (Age: 3.0 - 5.9; Mage=4.5), 48 of each year, participated online via Zoom videoconferencing software. An additional 22 children were tested but excluded (8 for parental interference, 6 for failing to provide a clear response, 3 for experimenter error, 3 for technical difficulties, and 2 for significant lapses in attention).

Materials The materials were the same as Study 2.

Procedure The procedure was nearly identical to Study 2, except participants were assigned to either the Communicative or Non-communicative condition. In the Communicative condition, the speakers in the videos spoke Jabberwocky as before, while the speakers in the Non-communicative condition sneezed three times instead. The sequence of events in the Effective Speaker scenario was identical to Study 2. The Ineffective Speaker scenario was similar to Study 2, but the speaker entered the scene before the listener pressed the right-most (green, inert) button, spoke or sneezed, and left. Then the listener, moving right to left, pressed two inert buttons, expressing dissatisfaction after each, before pressing the pink button which caused the toy to activate. After playing the two videos, children were asked which wub (blue or orange) knew how to make the toy ring in the same way as Study 2.

Results

A logistic regression predicting speaker choice by age (years) and condition found a marginally significant interaction between condition and age, suggesting that older children were more likely to choose the Effective Speaker in the speech condition than in the sneeze condition ($\beta = .79, p = .062$). We then conducted separate logistic regressions in each age group predicting speaker choice by condition, followed by binomial tests (two-tailed against 50%). Neither 3-year-olds nor 4-year-olds showed a difference between conditions (all $ps > .5$), and their choices did not differ from chance in

either condition (all $ps > .5$). However, 5-year-olds were more likely to attribute knowledge to the Effective Speaker in the speech condition (18/24) compared to the sneeze condition (9/24), ($\beta = 1.61, p = .011$). Furthermore, their preference for the Effective Speaker was significantly greater than chance in the speech condition ($p = .023$) but did not differ from chance in the sneeze condition ($p = .31$).

Consistent with prior results, only 5-year-olds showed a clear preference for the Effective Speaker. Critically, this preference disappeared when the speaker merely produced a non-communicative vocalization (See Figure 3).

Discussion

The current work investigates children’s understanding of communication as a causal force that influences listeners and reflects the minds of speakers. By studying young children who are still developing many of the key social-cognitive skills that enable them to reason about social interactions, this research examines the developmental origins and trajectory of our understanding of communication as a means for social influence.

Across three studies, 5-year-olds, but not younger children, consistently attributed knowledge to a speaker whose communication caused a listener to succeed (Study 1) or altered the course of their actions (Study 2). Critically, these inferences were appropriately limited to communicative acts that transmit information, and disappeared when speakers produced non-communicative vocalizations (Study 3). Importantly, children made these inferences without meaningful information about the speaker’s perceptual access, prior knowledge, or the semantic content of their speech. While previous research has shown that young children, and at times infants, can identify epistemic changes in listeners (Jin et al., 2019; Song et al., 2008), the current results suggest that such epistemic reasoning goes beyond the target of communication (listener) and extends even to its source (speaker).

Children in our study faced a rather simple choice: Which of two agents knows how a toy works? This choice, however, likely reflects the results of several inferential steps: (1) identifying a change in a listener's epistemic state (e.g., going from not knowing to knowing how to activate a toy), (2) understanding that the change was caused by a speaker's communication, and (3) inferring what knowledge the speaker needed to bring about this change in the listener. However, our findings are also consistent with the possibility that children reasoned about what the *listener* thinks about the speaker; instead of directly inferring the speaker's knowledge, they might have considered whether the listener thought the speaker is knowledgeable or ignorant based on whether the listener seems to have "listened" to the speaker. This inference, while even more sophisticated, still requires reasoning about the speaker's influence over the listener.

More generally, the forced choice measure leaves open questions about the precise inferences children made or the degree of knowledge children assigned to either the effective or ineffective speakers. Future work could examine both children's precise inferential process and epistemic attributions by asking children to make epistemic judgements throughout each video and also to recover the contents of their utterances. These additional studies could provide further insight into how children infer and update others' knowledge based on the events they observe.

While the current work focused on children's inferences about speakers' knowledge, children may also be able to reason about other mental states, such as *intent*, from these scenarios. For instance, in the Reykjavik example, you might infer that Barry whispered to Gus because he wanted to help or inform him. Yet, the real world is also full of cases of disinformation, where speakers intentionally provide misleading or false information to others to engender false beliefs. Whether children make nuanced evaluations of others' pedagogical intent by observing communicative exchanges remains an open question for future work.

More broadly, identifying knowledgeable informants based on their causal influence over others could support children's ability to assess and learn from others in pedagogical contexts (Bass et al., 2022; Einav & Robinson, 2011). Indeed, prior work shows that even infants track the reliability of informants (Tummeltshammer et al., 2014) and selectively query them by pointing to novel objects and points of interest (Bergus & Southgate, 2012; Lucca & Wilbourn, 2019). Being able to assess the knowledge of speakers via their influence on others could complement these abilities by allowing infants to identify knowledgeable informants without needing to understand the content of their speech.

Although we observed relatively late-emerging competence in the current studies, a number of factors could have obscured earlier competence. The current studies were quite cognitively demanding, requiring children to remember the contents of two similar videos, attend to the actions of multiple agents, and represent the epistemic states of at least

two agents at a time. Therefore, a number of cognitive processes, including mental state representation, executive function, memory, or a combination could be responsible for younger children's failure in these tasks.

The current findings also extend our understanding of young children's causal reasoning abilities more broadly. Previous research has shown that preschool-aged children can track changes to causal systems, predict downstream effects, and even identify experts based on their causal knowledge (e.g., Buchanan & Sobel, 2011; Kushnir et al., 2013; Lockhart et al., 2019). Yet, far fewer studies have investigated children's ability to reason about communication between two agents in terms of one's causal influence over the other. In the current studies, children inferred what a speaker knew based on the degree to which their communication appeared to cause a listener to perform either a successful action or an action they wouldn't have otherwise. Although, further research is needed to determine to what extent, if any, children in the current studies engaged in genuine counterfactual reasoning. Nonetheless, the ability to identify "causal influencers" in social settings dovetails with young children's capacity to identify causal difference makers and root causes in physical systems (Bonawitz et al., 2010; Goddu & Gopnik, 2020).

While prior work has focused primarily on how children reason about other minds in isolated contexts, recent research has begun to examine how they reason about other minds in more social environments with multiple agents. For example, socially evaluative contexts may facilitate infants' reasoning about others' beliefs (Woo et al., 2022; Woo & Spelke, 2022), and they can even use the kinds of actions agents perform with one another to infer the kind of relationship they possess (Thomas et al., 2022). The current work raises important questions about how such a nascent understanding might develop and support more sophisticated inferences in communicative contexts. Indeed, children in our studies applied their understanding of communication and information exchange to reason about the knowledge of communicative agents who exert influence on others rather than on the world itself.

Overall, the current results suggest that by 5-years, children can use the actions of one agent (listener) to infer what another agent (speaker) knows by reasoning causally about the consequences of communication. These abilities could play an important role in a world filled with ever-growing social networks and evolving manners of communication.

Acknowledgments

We thank the Stanford Social Learning Lab for helpful feedback on this project, Juelle Ford for assistance with data collection, and our funding sources: McDonnell Scholars Award and NSF BCS-2019567 (HG).

References

Aboody, R., Huey, H., & Jara-Ettinger, J. (2022). Preschoolers decide who is knowledgeable, who to inform, and who

- to trust via a causal understanding of how knowledge relates to action. *Cognition*, 228, 105212.
- Baron-Cohen, S., Leslie, A. M., & Frith, U. (1985). Does the autistic child have a “theory of mind”? *Cognition*, 21(1), 37–46.
- Bass, I., Bonawitz, E., Hawthorne-Madell, D., Vong, W. K., Goodman, N. D., & Gweon, H. (2022). The effects of information utility and teachers’ knowledge on evaluations of under-informative pedagogy across development. *Cognition*, 222, 104999.
- Begus, K., & Southgate, V. (2012). Infant pointing serves an interrogative function. *Developmental science*, 15(5), 611–617.
- Bonawitz, E. B., Ferranti, D., Saxe, R., Gopnik, A., Meltzoff, A. N., Woodward, J., & Schulz, L. E. (2010). Just do it? investigating the gap between prediction and action in toddlers’ causal inferences. *Cognition*, 115(1), 104–117.
- Buchanan, D. W., & Sobel, D. M. (2011). Mechanism-based causal reasoning in young children. *Child development*, 82(6), 2053–2066.
- Chuey, A., Williams, R. M., Frank, M. C., & Gweon, H. (2022). Young children’s reasoning about the epistemic consequences of auditory noise. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 44).
- 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.
- Gerstenberg, T. (2022). What would have happened? counterfactuals, hypotheticals and causal judgements. *Philosophical Transactions of the Royal Society B*, 377(1866), 20210339.
- Gerstenberg, T., Goodman, N. D., Lagnado, D. A., & Tenenbaum, J. B. (2021). A counterfactual simulation model of causal judgments for physical events. *Psychological review*, 128(5), 936.
- Goddu, M. K., & Gopnik, A. (2020). Learning what to change: Young children use “difference-making” to identify causally relevant variables. *Developmental psychology*, 56(2), 275.
- Hogrefe, G.-J., Wimmer, H., & Perner, J. (1986). Ignorance versus false belief: A developmental lag in attribution of epistemic states. *Child development*, 567–582.
- Jin, K.-S., Kim, Y., Song, M., Kim, Y.-J., Lee, H., Lee, Y., ... Song, H.-J. (2019). Fourteen-to eighteen-month-old infants use explicit linguistic information to update an agent’s false belief. *Frontiers in psychology*, 10, 2508.
- Kim, E. Y., & Song, H.-j. (2015). Six-month-olds actively predict others’ goal-directed actions. *Cognitive Development*, 33, 1–13.
- Knudsen, B., & Liszkowski, U. (2012). Eighteen-and 24-month-old infants correct others in anticipation of action mistakes. *Developmental science*, 15(1), 113–122.
- 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.
- Kushnir, T., Vredenburg, 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.
- Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, 25(3), 265–288.
- Lockhart, K. L., Chuey, A., Kerr, S., & Keil, F. C. (2019). The privileged status of knowing mechanistic information: An early epistemic bias. *Child development*, 90(5), 1772–1788.
- Lucca, K., & Wilbourn, M. P. (2019). The what and the how: Information-seeking pointing gestures facilitate learning labels and functions. *Journal of experimental child psychology*, 178, 417–436.
- Luo, Y., & Baillargeon, R. (2007). Do 12.5-month-old infants consider what objects others can see when interpreting their actions? *Cognition*, 105(3), 489–512.
- Luo, Y., & Johnson, S. C. (2009). Recognizing the role of perception in action at 6 months. *Developmental science*, 12(1), 142–149.
- Magid, R. W., Yan, P., Siegel, M. H., Tenenbaum, J. B., & Schulz, L. E. (2018). Changing minds: Children’s inferences about third party belief revision. *Developmental science*, 21(2), e12553.
- Michotte, A. (2017). *The perception of causality*. Routledge.
- Moll, H., Koring, C., Carpenter, M., & Tomasello, M. (2006). Infants determine others’ focus of attention by pragmatics and exclusion. *Journal of Cognition and Development*, 7(3), 411–430.
- Nyhout, A., & Ganea, P. A. (2019). Mature counterfactual reasoning in 4-and 5-year-olds. *Cognition*, 183, 57–66.
- Scott, R. M. (2017). Surprise! 20-month-old infants understand the emotional consequences of false beliefs. *Cognition*, 159, 33–47.
- Song, H.-j., Onishi, K. H., Baillargeon, R., & Fisher, C. (2008). Can an agent’s false belief be corrected by an appropriate communication? psychological reasoning in 18-month-old infants. *Cognition*, 109(3), 295–315.
- Surian, L., Caldi, S., & Sperber, D. (2007). Attribution of beliefs by 13-month-old infants. *Psychological science*, 18(7), 580–586.
- Tauzin, T., & Gergely, G. (2018). Communicative mind-reading in preverbal infants. *Scientific reports*, 8(1), 1–9.
- Thomas, A. J., Woo, B., Nettle, D., Spelke, E., & Saxe, R. (2022). Early concepts of intimacy: Young humans use saliva sharing to infer close relationships. *Science*, 375(6578), 311–315.

- Tomasello, M., & Haberl, K. (2003). Understanding attention: 12- and 18-month-olds know what is new for other persons. *Developmental psychology*, 39(5), 906.
- Tummeltshammer, K. S., Wu, R., Sobel, D. M., & Kirkham, N. Z. (2014). Infants track the reliability of potential informants. *Psychological Science*, 25(9), 1730–1738.
- Vouloumanos, A., Martin, A., & Onishi, K. H. (2014). Do 6-month-olds understand that speech can communicate? *Developmental Science*, 17(6), 872–879.
- Vouloumanos, A., Onishi, K. H., & Pogue, A. (2012). Twelve-month-old infants recognize that speech can communicate unobservable intentions. *Proceedings of the National Academy of Sciences*, 109(32), 12933–12937.
- Woo, B. M., & Spelke, E. (2022). Eight-month-old infants' social evaluations of agents who act on false beliefs. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 44).
- Woo, B. M., Tan, E., Yuen, F. L., & Hamlin, J. K. (2022). Socially evaluative contexts facilitate mentalizing. *Trends in Cognitive Sciences*.
- Wu, Y., Schulz, L., & Saxe, R. (2018). Toddlers connect emotional responses to epistemic states. In *Cogsci*.