UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

How you learned matters: The process by which others learn informs young children's decisions about whom to ask for help

Permalink

https://escholarship.org/uc/item/2mr04831

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 40(0)

Authors

Bridgers, Sophie

Gweon, Hyowon

Bretzke, Maria

et al.

Publication Date

2018

How you learned matters: The process by which others learn informs young children's decisions about whom to ask for help

Sophie Bridgers[1] **(sbridge@stanford.edu), Hyowon Gweon**[1], **Maria Bretzke**[2], **& Azzurra Ruggeri**[2]

[1]Department of Psychology, Stanford University, California, USA

[2]Max Planck Research Group iSearch, Max Planck Institute for Human Development, Berlin, Germany

Abstract

Prior work suggests that young children consider others' knowledge and expertise to decide from whom to learn. Do children also consider *how* others came to know what they know? Here we investigate young children's sensitivity to the process by which people have learned. In Exp.1, 3- to 6-year-olds preferentially sought help from an *active learner*, who had figured out how to solve a problem by herself, over learners who had learned through passive observation or direct instruction. Yet, this preference emerged only when the problem children needed to solve was related to the one the learners had previously solved (i.e., when they thought the active learner's competence would be relevant). These findings suggest children inferred competence from the process of active learning, but considered this competence to be constrained to a particular task rather than more broadly generalizeable. The results of Exp.2 (3- to 7-year-olds) suggest that younger children's learner preference might be driven by more superficial cues related to active learning such as being alone and that a more abstract understanding of the process of active learning might develop with age.

Keywords: help-seeking; selective trust; active learning; knowledge acquisition; problem-solving

Introduction

Children are selective social learners. Even very young children have mental models of what constitutes a knowledgeable informant and strategically choose from whom to learn (Shafto, Eaves, Navarro, & Perfors, 2012) . For instance, children as young as 3 years preferentially learn from other people who are knowledgeable rather than ignorant, and the sophistication of their ability to identify trustworthy informants develops across the preschool years (e.g., Birch, Vauthier, & Bloom, 2008; Corriveau, Fusaro, & Harris, 2009; Einav & Robinson, 2011; Koenig & Harris, 2005; Sabbagh & Baldwin, 2001, see Sobel & Kushnir, 2013 for a review).

Much of the research on children's selective social learning has focused on children's ability to identify *good teachers*. Collectively, this literature suggests that children are sensitive to the quality of information different teachers provide and use it to identify who is more knowledgeable and helpful. However, learning involves more than receiving useful information from others; it also involves acquiring the abilities to actively explore and gather information yourself. By identifying *good learners* and learning from them, children not only can learn *about* the world from others, but also *how* to effectively learn about the world. Indeed, a good learner offers a model of how to perform goal-directed actions that generate useful evidence or how to ask questions that elicit informative answers from others (Frazier, Gelman, Kaciroti, Russell, & Lumeng, 2011; Markant & Gureckis, 2014; Mills, Legare, Bills, & Mejias, 2010; Schulz & Bonawitz, 2007).

Similar to how knowledgeable teachers can be identified by the quality of information they provide, competent learners could be identified by the quality of their explorative actions or questions, or by their ability to make novel discoveries on their own. Prior work suggests that preschoolers are sensitive to the effectiveness of learners' inquiry strategies before being able to implement efficient strategies on their own. For instance, children as young as 4 years are already able to identify which agents ask the most informative questions, despite not being able to *generate* such questions on their own (Ruggeri, Sim, & Xu, 2017). In this sense, identifying competent learners might be useful not only to identify who can best help us learn new things, but also to serve as a stepping stone to developing successful inquiry strategies ourselves.

Do young children recognize good, competent learners? In this paper, we investigate the extent to which young children infer competence based on the nature of others' learning. In particular, do children prefer to learn and seek help from *active* problem-solvers, and does this preference depend on the kind of problem children themselves have to solve? In other words, is the learner's competence, inferred from her learning process (i.e., figuring out a problem from her own independent exploration v. with help from another person), generalizable to other tasks or constrained to the tasks on which this process was observed? We designed a study implementing an experimental paradigm similar to that used in prior work on selective trust. We presented children with multiple agents and asked them to choose whom they wanted to ask for help. But rather than contrasting *informants* who differed in knowledge, we presented children with *learners* who eventually acquired the same knowledge (i.e., how to activate a novel toy), but differed in *how* they acquired this knowledge. Our critical question was from whom children would seek help when given a problem that varied in its similarity to the one the learners had figured out.

Experiment 1

In Experiment 1, children observed three learners: (1) the Active learner who figured out how to activate a causal toy on her own, (2) the Instructed learner who learned from another person through direct instruction, and (3) the Passive learner who learned by watching an active learner figure out the toy on her own. Children were then given a causal toy to figure out themselves, and the opportunity to seek help from one of the three learners.

To examine whether children's choice was influenced by the similarity between the problem on which children needed
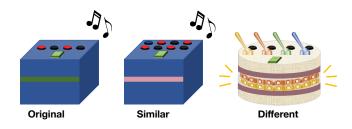
Figure 1: Schematic of toys: The Original and Similar Toys were the same shape and color, and both played music, but the Similar Toy was more complex. The Different Toy was a different shape, color, and texture; when activated, it lit up.

help and the problem the learners had previously solved, children were presented with three different toys: (1) the Original toy, identical to the one the learners' had figured out, (2) the Similar toy, which looked similar but was visually more complex and seemed harder to figure out, and (3) the Different toy, which was visually and functionally dissimilar from the Original and Similar toys, suggesting that it was part of a different set of toys and thus likely worked in a different way. Following prior work on children's selective trust and help-seeking, as well as that on children's question selection, we investigated children aged 3 to 6 years.

Methods

Participants We recruited 120 3-, 4-, 5-, and 6-year-olds ($n = 30$ per age group; $M(SD) = 60(14)$ months; 53% female) from local museums in Berlin, Germany. An additional 5 children were excluded from analysis because they refused to seek help ($n = 4$) or due to experimental error ($n = 1$).

Materials The Original toy was a blue cardboard box with a single row of 4 buttons (2 red, 2 black) alternating in color, and a green push-switch. When activated, this toy played music. The Similar toy was the same color and shape as the Original toy and had similar causal affordances; the only difference was that it had *two* rows of 4 buttons (red and black alternating in color). When activated, this toy also played music. The Different toy looked different from the Original toy and had additional causal affordances: It was made of a white round box with a wicker texture with 4 black buttons, a row of 4 different colored flip-switches, and a green push-switch. When activated, an LED strip that was wrapped around the toy lit up (see Figure 1). The toys were not actually functional, but surreptitiously activated by a remote control hidden from the child's view.

Three videos were used to introduce the different learners. The actors were three Caucasian women with brown hair, each wore a different colored t-shirt (blue, yellow, or red) and was referred to by that color (e.g., "My friend Blue"). Each actress always wore the same colored t-shirt, but the learner each actress portrayed (Active, Instructed, or Passive) was counterbalanced across children.

Procedure The experimenter told children about her friends who had learned earlier that day how to activate a toy. Children watched videos of the three different learners (Active, Instructed, Passive) in a pseudo-randomized order, on a tablet. All videos involved the Original toy and consisted of 4 phases: (1) Introduction, (2) Exploration, (3) First Activation, and (4) Second Activation. The Introduction and Second Activation were the same for all three learners, while the Exploration and First Activation differed.

In the *Introduction* phase, the learner sat at a table behind a black screen and lifted the Original toy from behind the screen. She rotated it forward such that children could see the top of the toy with the buttons and switch, while saying "Look at this cool toy. I wonder how it works." The learner then placed the toy back behind the screen. In subsequent phases, the toy remained behind the screen out of children's view; children could see whether someone was acting on the toy, but not what those actions were.

The *Exploration* and *First Activation* phases differed by learner. In the Active learner video, the Active learner apparently pushed buttons on the toy behind the screen. She explored the toy for 5 seconds, said "Hmm" to herself (2 seconds), and then explored for another 5 seconds (*Exploration*). Music then played, and the learner responded to this activation: "Aha! So, that is how this toy works!", indicating that she had discovered how to activate the toy (*First Activation*).

In the Instructed learner video, the *Exploration* phase was the same as the Active learner video. In the *First Activation* phase, however, a second actor wearing a grey t-shirt (henceforth the teacher) entered, made eye-contact with the learner, and apparently pushed a particular combination of buttons on the toy behind the screen. The music played, indicating that the teacher had activated the toy. The Instructed learner said, "Aha! So that is how this toy works!", indicating that she had learned from the teacher's action how to make the toy go.

In the Passive learner video, the *Exploration* phase started with the entrance of a second actress (black t-shirt, henceforth the model learner). The model learner apparently pushed buttons on the toy behind the screen. She explored the toy (5 seconds), said "Hmm" to herself (2 seconds), and then further explored the toy (5 seconds). The Passive learner watched the model learner's actions on the toy, but never interacted with the toy. In the *First Activation* phase, the music played from behind the screen, indicating that the model learner had successfully figured out how to activate the toy. The Passive learner said, "Aha! So that is how this toy works!", indicating that she had learned from the model learner's action how to make the toy go.

In the *Second Activation* phase (identical for all learners), the learner successfully activated the toy by herself. This phase made clear that all learners knew how to activate the toy, regardless of how they learned the solution.

After watching the three videos, the experimenter brought out the three toys, one at a time (order counterbalanced). For each toy, the experimenter explicitly stated its relation to the

toy the learners had interacted with in the videos. For the Original toy, she said, "This is the same toy as the one in the videos. Yellow, Red, and Blue have seen it before"; Similar toy: "This toy is similar to the one in the videos, but it looks more complicated. Yellow, Red, and Blue have never seen it before."; Different toy: "This toy is completely different than the one in the videos. Yellow, Red, and Blue have never seen it before." Children were given 10 seconds to explore the toy, but were unable to activate it. Then the experimenter presented photos of the three learners on the tablet and said, "Hmm, it's hard, isn't it? Maybe we should ask for help. Whom do you want to ask for help?" Children responded by tapping a photo on the tablet, which recorded their response. The same procedure was repeated for the two remaining toys. At the end, the experimenter showed children how to activate the Different toy and gave them the opportunity to activate it.

Predictions and Results

By the end of the videos, all three learners knew how to activate the Original toy and did so successfully. Thus, in the Original toy trial, we predicted children would have no preference for any of the learners because they each had the necessary knowledge to make that toy go. Only the Active learner video, however, provided clear evidence that the learner was capable of discovering the correct solution by herself. If children used the process of learning to guide their decisions about from whom to seek help, they should prefer the Active learner over the other learners, but only when they think her competence is likely to be helpful. Thus, we predicted that children would show a preference for the Active learner in the Similar toy trial, because this toy appears to be of the same type as the Original toy but more complex. If children think the Active learner's competence for figuring out toys on her own generalizes to all sorts of toys, then they might also prefer to seek her help in the Different toy trial. However, if they have inferred a more narrow competence specific to a certain kind of toy (e.g., blue toys, or music toys), children might again show no preference because they might not consider the active learner's skills to be relevant.

We fit a MCMC generalized linear mixed model (GLMM) predicting learner choice (Active, Instructed, Passive) with fixed effects of toy trial (categorical, 3-levels: Original, Similar, Different) and age (continuous) with a random intercept for subject.[1] This model revealed a main effect of toy trial: Children's choice of learner for whom they wanted to ask for help differed on the Similar toy trial compared to the Original ($\beta = -0.938$, 95% CIs [-1.556, -0.345], $p = .006$), but not on the Different toy trial compared to the Original ($\beta = 0.069$, 95% CIs [-0.540, 0.689], $p = .814$). Fitting this model again, dummy coding the Similar toy as the baseline variable, revealed that children's learner choice also differed on the Similar toy trial compared to the Different ($\beta = 0.966$, 95% CIs [0.330, 1.562], $p = .0008$).

We further analyzed children's choices by collapsing

[1]This model was run using the `MCMCglmm` package, version 2.25.

across age and conducting a series of chi-square goodness of fit tests comparing learner choice on each toy trial against chance (33%). As can be observed in Figure 2(a), children had no preference for which learner to ask for help when presented with the Original toy ($\chi^2(2) = 2.45$, $p = .294$) or the Different toy ($\chi^2(2) = 2.15$, $p = .341$). When presented with the Similar toy, however, children selected the Active learner more often than the other two ($\chi^2(2) = 18.6$, $p < .001$).

The refactored MCMC GLMM model also showed a marginal main effect of age ($\beta = -0.218$, 95% CIs [-0.451, 0.027], $p = .074$), providing suggestive evidence that children's preference for the Active learner on the Similar toy increased with age. To explore this potential age difference, we performed a median age split (median age = 60 months) and analyzed the younger ($n = 60$) and older ($n = 60$) children separately. For each age group, we fit a MCMC GLMM predicting learner choice with a fixed effect of toy trial and a random intercept for subject. These analyses revealed that older children's choice of learner differed on the Similar toy trial compared to the Original ($\beta = -1.147$, 95% CIs [-2.080, -0.309], $p = .011$) and Different ($\beta = 1.118$, 95% CIs [0.212, 2.092], $p = .022$). Younger children's learner choice, however, did not significantly differ across toys, though there was a trending difference between the Original and Similar toy trials ($\beta = -0.803$, 95% CIs [-1.688, -0.016], $p = .066$) and the Similar and Different toy trials ($\beta = 0.836$, 95% CIs [-0.089, 1.735], $p = .075$). (See Figure 2(b).)

Discussion

These results suggest that children consider the process of learning (i.e., active learning versus instruction or passive observation) as a relevant cue for deciding from whom to seek help. Note that all learners knew how and were able to activate the Original toy. Indeed, when faced with a problem identical to the one the learners had previously solved (the Original toy), children showed no preference for any learner. However, when faced with a novel problem, children's choice of learner showed a clear pattern: They preferred to seek help from the Active learner on a harder, related problem (the Similar toy), but not on a more distant problem (the Different toy). There is also suggestive evidence that children's selective preference for the Active learner on the Similar toy might increase with age.

Notably, children had no evidence that the Instructed or Passive learners would have *failed* to solve the original problem. In fact, both learners were interrupted: the Passive learner before she began exploring the toy, and the Instructed learner during her exploration. It is possible that if left to their own devices, they too would have figured out the toy by themselves. Yet, children still preferred to seek help from the Active learner; the only one for whom they had positive evidence that she could figure out the toy via her own exploration.

The selectivity observed in children's preference for the Active learner suggests they did not just form a positive association with her or a negative association with the other learn-
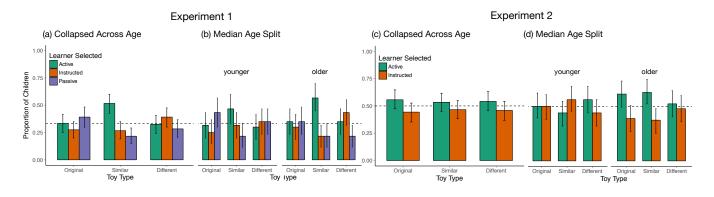
Figure 2: Experiment 1: Proportion of children who selected the Active, Instructed, or Passive learners as helpers for each toy (Original, Similar, and Different) collapsed across age (a) and by age group (b). Dashed line represents chance, 33%. Experiment 2: Proportion of children who selected the Active or Instructed learners as helpers for each toy collapsed across age (c) and by age group (d). Dashed line represents chance, 50%. Error bars indicate 95% bootstrapped confidence intervals.

ers. Rather, children seemed to infer from the Active learner's successful exploration of the Original toy some kind of competence or knowledge that is helpful for completing particular tasks but not others.

Experiment 2

Although results from Experiment 1 were consistent with our hypothesis, they leave open an alternative possibility: Children could have attributed competence to the Active learner simply because she was *alone* while learning, not because she actively discovered the solution by herself. Being alone makes it clear that the learner did not receive help; it also suggests that no one else *thought* the learner needed help, because no one stepped in as they did for the Instructed and Passive learners. Thus it is possible children thought the Active learner was more competent due to the absence of other people in the video, rather than by virtue of the process of learning itself.

We address this possibility in Experiment 2, introducing a small change to the Active Learner video, where now another person comes in and offers help to the learner. Because we did not observe any differences in children's preference for the Passive versus the Instructed learner in Experiment 1, in Experiment 2 we just compared the Active learner to the Instructed learner whose intention to explore and actions on the toy are more closely matched to those of the Active learner. Given the marginal effect of age we found in Experiment 1, we recruited a broader age range (3- to 7-year-olds) to explore whether there is a developmental difference in the cues (i.e., learning independently vs. being alone while learning) children use to infer competency.

Methods

Participants We recruited 133 3- to 7-year-olds ($M(SD) =$ 65(16) months; 47% female; planned sample is 150) from local museums in Berlin, Germany. An additional four children were excluded from analysis because they refused to seek help ($n = 3$) or due to distractions at the museum ($n = 1$).

Materials All materials were identical to those used in Experiment 1, except that the Passive learner video was not used and the Active learner video was replaced with a *Modified* Active learner video.

Procedure The procedure was identical to Experiment 1 except that children only saw two learner videos, the Modified Active learner and the Instructed learner videos (order counterbalanced). The Modified Active learner video was similar to the Active learner video used in Experiment 1, but a second actress (henceforth the teacher) entered at the end of the *Exploration* phase and looked at the learner, as if to offer help. The Active learner shook her head, suggesting that she refused the help, and acted on the toy, successfully activating it as the teacher watched. Thus, both the (modified) Active learner and the Instructed learner explored the toy for the same amount of time, and were both offered help by the teacher. The only difference was *who* eventually activated the toy successfully for the first time: the Active learner in the Modified Active learner video, and the teacher in the Instructed learner video.

Results and Discussion

We fit a mixed effects logistic regression model predicting learner choice (Active vs. Instructed) with fixed effects of toy trial (categorical, 3-levels: Original, Similar, Different) and age (continuous) with a random intercept for subject.[2] Unlike Experiment 1, this model revealed no main effect of toy trial (largest $\beta = 0.0915$, $SE = 0.247$, $z = 0.370$, $p = .711$) or of age ($\beta = -0.115$, $SE = 0.075$, $z = -1.543$, $p = .123$).

As in Experiment 1, we performed a median age split (median age = 66 months) and ran an exploratory analysis fitting a mixed effects logistic regression predicting learner choice with the single fixed effect of toy and a random intercept for subject, for each age-group separately (younger: $n = 66$; older: $n = 67$). These analyses revealed no difference in learner selected across toy trials for younger children (largest

---

[2]This model was run using the `lme4` package, version 1.1-15.

β = 0.244, $SE = 0.350$, $z = 0.697$, $p = .486$) or older children (largest β = 0.366, $SE = 0.350$, $z = 1.045$, $p = .296$).

Yet, as can be observed in Figure 2(d), the younger children appear to not have a preference for either learner across toys, while older children seem to prefer the Active learner on both the Original and Similar toys. A series of exploratory binomial tests looking at children's preference for the Active learner within each toy trial and age group were in line with this observation (two-tailed binomial tests, all $ps > .1$, except older children on the Original toy: $p = 0.086$ and Similar toy: $p = .0498$).

The analysis of our full sample from Experiment 2 suggests that children's preference for the Active learner in Experiment 1 may have been driven by the fact that this learner was alone while learning and not necessarily that she figured out the toy independently. However, our exploratory age analyses from Experiments 1 and 2 suggest that preference for the Active learner is emerging between age 3 and 7. Older children may indeed infer competence from the process of successful active learning itself, while younger children appear to have a weaker preference for the active learner and may be using more superficial cues, such as being alone. Data collection for Experiment 2 is ongoing, so we will see if the hint of a developmental change we observed here holds.

General Discussion

Children are selective in whom they decide to learn and get help from (e.g., Sobel & Kushnir, 2013). Building on past research exploring how children evaluate informants and teachers based on their knowledge, we find that young children are also sensitive to the process by which people come to know what they know. In Experiment 1, children differentiated a learner who solved a problem through independent, active exploration from a learner who was taught by someone else or watched someone else solve the same problem. In particular, they preferred to seek help from the Active learner on a related, seemingly more complex problem. However, when faced with a problem that was identical to or different from the problem for which all learners knew the solution, children showed no preference in whom they asked for help.

Why did children prefer the Active learner only when the task they had to solve was similar to the one that the learner was able to solve on her own? It is possible that children attributed a competence for problem-solving to the learner, and understood that it would likely apply to a near-transfer problem (the Similar toy). Interestingly, children did not generalize this expectation to a far(ther)-transfer problem, where the task was less related to the one that the Active learner was able to solve (the Different toy). In this sense, children's preference for the Active learner suggests that they did not just form a positive or "warm glow" association with her. Rather, children seemed to infer from the Active learner's independent, successful exploration of the toy a specific competence or knowledge constrained to a particular class of problems.

Defining the boundaries of a class of problems and decid-

ing when competence in one task is likely to transfer to another is challenging. In our experiments, the similarity between the toys was visually evident, but we would not necessarily expect children to preferentially seek help from the Active learner in domains where they struggle to assess the similarity between problems. Children were actually quite conservative in their willingness to extend the Active learner's competence. The ability to solve one mechanical toy might very well reflect a knack for solving mechanical toys in general, but children did not prefer the Active learner on the Different toy. In another sense, this conservatism seems justified given that children have only seen the Active learner figure out one toy, and they do not have any evidence that the Instructed and Passive learners are incapable of figuring out toys on their own.

Our findings also begin to shed light on what components of the active learning process children are using to assess the learner's competency. All learners knew how to activate the Original toy, so their preference for the Active learner likely comes from the process by which she arrived at this knowledge state and not from her declarative knowledge of how to make the toy play music. Furthermore, children's preference for the Active learner cannot purely be attributed to her experience directly interacting with the toy, since the Instructed learner also had such direct experience. Nor is it solely dependent on observing a pattern of actions that eventually led to successfully activating the toy, since the Passive learner observed the model learner's successful exploration. We can thus be more confident that this preference is due to the actual process of effective active learning (i.e., the ability to generate interventions that eventually led to successful activation), and not other factors shared with learning from others.

In Experiment 2, we began to decompose the cues that could have led children to consider the Active learner more competent than the Instructed and Passive learners – namely the fact that she was the only learner *alone* while learning. Children did not prefer an Active learner for whom a teacher stepped in and offered help to an Instructed learner who was taught how to make the toy go. Exploratory analyses, however, suggested that older children but not younger children may prefer the Active learner's help even when she is not alone while learning. These findings suggest that being alone while learning contributed to children's evaluation of the Active learner's competence in Experiment 1, though older children may also be able to infer competence from the process of active learning itself.

In both experiments, we found insignificant but suggestive evidence that older and younger children were performing differently. Older children seem to have a stronger preference for the active learner (Experiment 1), and be better able to infer competence from the actual process and effectiveness of a learner's exploration, rather than from simpler cues such as being alone (Experiment 2). It is important to note that learning alone is correlated with active learning, and if someone learns alone you can be more confident that they actually

figured out the problem on their own. In fact, a preference for a learner who is learning alone may be a stepping stone to drawing more sophisticated inferences about what successful active learning implies. We are cautious to draw conclusions about developmental differences at this time. Replicating and extending this work will be critical to identify the cues children rely on when inferring a learner's competence and how these cues might be changing in early childhood.

Moving forward, we also plan to further interrogate the inferences children are drawing about the Active learner. We interpret children's preference for the Active learner's help as reflecting an attribution of competence, but what the exact nature of this competence might be remains unclear. Children consider this competence to be bounded but do they think the Active learner just knows more about this particular class of toys or do they really think she knows more about how to solve this class of toys? Manipulating the process by which the Active learner explores the toy and seeing how this affects children's help-seeking would begin to answer these questions. For example, would children prefer an active learner who *deliberately* figured out the toy to an active learner who *stumbled* upon the correct activation sequence?

Similarly, our experiments only investigate whom children select to learn from, but it remains an open question what exactly they want to learn. If children could see the actual actions different learners took on the toy, would they choose to imitate certain learners over others? We know that toddlers exert different amounts of effort to achieve a goal depending on how much effort an adult exerts (Leonard, Lee, & Schulz, 2017), and preschoolers imitate intentional agents more faithfully than agents who act accidentally (Buchsbaum, Gopnik, Griffiths, & Shafto, 2011). It thus seems plausible that children would also selectively imitate the actions of certain learners depending on the competence inferred from their learning process.

Our work begins to provide insights into children's internal models of *good learners*. The current investigation focused on one quality of good learners: the ability to effectively query their environment. But competent learners also effectively query other people. Children can identify people who ask good questions (Ruggeri et al., 2017), would they also prefer to seek their help? It is also possible that children's understanding of what makes a good learner is mediated by the learning they see modeled and encouraged within their community and culture (Legare & Harris, 2016).

Here, we find preliminary evidence that children are sensitive to the benefits of active learning and selectively use it as a cue to identify helpers with abilities relevant for the task at hand. Even if two people have the same knowledge content, children can differentiate between them based on the learning process by which they have acquired this knowledge and consider some processes to reflect more competence than others. Even in childhood we understand that *how* you know, not just *what* you know, matters.

References

Birch, S. A. J., Vauthier, S. A., & Bloom, P. (2008). Three- and four-year-olds spontaneously use others' past performance to guide their learning. *Cognition*, *107*(3), 1018–1034.

Buchsbaum, D., Gopnik, A., Griffiths, T. L., & Shafto, P. (2011). Children's imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition*, *120*(3), 331–340.

Corriveau, K. H., Fusaro, M., & Harris, P. L. (2009). Going With the Flow: Preschoolers Prefer Nondissenters as Informants. *Psychological Science*, *20*(3), 372–377.

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.

Frazier, B. N., Gelman, S. A., Kaciroti, N., Russell, J. W., & Lumeng, J. C. (2011). I'll have what she's having: the impact of model characteristics on children's food choices. *Developmental Science*, *15*(1), 87–98.

Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, *76*(6), 1261–1277.

Legare, C. H., & Harris, P. L. (2016). The Ontogeny of Cultural Learning. *Child Development*, *87*(3), 633–642.

Leonard, J. A., Lee, Y., & Schulz, L. E. (2017). Infants make more attempts to achieve a goal when they see adults persist. *Science*, *357*(6357), 1290–1294.

Markant, D. B., & Gureckis, T. M. (2014). Is it better to select or to receive? Learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General*, *143*(1), 94-122.

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.

Ruggeri, A., Sim, Z. L., & Xu, F. (2017). Why is Toma late to school again? Preschoolers identify the most informative questions. *Developmental Psychology*, *53*(9), 1620–1632.

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.

Schulz, L., & Bonawitz, E. (2007). Serious fun: Preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology*, *43*(4), 1045–1050.

Shafto, P., Eaves, B., Navarro, D. J., & Perfors, A. (2012). Epistemic trust: modeling children's reasoning about others' knowledge and intent. *Developmental Science*, 436-447.

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), 779–797.