



Children seek help based on how others learn

Sophie Bridgers¹  | Costanza De Simone² | Hyowon Gweon¹ | Azzurra Ruggeri² 

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

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

Correspondence

Sophie Bridgers, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Room 46-2005, 77 Massachusetts Avenue, Cambridge, MA 02139-4307, USA.

Email: secb@mit.edu

Funding information

NSF

Abstract

Do children consider how others learned when seeking help? Across three experiments, German children ($N = 536$ 3-to-8 year olds, 49% female, majority White, tested 2017–2019) preferred to learn from successful active learners selectively by context: They sought help solving a problem from a learner who had independently discovered the solution to a previous problem over those who had learned through instruction or observation, but only when the current problem was novel, yet related, to the learners' problem (Experiment 1). Older, but not younger, children preferred the active learner even when she was offered help (Experiment 2), though only when her discovery was deliberate (Experiment 3). Although a preference to learn from successful active learners emerges early, a genuine appreciation for *process beyond outcome* increases across childhood.

Humans enter the world fairly helpless and with a great deal to learn. Fortunately, they also enter the world as powerful learners: Infants are able to draw rich inferences from the sparse data they observe in their environment (e.g., Denison et al., 2013; Denison & Xu, 2010; Gweon et al., 2010; Xu & Garcia, 2008), and young children actively and effectively test their beliefs and learn through their own exploration and play (e.g., Bonawitz et al., 2011; Gweon & Schulz, 2008, 2019; Legare, 2011; Piaget & Cook, 1952; Ruggeri et al., 2019; Schulz, 2012; Schulz & Bonawitz, 2007; Sim et al., 2017; Stahl & Feigenson, 2015; Xu & Kushnir, 2013). Yet, because of children's still relatively limited prior knowledge and cognitive/physical capabilities, they are often faced with problems they want to, but do not know how to, solve. Luckily, there are other individuals in the environment from whom children can seek help on such problems. Learning from others not only makes one's own learning less costly, but also enables learning beyond what one's direct, hands on experience could ever allow (Bandura & Walters, 1977; Boyd et al., 2011; Bridgers et al., 2020; Csibra & Gergely, 2009; Gweon, 2021; Kline, 2015; Tomasello, 2009; Tomasello et al., 1993; Vygotsky, 1997).

However, not all individuals are equally helpful. Indeed, children do not indiscriminately seek help from anyone, rather they leverage what they know about

others to make informed decisions about from whom to learn (Gweon, 2021; Koenig & Harris, 2005; Kushnir et al., 2008; Mills et al., 2010; Pasquini et al., 2007). For example, 3- and 4-year olds more faithfully and persistently imitate the actions of someone who claimed to be knowledgeable and intentionally demonstrated how to achieve a goal, compared to someone who communicated ignorance or accidentally achieved a goal (e.g., Bonawitz et al., 2011; Buchsbaum et al., 2011; Butler & Markman, 2012; Kushnir et al., 2008; Shafto et al., 2012). Also, from 3 to 5 years of age children become increasingly able to direct their questions to individuals who are knowledgeable, accurate, and have relevant expertise over individuals who are naïve, inaccurate, or have irrelevant knowledge (e.g., Koenig & Harris, 2005; Mills et al., 2010, 2011; Sobel & Corriveau, 2010).

Sometimes, however, children are faced with problems that no one around them has already solved. In such contexts, children need to identify individuals who can help them discover the solution to the problem, rather than individuals who already possess knowledge of it. Figuring out who can help on a completely novel problem is not straightforward, as direct evidence of the potential helpers' abilities to solve the problem is not available. In these cases, children must infer an individual's potential to figure out the novel problem based on

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2023 The Authors. *Child Development* published by Wiley Periodicals LLC on behalf of Society for Research in Child Development.

the individual's previous behavior. Observing how this individual solved other problems and discovered information in the past, that is their *learning history*, may be critical for informing this inference.

Prior research suggests that children use whether or not an individual was successful at solving previous problems to guide their help seeking on new problems (Cluver et al., 2013; Kushnir et al., 2013). In addition to considering the outcome of others' learning (e.g., action success), might children also be sensitive to the *process* by which an individual has learned or solved a problem in the past? For example, in their past history of learning, was this individual able to figure out the solution to a problem through their own actions, or did they receive help? If the discovery was made independently, was the underlying process deliberate (e.g., goal oriented) or accidental? Someone who has demonstrated both successful and deliberate active learning signals that they are capable of finding out new information about the world and solving novel problems. The relevance of this prior observation to the problem at hand, however, is also a crucial cue: Is the competence that was demonstrated in solving the prior problem likely to generalize to the present problem?

Across three experiments, we explore whether children's decisions about from whom to seek help take into account (1) the process by which an individual has learned the solution to a prior problem, as well as (2) the similarity between the problem children need to solve and the observed problem. We propose that how an individual has learned to solve a problem in the past licenses different inferences about their potential to solve future problems. In particular, someone who has demonstrated an ability to solve a prior problem through their own independent, intentional and goal-directed exploration provides stronger evidence for their ability to provide help on a new problem than someone who has learned the solution from someone else or someone who has accidentally discovered the solution.

Children's abilities to selectively seek help from effective, competent active learners may depend on their abilities to figure out when and from whom to seek help in problem-solving contexts, as well as on their understanding of the origins of their own and others' knowledge (i.e., how people come to know what they know). In what follows, we discuss these branches of prior literature in more detail and their relation to the question of whether children seek help on novel problems based on others' learning history. We then present the current experiments and hypotheses.

Children's help seeking when solving problems

Even though children are capable of independent, successful hypothesis testing and learning, humans do not (and cannot) learn in isolation (Bandura & Walters, 1977;

Vygotsky, 1997). Indeed, children integrate their own learning abilities with the social support available to them, making rational inferences about when to explore versus seek help (e.g., Goupil et al., 2016; Gweon & Schulz, 2011; Vredenburg & Kushnir, 2016). Infants as young as 8 months are more likely to reach for an out-of-reach object when another person is present than when no one is around (Liszkowski, 2013), and toddlers can decide when it is worth it to persist on their own versus ask for help, based on their own uncertainty about the task at hand (Goupil et al., 2016) or whether they are likely struggling due to an external versus internal cause (Gweon & Schulz, 2011). Four-year-olds' tendency to ask for assistance depends both on the difficulty of the task and their own problem-solving competence, lending further support to children's help-seeking behavior as information gathering (Vredenburg & Kushnir, 2016).

Children can evaluate their own problem-solving abilities in order to determine when they need help. Is children's help seeking also driven by their assessment of others' problem-solving abilities? As discussed, this evaluation might be particularly relevant when faced with problems no one in the environment has yet solved. When the potential helpers only have partial knowledge of a problem's solution (e.g., one knows the color of the key needed to unlock a box and the other knows about the key's shape), 4- and 5 year olds (but not 3 year olds) direct their questions to the appropriate helper depending on the information they want (Mills et al., 2011). There is also some evidence that children use the success of a person's past actions as a cue for their ability to successfully solve a novel problem. In one study, 2- and 3 year olds observed an individual who retrieved two different objects from two different puzzle boxes in a socially engaged, pedagogical manner, and another individual who communicated ignorance about the puzzle boxes, was socially disengaged, and failed to retrieve the objects (Cluver et al., 2013). Children were then faced with the same puzzle boxes the potential helpers had faced, as well as two new puzzle boxes for which children had no direct evidence of the helpers' abilities to solve. On the new puzzle boxes, both 2- and 3 year olds were more likely to seek help from the previously successful actor than from the previously unsuccessful actor. Three- and 4 year olds also consider past success in fixing toys as indicative of an individual's ability to fix a new toy, and selectively seek help from an individual who fixed two other toys (but did not know what the tools they used were called) over an individual whose actions failed to fix the toys but who knew the name of the tools (Kushnir et al., 2013).

From this research, we know that children make systematic help-seeking decisions based on whether or not (1) they need help (Vredenburg & Kushnir, 2016) and (2) the potential helpers were appropriately knowledgeable about and successful at solving past problems (Cluver et al., 2013; Kushnir et al., 2013). These studies

reveal that children pay attention to the outcomes of others' problem-solving attempts but what about the process by which they achieved these outcomes? What's more, it was not clear that the successful problem solvers were actually attempting to solve a problem in real-time, rather they appeared to be demonstrating knowledge they already possessed (e.g., in Cluver et al. (2013), the successful actor pedagogically demonstrates how to unlock the puzzle boxes). What if these helpers were all a-priori ignorant about a problem's solution and eventually successful at figuring it out, but the *avenues* by which they arrived at that success differed? Thus, the question still remains whether children are sensitive to others' history of *learning*, and more specifically, whether children would draw different inferences about an individual's potential to help on a new problem based on *how* that individual discovered the solution to a previous problem (e.g., independently versus from someone else).

Children's reasoning about the origins of their own and others' knowledge

To selectively seek help based on the process of how others have learned, one needs to be able to recognize that the same knowledge can be acquired via different means. For example, appreciating that if one person presses a lever on a toy and observes that it lights up and another person watches that first person press the lever and also sees the toy light up, then they both now know that pressing the lever makes the toy light up, but how they learned this causal relationship differed. One person learned directly through their own actions, while the other learned indirectly through another person's actions.

Distinguishing between direct and indirect discovery of information may initially be challenging for young children who struggle to keep in mind the source of their own knowledge. Three- and 4-year olds report that they have always known facts that in reality, they had just been told by someone else (Gopnik & Graf, 1988; Taylor et al., 1994). Children this age also struggle to articulate when someone else learned from their own interventions versus from another person's actions and verbal instruction: They underestimate the role of instruction and tend to say both individuals learned via their own actions (Sobel & Letourneau, 2018). By age five, children are better able to make this distinction about their own and others' knowledge (Gopnik & Graf, 1988; Sobel & Letourneau, 2018; Taylor et al., 1994), and from 4 to 8 years, children begin to describe learning as a process involving a source (e.g., a teacher) or strategy (e.g., practice) that results in the acquisition of knowledge (Sobel & Letourneau, 2015). Beyond age eight and into middle childhood, children become increasingly able to identify what knowledge is even possible to acquire from direct

experience versus knowledge that must be learned from others (e.g., "that birds fly" or "how to walk" vs. "how to fly a helicopter"; Lockhart et al., 2016).

Though 3- to 5-year olds are sometimes confused about the source of their own knowledge, research on selective trust has revealed that children this age draw different inferences about the reliability of others' testimony based on how the information being communicated was acquired. For example, in one study, 4-year olds were presented with opposing endorsements about which of the contents of two boxes were better: one endorsement came from a group of informants who each looked inside the boxes and the other came from a group where only one person looked and told the other members what was inside. Children were more likely to explore the contents of the box that was endorsed by the informants with direct, visual access than by the informants with hearsay (Hu et al., 2015). In another related study, 4- and 5-year olds (but not 3-year olds) were less likely to learn novel labels for animals from someone who in the past needed help to label familiar animals (i.e., another person whispered the labels to them) than from someone who correctly named the familiar animals without help (Einav & Robinson, 2011).

Taken together, these findings suggest that children are sensitive to how others learn and may very well privilege direct exploration and discovery of a problem's solution over more indirect or assisted means of discovery (e.g., observation or instruction). Indeed, children themselves are capable of learning from their own active exploration (e.g., Legare, 2011; Piaget & Cook, 1952; Schulz, 2012; Schulz et al., 2007; Sim & Xu, 2017; Xu & Kushnir, 2013), and there are even cases where children learn more from their direct actions in the world than from observation or instruction (Kuhn & Ho, 1980; Kushnir et al., 2009; Sobel & Sommerville, 2010; Sommerville et al., 2005), providing even further evidence that children might be sensitive to others' success as explorers.

However, appreciation of the difference between direct and indirect routes to discovery might become more robust across early to middle childhood. Recognizing and/or keeping in mind this distinction may be particularly difficult for 3- and 4-year olds, as they often struggle even to remember whether they, as well as others, learned something on their own versus from someone else (Gopnik & Graf, 1988; Sobel & Letourneau, 2018; Taylor et al., 1994). Moreover, in the studies on children's help seeking depending on the sources of others' knowledge or advice, perception and access were sufficient for knowledge acquisition, raising the question of how children reason about the discovery of information that cannot be gleaned simply via perceptual access, such as the latent structure of a causal system. Are children sensitive to different problem-solving processes and are they able to draw inferences about an agent's active-learning competence based on how the agent discovers more abstract information?

Current experiments

Do children use the *process* by which individuals have learned to solve problems in the past to inform their help seeking when presented with novel problems, which no one around them has yet solved? More specifically, do children draw different inferences about an individual's ability to provide help on a given problem, depending on (1) how the individual discovered the solution to a previous problem, and (2) the similarity between the previous and the current problem?

Across three studies, we present 3- to 8-year-old children with learners who have acquired the same knowledge but differ in how they acquired this knowledge, and ask children from whom they would like to seek help when faced with a problem that is identical, that is similar (but more complicated), or that is quite different from the one the learners solved. In addition to varying whether the learner solved the problem via active exploration, direct instruction, or passive observation (Experiment 1), we also varied whether a learner was physically alone or with someone else while solving the problem (Experiment 2), and whether or not the learner's actions and discovery were deliberate or accidental (Experiment 3). Thus, beyond investigating whether children prefer to learn from effective active learners, we explore the cues to which children attend when making inferences about the learner's competence.

When presented with a problem to which all learners already know the solution, it might not matter from whom you seek help: All learners can provide the solution to the problem you want to solve, regardless of how they originally discovered it. When presented with a problem to which none of the learners yet know the solution, it might be more beneficial to seek help from the active learner than the indirect learners, as only the active learner has demonstrated competence in independent exploration and successful discovery. The inference that the active learner has the greatest potential to provide effective problem-solving assistance, however, might vary in strength depending on the similarity between the problem at hand and the problem the learners previously solved. For example, this inference might be particularly strong when the novel problem is clearly related to the problem you observed the active learner solve, as the problem-solving ability demonstrated would have a high probability of being relevant. However, if the novel problem looks completely different from the problem previously solved, you may be less certain about whether the learner's demonstrated problem-solving competence would apply or about the learner's ability to transfer their skills to the new situation. Thus, we predicted that if children preferred to seek help from the active learner over the other learners, this preference would be strongest when the problem was analogous to the problem the learners previously solved, than when the problem was identical to or quite different from the observed problem.

We recruited children in the range of 3 to 8 years. Prior related work on children's help seeking when problem solving (Cluver et al., 2013; Kushnir et al., 2013; Mills et al., 2010, 2011; Vredenburg & Kushnir, 2016) and on their understanding of direct versus indirect sources of knowledge (Einav & Robinson, 2011; Hu et al., 2015; Lockhart et al., 2016; Sobel & Letourneau, 2018) has also focused on this age range, and identified it as a crucial time of both early competence and development in these domains. In our experiments, for children to selectively seek help from the active learner, they need to (1) recognize that the learners learned the solution to a past problem in different ways, (2) draw informed inferences about the various learners' potentials to solve a given problem, and (3) actually select the active learner as the one from whom they want help. On the one hand, given that children as young as 2 years can use a person's success on a past problem to infer their success on future problems (Cluver et al., 2013), and that children as young as 4 years choose to learn from individuals with direct perception and independent knowledge over individuals with hearsay or who needed help (Einav & Robinson, 2011; Hu et al., 2015), it is quite possible that by age four or five children would also be able to distinguish someone who solved a problem independently versus with assistance, and selectively seek their help on a novel, similar problem (Experiments 1 and 2). Also, given that children this age are more likely to trust the actions of a knowledgeable and intentional actor over an ignorant or accidental one (e.g., Bonawitz et al., 2011; Buchsbaum et al., 2011; Kushnir et al., 2008), they may also prefer to seek the help of someone who previously discovered the solution to a problem through intentional exploratory actions, rather than someone who stumbled upon the solution accidentally (Experiment 3). On the other hand, it is also quite possible that children's selective preference for the active learner would become more robust with age. Prior research indicates that from 3 to 5 years of age, children become more skilled at directing their questions to more knowledgeable individuals (Einav & Robinson, 2011; Mills et al., 2010, 2011), and their ability to keep track of the likely source of their own and others' beliefs continues to improve beyond age five (Gopnik & Graf, 1988; Lockhart et al., 2016; Sobel & Letourneau, 2018; Taylor et al., 1994).

The current work began as a relatively exploratory effort in this domain, inspired by the literature reviewed above, which suggested both the possibility of an early preference to seek help from active learners (over more passive learners) and of developmental change in this preference. However, all experiments include relatively large sample sizes, and the work became increasingly confirmatory from Experiment 1 to 3. We explicitly specify the confirmatory versus

exploratory nature of each analysis in the results sections.

EXPERIMENT 1

Experiment 1 examined (1) whether children prefer to seek help from successful active problem solvers, who solved a problem through their own independent exploration rather than another person's help, and (2) whether this preference is mediated by how similar the problem the active learner solved is to the problem children themselves have to solve. We recruited 3- to 6-year olds for this initial experiment because research on selective trust and help seeking point to it as a period where children might begin distinguishing independent exploration from assisted discovery (Cluver et al., 2013; Einav & Robinson, 2011; Hu et al., 2015; Kushnir et al., 2013; Sobel & Letourneau, 2018).

Children were presented with three learners: an *Active learner*, who figured out how to activate a causal toy through her own independent exploration, an *Instructed learner*, who first explored the toy, but did not activate it and instead watched another person demonstrate how to activate it, and a *Passive learner*, who observed another person explore and successfully activate the toy. All learners then activated the toy demonstrating that they each now knew how to make the toy go. Children were then presented with three toys in a counterbalanced order: the *Original toy* (the same toy as the learners' toy), the *Similar toy* (visually similar to but more complicated than the Original toy, i.e., with more buttons), and the *Different toy* (visually and functionally very distinct from the Original toy). On each toy, children were given the opportunity to explore the toy and try to make it go, but when they could not figure it out, they were prompted to seek help from one of the learners.

On the Original toy, we predicted that children would be equally likely to seek help from all three learners, since all of them knew the solution to this toy and had demonstrated the ability to activate it. On the Similar toy, a novel yet related problem for all learners, children should seek help from the learner who is most likely able to figure it out. We thus predicted they would seek help from the Active learner, the only one who provided clear evidence that she was independently capable of discovering how to activate a related toy (the Original toy). On the Different toy, children's preference may depend on how similar they think this novel problem is to the one posed by the Original toy. That is, if children think the Active learner's problem-solving competence generalizes to all sorts of toys, then they should again prefer her help. If children, however, infer a more narrow competence, specific to solving only a certain kind of toy, they might not consider the Active learner's skills relevant for solving the Different toy, and therefore show no preference in their help seeking. If there is a

developmental difference in children's help seeking, we expect the selective preference for the Active learner, that is, an appreciation for her competence, should increase with age (Einav & Robinson, 2011; Gopnik & Graf, 1988; Lockhart et al., 2016; Sobel & Letourneau, 2018; Taylor et al., 1994).

However, given that past research has shown that children in our age range consider successful action on past problems as a cue for future problem-solving success (Cluver et al., 2013; Kushnir et al., 2013), it is possible that even the 6-year-olds' help seeking would rely solely on successful activation of the Original toy. If this is the case, children would consider the three learners equally likely to be helpful on all three toys and not exhibit selectivity in their help seeking. Still another possibility is that children might indeed appreciate the Active learner's greater potential to provide effective help but not take into account the nature of the problem at hand. In this case, we would expect to observe a preference for the Active learner on all three toys.

Finally, children may ground their help-seeking preference on different aspects of the learning process. For example, if children consider *successful* exploration of the Original toy, regardless of whether it was direct or indirect, as a sign of competence, then they might seek help equally from the Active and Passive learners over the Instructed learner, who did not conduct nor observe a successful exploration. If children consider *direct, first-hand* exploration, regardless of its demonstrated success, as evidence of competence, they might instead be equally likely to seek help from the Active and Instructed learners over the Passive learner. However, if children consider both active exploration and its success, we should observe a preference for the Active learner over the other two.

Methods

Participants

In Experiment 1, 121 3- to 6-year olds (M_{age} (SD) = 60.66 (13.61) months; range: 36.69–83.74 months, 53% female) were recruited from and tested at the Natural History Museum in Berlin, Germany from August to October 2017. Participants in Experiment 1 were mostly of White European descent and were native or fluent speakers of German. We did not collect specific data on sample characteristics, so further demographic detail cannot be provided. Museums in Berlin, however, are generally affordable and accessible to people from a variety of economic backgrounds. The Natural History Museum offers appealing attractions for young kids and is of broad interest to people from differing educational and social backgrounds. Informed consent was obtained from the parents or legal guardians before commencing the study. The experimental design and procedure of all

experiments presented here were approved by the ethics committee of the Max Planck Institute for Human Development.

An additional 5 children were excluded from analyses because they refused to seek help ($n = 4$) or due to experimental error ($n = 1$). The experiment was a within-subject design: All children participated in three conditions or trials (the Original, Similar, and Different toy trials) in an order that was counterbalanced across participants. We selected this sample size so that we would have roughly 30 children per age bin by year, yielding a large sample per condition ($n = 121/\text{condition}$), as well as a fairly even distribution and medium to large sample across the age range to explore potential effects of age (Agresti & Min, 2002).

Materials

As illustrated in Figure 1, 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. The toys were not actually functional, but surreptitiously activated by a remote control hidden from children's view.

Children watched videos of three different learners (the Active learner, the Instructed learner, and the Passive learner) figuring out how to activate the Original toy. We took a number of steps to minimize the possibility that children's help-seeking behaviors might be influenced by confounding physical characteristics of the

actors. First, the three actors who portrayed the different learners were recruited from a larger pool of research assistants and applicants based on their similarity in appearance (e.g., similar in ethnicity, hair color, and attractiveness); they were all White women with brown hair. Second, in the videos, the actors wore identical t-shirts that only differed in color (blue, yellow, or red) and were referred to by that color (e.g., "My friend Blue"). Finally, each actor always wore the same colored t-shirt, but the learner each actor portrayed (Active, Instructed, or Passive) was counterbalanced across children. The videos are described in detail below.

Procedure

Testing took place in a quiet area separate from the main exhibits. The experimenter sat next to children at a table with a tablet between them. The experiment began with the experimenter telling children about her three friends who had learned earlier that day how to activate a toy. Participants were shown videos of the three different learners, each of whom tried to figure out how to activate a novel toy, in a pseudorandomized order on the 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 learners, while the Exploration and First Activation differed (see Figure 2a).

In the Introduction phase of all three videos, a learner sat at a table behind a black screen and held the toy above the screen so that it was visible to the children viewing the video. She rotated it forward such that children could see the buttons and switch on the top of the toy, and she expressed curiosity about how the toy worked, stating "Look at this cool toy. I wonder how it works." She then placed the toy in front of her but behind the black screen so that it was out of children's view.

Next came the Exploration and First Activation phases which differed by learner. In the *Active* learner video, the




FIGURE 1 Schematic of the toys used in all experiments: The Original and Similar toys had the same shape and color and both played music, though the Similar toy looked more complex (i.e., had more buttons). The Different toy had a different shape, color, and texture; when activated, instead of playing music, it lit up.

Experiment 1




FIGURE 2 Experiment 1: Study design and results. (a) Screenshots from videos used in Experiment 1. All learners (Active, Instructed, Passive) explored and learned how to activate the Original toy differently (*Exploration* and *First Activation*): the *Active* learner explored the toy independently and activated it for the first time on her own, the *Instructed* learner explored the toy independently but watched as another actor first activated it, and the *Passive* learner watched as another actor explored the toy independently and first activated it. All learners successfully activated the toy for a second time (*Second Activation*). (b) Proportion of children who selected each learner for help on each toy (Original, Similar, Different) collapsed across age. Dashed horizontal line represents chance, 33%. Error bars are bootstrapped 95% CIs.

learner explored the toy on her own. She appeared to press the buttons and flip the switch for 5s, said “Hmm” to herself (2s), and then explored for another 5s. Music then played, seeming to come from the toy and indicating that the learner had just figured out how to activate it (First Activation). The learner expressed surprise and responded to this activation: “Aha! So, that is how this toy works!”, confirming that she had discovered how to activate the toy.

In the *Instructed* learner video, the learner explored the toy on her own for 10s just like the Active learner. In the First Activation phase, however, a second actor entered the scene and demonstrated how to activate the toy from behind the screen. The music played, indicating that this second actor had activated the toy. The Instructed learner expressed surprise and said, “Aha! So that is how this toy works!”, confirming that she had learned from this other person's intervention how to activate the toy.

In the *Passive* learner video, the learner did not explore the toy nor did she first activate it; instead she observed someone else explore and activate the toy. The Exploration phase started with the entrance of a second actor who sat down next to the Passive learner and explored the toy for 10s in the same way as the Active and Instructed learners. The Passive learner watched this other actor's actions on the toy, but never interacted with the toy. In the First Activation phase, the music played, and the Passive learner expressed surprise and said, “Aha! So that is how this toy works!”, indicating that she had learned from this other person's actions how to activate the toy.

All three videos ended with the Second Activation phase, which was identical across learners: the learner

intervened on the toy and music played. This phase made clear that all learners knew how to activate the toy, regardless of how they learned the solution. Because exploration and both activations took place from behind the black screen, children never observed the actual sequence of actions the learner (or second actor) performed on the toy during exploration, nor did they observe whether the switch needed to be flipped nor which buttons needed to be pressed to activate the toy. Thus, from the videos, children observed how each learner acquired the solution to the toy and that each learner could successfully activate the toy on her own but did not learn how to make the toy play music themselves.

After showing children the videos, the experimenter brought out three toys, one at a time (order counterbalanced) and explicitly stated its relation to the toy with which the learners had interacted in the videos: (1) the *Original* toy, identical to the one the learners had figured out in the videos (“This is the same toy as the one in the videos. [The learners] have seen it before.”), (2) the *Similar* toy, which looked similar but had more buttons, suggesting that it was related to the toy in the video but more complicated (“This toy is similar to the one in the videos, but it looks more complicated. [The learners] have never seen it before.”), and (3) the *Different* toy, which was visually and functionally dissimilar from the Original and Similar toys, suggesting that it was a different kind of toy and thus likely worked in a different way (“This toy is completely different than the one in the videos. [The learners] have never seen it before.”). Children were given 10s to figure out how to activate each toy, but could not succeed, as the toys were in reality activated by hidden remotes.

At this point, the experimenter presented photos of the 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?” The experimenter prompted children to indicate which learner they would like to ask for help by pointing to a picture of the learner on the tablet. Children selected from all three learners for each toy (i.e., they could select the same learner more than once) and did not receive feedback on their choice, rather the experimenter said, “Ok, we will ask her for help later.”

After choosing whom to ask for help on each toy, as an exploratory measure, children were asked to describe how each learner figured out the Original toy. Children were shown portrait photos of the three learners, one by one, and were told: “Look at my friend Blue/Yellow/Red! Do you remember how she learned to activate this toy?” (pointing at the Original toy). As it may have been particularly challenging for younger children to find the right words to describe the learning process, when children seemed reluctant to respond, the experimenter prompted an answer saying: “Did she figure it out independently, by watching someone else, or did someone else tell her how to [do it]?” (see [Supporting Information](#) for details on reliability coding and analyses for all experiments).

At the very end, the experimenter explained that the learners were unfortunately unavailable. As consolation, the experimenter showed children how to activate the Different toy and gave them the opportunity to activate it.

Notably, a number of steps were taken to ensure that the only difference across the videos was the Exploration and First Activation phases: All three learners claimed ignorance at the start of the video in the same manner and with the same statement, expressed surprise and made the same verbal statement (“Aha! So that is how this toy works!”) upon the first activation, demonstrated that they could each successfully activate the toy at the end, and were equally socially engaged with the video camera in terms of eye contact and tone throughout.

Results

We fit a mixed effects Bayesian multinomial (categorical) regression predicting children's learner choice (categorical, 3-levels: Active, Instructed, Passive with Active as the reference category) with a fixed effect of toy trial (categorical, 3-levels: Original, Similar, Different with Original dummy coded as the reference category) and a random intercept by subject. This analysis is confirmatory for the Original and Similar toys based on our clear, a-priori predictions. However, this same analysis is exploratory for the Different toy, where we were agnostic to whether children would exhibit a preference for the Active learner or no preference for any of the learners. In addition to this model, we also fit a model with simple effects of toy and age (continuous), as well as a

model with simple effects of toy and age, and their interaction. Formal model comparison, as assessed by an approximate Bayesian model comparison method, preferred the model with the single predictor of toy, so we report the results of this analysis (i.e., learner choice \sim toy + (1|subject)). The R-package *brms::brm* was used to run all analyses reported in this paper. All formal model comparisons presented were assessed using expected log-posterior densities, which index the quality of the predictions of each model, and were estimated using leave-one-out cross validation using the *loo* package in R; we report the outcomes of model comparison in terms of (posterior) model weights (McElreath, 2020; Nicenboim et al., 2021). (See [Supporting Information](#) for more details on these model comparisons.)

We first look at children's learner choice within each toy trial. As a decision rule here and throughout the paper, we see if the 95% CI for the parameter of the relevant contrast crosses 0. Please note that for Bayesian regression models, the parameters can be interpreted in the same way as the parameters of frequentist regression models. The only difference is that for the Bayesian regression model the 95% CIs can be interpreted as there being a 95% chance that the parameter falls within that interval (as opposed to the frequentist interpretation of confidence intervals, which has a different, less intuitive interpretation, Morey et al., 2016). The model revealed that children did not exhibit a preference to seek help from any of the three learners when presented with the Original toy; specifically, their tendency to select the Active learner over the Instructed learner or Passive learner was not significant (Active vs. Instructed: $\beta = 0.23$, 95% CI $[-0.24, 0.70]$; Active vs. Passive: $\beta = -0.13$, 95% CI $[-0.56, 0.29]$). Children also had no preference when presented with the Different toy (Active vs. Instructed: $\beta = -0.16$, 95% CI $[-0.59, 0.26]$; Active vs. Passive: $\beta = 0.17$, 95% CI $[-0.28, 0.63]$). When presented with the Similar toy, however, children selected the Active learner more often than the other two (Active vs. Instructed: $\beta = 0.63$, 95% CI $[0.21, 1.06]$; Active vs. Passive: $\beta = 0.83$, 95% CI $[0.39, 1.30]$).

Second, we examine whether children's tendency to seek help from the Active learner over the other learners differed across the toy trials. Children were more likely to select the Active learner (over the Passive learner) on the Similar toy compared to both the Original toy ($\beta = 0.97$, 95% CI $[0.34, 1.60]$) and the Different toy ($\beta = 0.66$, 95% CI $[0.03, 1.30]$). Children's tendency to select the Active learner over the Instructed learner did not significantly differ on the Similar compared to the Original toy ($\beta = 0.40$, 95% CI $[-0.23, 1.03]$), but did differ on the Similar compared to the Different ($\beta = 0.79$, 95% CI $[0.20, 1.39]$). Across the Original and Different toys, children's tendency to select the Active learner over the Instructed learner did not differ ($\beta = -0.39$, 95% CI $[-1.03, 0.24]$) nor did their tendency to select the Active learner over the Passive learner differ ($\beta = 0.31$, 95% CI $[-0.31, 0.93]$). If we

just look at children's preference for the Active learner, they preferred the Active learner significantly more often on the Similar toy compared to both the Original and the Different toys (Similar vs. Original: $\beta = 0.73$, 95% CI [0.48, 1.04]; Similar vs. Different: $\beta = 0.69$, 95% CI [0.46, 0.97]; see Figure 2b–d).

With respect to the memory questions, 5 children failed to remember how each learner figured out the toy (memory score 0: 4.13%), 33 correctly identified one out of the three learners (memory score 1: 27.27%), 47 correctly identified two out of the three learners (memory score 2: 38.84%) and 36 correctly identified all three learners (memory score 3: 29.75%). To explore potential relations between these questions and from whom children asked for help, we fit two exploratory Bayesian multinomial (categorical) regressions predicting children's learner choice (categorical, 3-levels: Active, Instructed, Passive with Active as the reference category) with (1) simple effects of toy trial (categorical, 3-levels: Original, Similar, Different with Original as the reference category) and memory score (coded as an integer from 0 to 3) and (2) simple effects of toy trial and memory score, and their interaction. We conducted a formal model comparison of the model above with the single predictor of toy and these models with memory score as a predictor; this comparison preferred the model with the single predictor of toy, indicating that children's memory score did not predict their learner choice. An exploratory Bayesian linear regression with age (centered and continuous) predicting memory score indicated that memory score did significantly improve with age ($\beta = 0.02$, 95% CI [0.02, 0.03]).

Discussion

As hypothesized, we found that children's help-seeking patterns depended on both (1) how others had learned in the past and (2) the similarity between the problem at hand and the one the learners had previously solved. In particular, when faced with a problem that all learners knew how to solve (the Original toy), children did not exhibit a preference for one learner over the others, suggesting that they considered all three learners as equally helpful. When faced with a novel problem that none of the learners knew how to solve but that was related to the problem the learners had previously solved (Similar toy), children preferentially sought the Active learner's help over the Passive and Instructed learners, suggesting that they inferred she was the most likely to provide effective assistance in figuring it out. When faced with a novel problem that was not clearly related to the one the learners had previously solved (the Different toy), however, children again were equally likely to seek help from all three learners. The fact that children's help seeking differentiated the learners on the Similar toy demonstrates that their relatively poor performance on the memory

questions was not due to an inability to distinguish the learners based on their learning process. Rather, it may have been due to the timing of the questions, which were asked at the end of the help-seeking task, and/or to a difficulty articulating and matching the descriptions provided to how each learner had learned. We return to this point in the [General Discussion](#).

Why did children preferentially seek the Active learner's help only on a novel problem that was analogous to the original problem, but not on a novel, apparently less related problem? On the one hand, children's inferences seem quite conservative: the ability to solve one mechanical toy might very well reflect a knack for solving mechanical toys in general (e.g., Cluver et al., 2013; Kushnir et al., 2013). On the other hand, given the specifics of our experiment and the particular learners contrasted, a conservative preference may have been warranted. First, children only had one data point they could use to compare the learners and assess their competence (i.e., the learners' performance on the Original toy). Second, the contrast in competence established was subtle. Although the Active learner was the only agent who clearly demonstrated that she was capable of discovering the solution to a problem on her own, the Passive and the Instructed learners had not clearly failed, but rather had not been given the opportunity to try (Passive) or had been interrupted in the attempt (Instructed); it is possible that if left to their own devices, they too would have figured out the Original toy by themselves. From this perspective, children's decision was not between a clearly helpful, competent agent and clearly unhelpful, incompetent agents. Rather, all three learners could have been somewhat helpful and children's task was to select which agent was likely to be the *most* helpful on each toy. Children seemed to infer that the Active learner was their best bet on some, but not on all, tasks. If children had observed more examples of the Active learner successfully figuring out different problems or if it was clear that the Passive and Instructed learners were unable to solve the problem on their own, we expect children might have exhibited an even stronger preference for the Active learner and been more willing to generalize this preference.

These findings also begin to elucidate the components of the active-learning process children might have considered to assess the learner's competency. All learners knew how to activate the Original toy and successfully activated it. Thus, children's preference for the Active learner likely stemmed from the process by which she had arrived at this knowledge state and action success, and not from her knowledge of how to generate this particular toy's music effect, nor from her physical ability to generate it. Furthermore, children's preference for the Active learner cannot purely be attributed to the Active learner's experience directly interacting with the Original toy, as the Instructed learner also had had such direct experience. Nor could it have been solely dependent on the Active learner observing a pattern of actions

that eventually led to successfully activating the toy, because the Passive learner also had observed a successful exploration, though performed by someone else. Rather, it seems important that the Active learner both directly explored the toy and independently discovered how to activate it. This preference is in line with prior work showing that 4-year olds endorse information from individuals with direct perceptual access to the state of the world versus individuals with hearsay (Hu et al., 2015), as well as individuals who did not need help providing information in the past over individuals who needed help (Einav & Robinson, 2011).

The results of Experiment 1 provide initial support for our key hypotheses: Children appeared to consider the *process* by which individuals have learned the solution to a problem (i.e., independent, active learning versus instruction or passive observation) as a relevant cue for deciding from whom to seek help on a novel problem. However, in Experiment 1 the Active learner's independence presents an unintended confound: She was the only learner who was physically alone in the video. Being alone while learning does indeed correlate with autonomous exploration and discovery and likely helped make it clear that the learner figured out the problem on her own. However, being alone is not necessarily causally related to one's independent problem-solving competence. Relatedly, being alone 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. Indeed, recent research suggests that 4- to 6-year olds consider groups of children who receive help from a teacher while solving a problem less smart than groups who do not receive help (Sierksma & Shutts, 2020). Thus, it is possible that children thought the Active learner was more competent just because there were no other people in the video, and specifically due to the absence of someone who offered assistance, rather than by virtue of the process of learning itself.

We addressed this potential confound in Experiment 2 by comparing an Active learner who is offered assistance but discovers the Original toy's solution on her own, with an Instructed learner. Because we did not observe any differences in children's preference for the Passive versus the Instructed learners in Experiment 1, we decided to just compare the new (no longer alone) Active learner to the Instructed learner, whose intention to explore and actions on the toy, as compared to the Passive learner, were more closely matched to those of the Active learner.

Finally, we did not find an effect of age in Experiment 1. This was somewhat surprising given that in prior research 3-year olds did not preferentially seek help from someone who was independently knowledgeable about animal labels over someone who received help (Einav & Robinson, 2011), and 4-year olds seemed to struggle to articulate the difference between exploration and instruction (Sobel & Letourneau, 2018), as well as to remember when they learned information from someone

else (Gopnik & Graf, 1988; Taylor et al., 1994). It is possible that being alone while learning helped younger children to differentiate the Active learner from the other two and that they might find it more challenging to distinguish the learners in Experiment 2. We therefore recruited a slightly 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 and not being offered help) children use to infer problem-solving competence.

EXPERIMENT 2

In Experiment 2, we presented children with a slightly modified version of the Active learner video (another actor entered the room toward the end of the Active learner's exploration and watched as the Active learner activated the toy), and the same Instructed learner video used in Experiment 1. If children no longer exhibit a preference for the Active learner, it would imply that, when making their inferences in Experiment 1, children relied on the fact that she was physically alone and/or no one offered her help while learning rather than her independent exploration and discovery.

Methods

Participants

In Experiment 2, 188 3- to 7-year olds (M (SD) = 67.63 (17.46) months; range: 36.66–96.03 months; 48% female) were recruited from and tested at the Natural History Museum in Berlin, Germany from November 2017–November 2018. An additional 4 children were excluded from analyses because they refused to seek help ($n = 3$) or due to distractions at the museum ($n = 1$). We selected this sample size so that we would have a large sample size per condition ($n = 188$ /condition), as well as a medium to large sample across the age range (we aimed for approximately 30–40 children per age bin by year) to explore potential relations between help seeking and age. A post-hoc power analysis simulating a mixed effects logistic regression model with *simr::powerSim* based on the experimental results below revealed that with 188 participants we achieved 63% power for the effects of interest.

Participants were mostly of White European descent and were native or fluent speakers of German. Informed consent was obtained from the parents or legal guardians before commencing the study.

Materials

The toy materials were identical to those used in Experiment 1. The same Instructed learner video from

Experiment 1 was used along with a modified Active learner video in which the learner was no longer alone when she discovered how to activate the toy (described below). The learners in the videos were again White women with brown hair, each wore a different colored t-shirt (blue or yellow) and were referred to by that color (e.g., “My friend Blue”). Each actor always wore the same colored t-shirt, but the learner each actor portrayed (Active (not alone) or Instructed) was counterbalanced across children.

Procedure

There were only two differences from the procedure used in Experiment 1: (1) children selected between two learners: the Active and the Instructed learners, and (2) the Active learner video was modified: After exploring on her own for 10s, a second actor entered the scene, but the Active learner shook her head signaling that she did not need help; the second actor then sat silently next to the learner and watched as the learner continued exploring and then successfully activated the toy. More details on the videos can be found in Figure 3a. Otherwise, the procedure was identical to that of Experiment 1: Children participated in the Original, Similar, and Different toy trials (order counterbalanced), in which they explored

a toy and then were prompted to ask a learner to help them figure it out. As in Experiment 1, after choosing which learner they would like to ask for help on each of the three toys, children were asked to describe how each learner had learned about the Original toy. When children hesitated to respond, the experimenter prompted an answer saying: “Did she figure it out independently or did someone else tell her how to [do it]?”

Results

We fit a confirmatory mixed effects Bayesian logistic regression predicting children's learner choice (Active vs. Instructed) with fixed effects of toy trial (categorical, 3-levels: Original, Similar, Different with Original dummy coded as the reference category), age in months (continuous and centered), and their interaction, including a random intercept by subject. In addition to this model, we also fit a model with simple effects of toy and age, as well as a model with toy as a single predictor. Formal model comparison preferred the model with simple effects of toy and age, and their interaction, so we report the results of this analysis (i.e., learner choice \sim toy \times age + (1|subject)).

Collapsing across age and looking at choice within each toy trial, this model revealed that children selected




FIGURE 3 Experiment 2: Study design and results. (a) Screenshots from videos used in Experiment 2. As in Experiment 1, the process of learning differed across learners but at the end they could both activate the toy by themselves. The *Instructed* learner was the same as in Experiment 1; the *Active* learner explored and discovered the solution through her own actions but in the presence of another individual. (b) Proportion of children who selected each learner by toy collapsed across age. (c) Children's learner choice by age in months (continuous) faceted by toy (each dot represents an individual participant; dashed gray vertical line represents median age, 67.79 months). (d) Discretized version of data in (c): Proportion of children who selected the Active learner by age in years faceted by toy. Dashed horizontal line represents chance, 50%. Error bars are bootstrapped 95% CIs.

the Active learner significantly more often than the Instructed on the Similar toy ($\beta = 0.34$, 95% CI [0.04, 0.65]), but their tendency to select the Active learner over the Instructed did not differ on the Original toy ($\beta = 0.20$, 95% CI [-0.09, 0.49]) or Different toy ($\beta = 0.22$, 95% CI [-0.08, 0.52]; see Figure 3b). Comparing learner choice across toy trials, this model also indicated that children's tendency to select the Active learner on the Similar versus Original toy increased with age ($\beta = 0.028$, 95% CI [0.003, 0.052]). There were no simple effects of toy (Similar vs. Original: $\beta = 0.14$, 95% CI [-0.27, 0.57]; Similar vs. Different: $\beta = 0.12$, 95% CI [-0.31, 0.55]; Different vs. Original: $\beta = 0.02$, 95% CI [-0.40, 0.44]), of age ($\beta = -0.01$, 95% CI [-0.02, 0.01]), nor other interactions with age (different \times age: $\beta = 0.01$, 95% CI [-0.02, 0.03]; see Figure 3c,d).

For the memory questions, 16 children failed to remember how either learner figured out the toy (memory score 0: 8.51%), 27 correctly identified one learner (memory score 1: 14.36%), and 145 correctly identified both learners (memory score 2: 77.13%). We again conducted exploratory analyses to see if memory score predicted learner choice. Formal model comparison of (1) a model with a simple effect of toy trial (Original is reference category), (2) a model with simple effects of toy trial and memory score (integer from 0 to 2), and (3) a model with simple effects of toy trial and memory score, and their interaction preferred (3) the interactive model. This model revealed that children's tendency to select the Active learner on the Similar compared to the Original toy increased with memory score ($\beta = 1.01$, 95% CI [0.33, 1.72]). An exploratory Bayesian linear regression with age (centered and continuous) predicting memory score indicated that memory score significantly improved with age ($\beta = 0.012$, 95% CI [0.009, 0.015]).

Discussion

When presented with an Active learner (for whom someone entered the room and watched as the learner activated the toy) and an Instructed learner (for whom someone entered the room and activated the toy for the learner), children did not preferentially seek help from either learner on the Original or Different toys but did prefer the Active learner on the Similar toy. Children's preference for the Active learner on the Similar toy (compared to the Original toy) also increased from ages 3 to 7 years. The current study was not designed to pinpoint precisely when this shift occurs between three and seven, but by examining the distribution of children's responses, the strengthening in preference appears fairly continuous across this age range and may continue to increase beyond age seven. At 3 years of age, children did not preferentially seek help from either learner on the Similar toy (45.8% selected the Active learner, 95% CI [30.0%, 62.1%]) but by age seven, they reliably selected

the Active learner (75.5% selected the Active learner, 95% CI [63.2%, 87.5%]; see Figure 3d).

The majority of children (77%) correctly identified how each learner figured out the Original toy (i.e., either independently or told by someone else), but exploratory analyses revealed that memory score improved with age and that children's tendency to select the Active learner on the Similar toy also increased with memory score. These analyses suggest that from 3 to 7 years of age both children's ability to recognize that these learners learned in different ways and to then draw the inference that the Active learner was the one who was more likely to provide effective problem-solving help may be improving (see Gopnik & Graf, 1988; Sobel & Letourneau, 2018; Taylor et al., 1994). Previous literature on selective trust suggests that from 3 to 5 years of age, children also more consistently differentiate informants based on their relative reliability in both explicit evaluations and their help-seeking behaviors (e.g., Einav & Robinson, 2011; Kushnir et al., 2013; Mills et al., 2010; Pasquini et al., 2007).

Together, the results of Experiments 1 and 2 suggest that even though children between 3 and 6 years of age showed a selective preference for the Active learner in Experiment 1, the ways in which children identify such learners may undergo a developmental change from age 3 to 7. That is, younger children in Experiment 1 may have based their preference on cues that correlate with independent exploration and discovery (i.e., being alone while exploring and no one thinking you need assistance), while older children may have considered the actual process of independent successful exploration and discovery, regardless of whether or not other people were physically present or offered help. Being alone while learning and not having others present who think you need help reflect more the external situation and others' assessment of your ability, than your actual, intrinsic problem-solving competence. It is possible that the confluence of such external and internal cues help younger children to identify a learner's ability to solve problems on their own, while older children's inferences are more robust to these correlated but causally unrelated cues.

In particular, younger children in Experiment 2 might have interpreted someone's stepping in as an act of offering help to the Active learner, leading children in this experiment to draw a weaker inference about the Active learner's problem-solving competence than in Experiment 1. As discussed, recent work suggests that children consider the act of offering help as an indication that the target of this offer is less competent or intelligent than someone who is not offered help (i.e., someone thought this person needed assistance, so they must not be good at figuring out this task on their own, Sierksma & Shutts, 2020). Another, not mutually exclusive, possibility is that younger children's preference for the Active learner was mitigated because she shook her head at the second actor who entered the scene. This gesture could have been seen as the Active learner expressing negativity

or refusing to socially engage. In prior work, 4- and 5 year olds were less likely to seek help from a previously successful actor if they had been socially disengaged (Rowles & Mills, 2018). Future work could test whether having the Active learner not shake her head would increase younger children's tendency to seek her help.

Another difference across the videos in Experiments 1 and 2 is that the Active learner is the only learner to first activate the toy on her own; in the other videos, another actor initially activates the toy and then the learner activates it. Thus, it is possible that children are only reasoning about the Active learner's independent *first* activation of the toy, rather than the process of exploration that led her to that discovery. In Experiment 3, we addressed this possibility by contrasting two Active learners: one who intentionally explores the toy and discovers its solution versus one who does not explore and accidentally discovers its solution.

EXPERIMENT 3

In Experiment 3, children were presented with the same Active learner video from Experiment 1 (here called the *Deliberate* learner) and with an *Accidental* learner, who does not explore the toy but still activates it (though unintentionally). It is possible children might prefer to seek help from anyone who makes a novel discovery, regardless of whether the actions leading to the discovery were intentional or accidental. Three- to 5 year olds, however, distinguish between intentional and unintentional demonstrations to guide their imitation and exploration (Bonawitz et al., 2011; Buchsbaum et al., 2011; Butler & Markman, 2012), so we hypothesized that children would consider intentional discoveries as stronger evidence for effective problem solving than accidental ones. Because in Experiment 2 younger children did not seem to differentiate between the learners, and 3 year olds are even less reliable than 4 year olds in their help seeking (e.g., Einav & Robinson, 2011; Kushnir & Gopnik, 2005; Mills et al., 2011), we recruited a slightly older age range (4- to 8 year olds) in Experiment 3. We also raised the upper limit by 1 year, as 7 year olds were not at ceiling and the increasing preference for the Active learner found in Experiment 2 might become even more robust beyond age seven.

Consistent with Experiments 1 and 2, we predicted that children would be equally likely to seek help from either learner on the Original toy since both knew how to make that toy play music. For the Similar toy we predicted children would prefer to seek help from the Deliberate learner because they would infer a competence for figuring out this toy from her successful, intentional exploration of the analogous Original toy. On the contrary, even though the Accidental learner's exploration was equally successful, we expected that children would not infer a problem-solving competence based on her stumbling

upon the solution without even trying, as she would not have demonstrated the same problem-solving or exploratory skills as the Deliberate learner. Thus, we thought children would infer that the Accidental learner would be less likely to provide effective help in figuring out the solution to a novel, related problem.

For the Different toy, one might intuitively predict to find a preference for the Deliberate learner. Unlike the Instructed and Passive learners in Experiments 1 and 2 who were interrupted and thus did not have an opportunity to explore, the Accidental learner had the opportunity to but chose not to explore. Thus, insofar as children attribute a lack of motivation to the Accidental learner, they might prefer the Deliberate learner instead. Given younger children's reliance on not being offered help as a cue to guide their help seeking in Experiments 1 and 2, such a preference for the Deliberate learner on the Different toy might also increase with age. However, given the subtlety of the cues to motivation and the lack of clear preference for the Active learner on the Different toy trial in previous experiments, we did not expect children to exhibit a preference for either learner on the Different toy trial in Experiment 3.

Methods

Participants

In Experiment 3, 227 4- to 8 year olds (M (SD) = 75.71 (16.06) months; range: 48.23–107.97 months 49% female) were recruited and tested at the Natural History Museum in Berlin, Germany from June 2018 to May 2019. An additional 8 children were excluded from analyses because they refused to seek help ($n = 2$), due to experimental error ($n = 4$), or due to cognitive disorders ($n = 2$). We selected this sample size so that we would have a large sample size per condition ($n = 227/\text{condition}$), as well as a medium to large sample across the age range (we aimed for approximately 40–50 children per age bin by year) to explore whether children's help seeking might change with age. A post-hoc power analysis simulating a mixed effects logistic regression model with *simr::powerSim* based on the experimental results below revealed that with 227 participants we achieved 66% power for the effects of interest.

Participants were mostly of White European descent and were native or fluent speakers of German. Informed consent was obtained from the parents or legal guardians before commencing the study.

Materials

The toy materials were identical to Experiments 1 and 2. The same Active learner video from Experiment 1 (here referred to as the *Deliberate* learner) was used along with

a new video of the *Accidental* learner described below. The learners in the videos were again White women with brown hair, each wore a different colored t-shirt (blue or yellow) and were referred to by that color (e.g., “My friend Blue”). Each actor always wore the same colored t-shirt, but the learner each actor portrayed (Deliberate or Accidental) was counterbalanced across children.

Procedure

Children were presented with two videos: The *Deliberate* (active) learner (the same Active learner video from Experiment 1) and a video in which an *Accidental* learner, after having expressed curiosity about how the toy works, instead of exploring it, stared at the ceiling and tapped her fingers on the table. Suddenly, she yawned and, while stretching her arms, she accidentally leaned on the toy and activated it after 10s had passed; she expressed surprise and remarked, “Oho! So that is how this toy works.” She said “Oho” instead of “Aha” to emphasize that she unintentionally activated the toy. At the end of the video, the Accidental learner activated the toy a second time on her own, just as in all the other learner videos, making clear that she knew how to activate it (see Figure 4a for more details). Critically, we matched the length of time before the first Activation, as well as the level of surprise expressed by both learners upon activating the toy, as both length of time to complete a task

and surprise have been shown to influence children's inferences about others' competence and prior knowledge (Leonard et al., 2019; Wu & Gweon, 2021). After the help-seeking trials, as in Experiments 1 and 2, children were asked to explain how each learner figured out the Original toy. If children hesitated, the experimenter prompted a response by asking: “Did she figure it out accidentally or deliberately?”

Results

We fit a confirmatory mixed effects Bayesian logistic regression predicting children's learner choice (Deliberate vs. Accidental) with fixed effects of toy trial (categorical, 3-levels: Original, Similar, Different with Original dummy coded as the reference category) and age in months (continuous and centered), including a random intercept by subject. In addition to this model, we also fit a model with simple effects of toy and age, and their interaction, as well as a model with toy as a single predictor. Formal model comparison preferred the model with simple effects of toy and age (i.e., learner choice \sim toy + age + (1|subject)).

If we collapse across age and look at children's choice within each toy trial, children did not exhibit a preference on the Original toy ($\beta = 0.23$, 95% CI $[-0.03, 0.49]$), preferred the Deliberate learner on the Similar toy ($\beta = 0.30$, 95% CI $[0.04, 0.57]$), and of note, preferred the

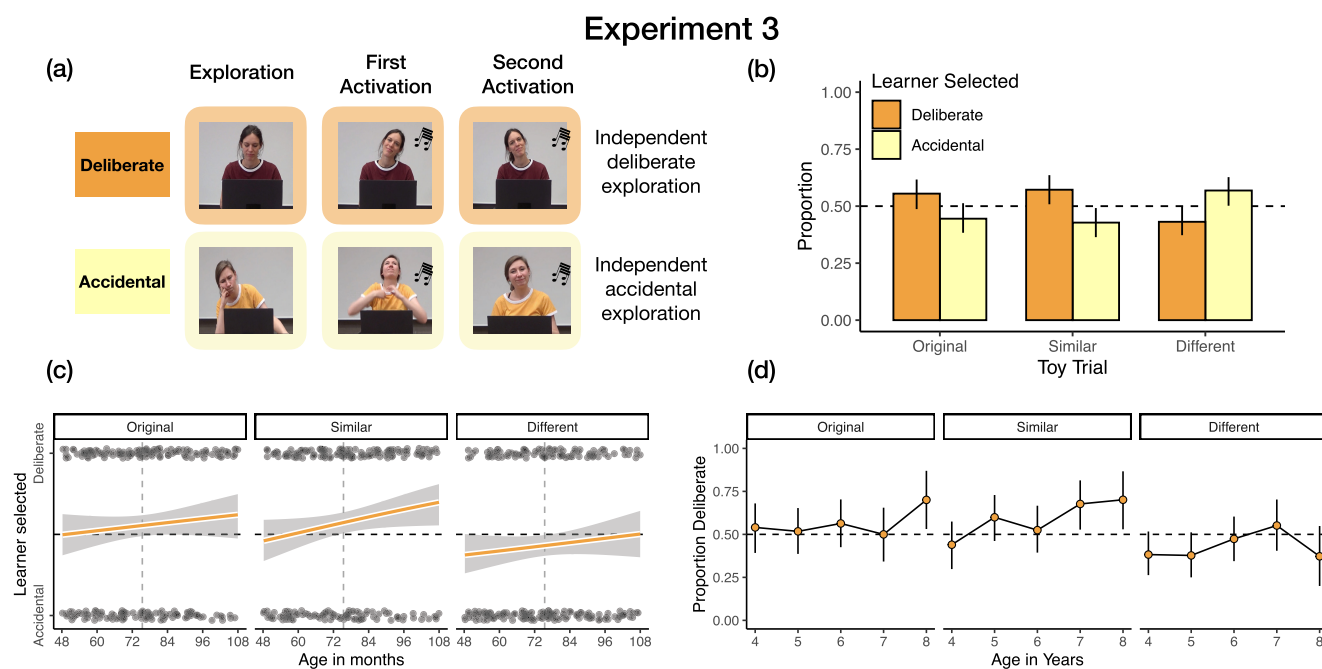


FIGURE 4 Experiment 3: Study design and results. (a) Screenshots from videos used in Experiment 3. The *Deliberate* learner was the same as the *Active* learner in Experiment 1; the *Accidental* learner tapped her hands down and leaned on the toy, it played music, suggesting she had activated it unintentionally. (b) Proportion of children who selected each learner by toy collapsed across age. (c) Children's learner choice by age faceted by toy (each dot represents an individual participant); dashed gray vertical line represents median age, 74.42 months). (d) Discretized version of data in (c): Proportion of children who selected Deliberate learner by age in years faceted by toy. Dashed horizontal line represents chance, 50%. Error bars are bootstrapped 95% CIs.

Accidental learner on the Different toy ($\beta = -0.28$, 95% CI $[-0.55, -0.02]$). This model also revealed that children were more likely to select the Deliberate learner over the Accidental learner on the Similar compared to the Different toy ($\beta = 0.58$, 95% CI $[0.21, 0.96]$) and on the Original compared to the Different ($\beta = 0.51$, 95% CI $[0.14, 0.89]$) but not on the Similar compared to the Original ($\beta = 0.07$, 95% CI $[-0.30, 0.45]$). Children's tendency to select the Deliberate learner also increased with age ($\beta = 0.011$, 95% CI $[0.002, 0.021]$; see [Figure 4b,c](#)).

For the memory questions, 33 children failed to remember how either learner figured out the toy (memory score 0: 14.54%), 42 correctly identified one learner (memory score 1: 18.50%), and 152 correctly identified both learners (memory score 2: 66.96%). Exploratory analyses revealed that a mixed effects Bayesian logistic regression predicting learner choice with the single predictor of toy trial was preferred by formal model comparison over a model with (1) simple effects of toy trial and memory score (integer from 0 to 2) and (2) simple effects of toy trial and memory score, and their interaction. An additional exploratory Bayesian linear regression with age (centered and continuous) predicting memory score indicated that memory score significantly improved with age ($\beta = 0.02$, 95% CI $[0.01, 0.02]$).

Discussion

Children differentiated between the Deliberate and Accidental learners in their help seeking. Like in Experiments 1 and 2, children did not preferentially seek help from either learner on the Original toy and preferred to seek help from the Deliberate learner on the Similar toy (though their tendency to select the Deliberate learner did not significantly differ across these toys). On the Different toy, contrary to our predictions, children preferred the Accidental learner. Children's preference for the Deliberate learner over the Accidental learner, however, increased with age across all three toys.

Upon inspecting the distribution of children's responses, it appears that their preference for the Deliberate learner increased fairly continuously from ages 4 to 8 years. Though an exploratory discretization of the data by age in years suggests that this increase might be most steady on the Similar toy where children reliably sought the Deliberate learner's help by age 7 (67.7% Deliberate, 95% CI $[52.8\%, 81.4\%]$). On the Original toy, children reliably sought the Deliberate learner's help by age 8 (70.1% Deliberate, 95% CI $[53.1\%, 87.0\%]$), but at no age did they reliably seek either learner's help on the Different toy. We are, however, cautious to overinterpret this representation of the data as our experiment was not designed with this analysis in mind (i.e., looking at choice grouped by age in years).

Children's preference for the Deliberate over the Accidental learner on the Similar toy and their overall

strengthening preference for this learner with age, suggests that children were not only paying attention to the outcome that the learner independently generated but also the process by which she achieved this outcome. Both the Deliberate and Accidental learners activated the toy on their own, so if children's help seeking were driven by independent successful action alone, they should not have exhibited a preference for either learner. Even children's unexpected preference for the Accidental learner on the Different toy indicates that children were considering the process and intentionality behind the discovery. It is possible that children interpreted the Accidental learner's discovery as lucky and serendipitous, and so attributed a positive trait to her (Olson et al., 2006) or gave her some extra credit, which weakened their preference for the Deliberate learner on the Similar toy and led to an effect in the opposite direction on the Different toy (at least when collapsing across age), a toy where in previous experiments children were unsure of whom to ask for help.

Though memory score increased with age, model comparison suggested performance on the memory questions did not correlate with children's help seeking. We thus speculate that where younger children differed from older children on this task may be in the inferences they drew from deliberate versus accidental discovery rather than an inability to recognize that these learners arrived at discovery by different means. This interpretation of the age difference is consistent with prior work showing that 3- and 4-year olds differentiate between accidental and intentional demonstrators in their imitation (Bonawitz et al., 2011; Buchsbaum et al., 2011; Butler & Markman, 2012).

Overall, Experiment 3's findings suggest that at least older children's (approximately 6- and 7-year olds') tendency to seek the Active learner's help in Experiments 1 and 2 was not solely based on independent discovery. Rather, it appears that children's preference for the Active learner in the prior experiments was likely increasingly grounded in the *process* of exploring the toy and the intentionality of the discovery, beyond the successful *outcome* that was achieved.

GENERAL DISCUSSION

The presented work provides crucial insights on children's help seeking when encountering novel problems, i.e., problems no one in their environment has yet solved. We find that children's help seeking is informed both by *how* the learner discovered a solution to a different problem in the past, as well as the *similarity* between this past problem and the present problem. Our work builds on prior studies revealing that when children are looking for someone to help them solve a novel problem, they not only use the success of individuals' past actions (i.e., the history of their actions' outcomes, Cluver et al., 2013;

Kushnir et al., 2013), but also the process by which those individuals learned to achieve this success, that is, the *history of their learning*. Beyond preferring individuals with direct versus indirect perceptual access and those who are already knowledgeable versus those who receive help (Einav & Robinson, 2011; Hu et al., 2015), children prefer individuals who may be ignorant a-priori, but are able to discover latent, causal information about the world through their firsthand interventions. Across three experiments, we gather evidence that from 3 to 8 years of age children are sensitive to the *process* by which people learn what they know and prefer to ask for help from agents who have demonstrated an ability to actively, independently, and intentionally discover new information, over those who have learned from someone else or accidentally uncovered the solution.

In particular, in Experiment 1, children preferred to seek help from a learner who had solved a problem through independent, active exploration over a learner who had been shown the solution by someone else and a learner who had watched someone else solve the problem. Children's preference for the Active learner suggests that their help seeking was not driven by the Active learner's knowledge or ability to activate the Original toy (which she shared with the other learners), but rather by the way she had acquired this knowledge/ability. What's more, children's preference for the Active learner over both the Instructed and Passive learners indicates that both direct exploration and successful discovery were important, and not just direct exploration without discovery (Instructed) or secondhand experience of exploration and discovery (Passive). Crucially, children did not indiscriminately prefer the Active learner, regardless of context. Instead, they made selective and flexible choices about whom to ask for help based on whether the Active learner's competence would likely be beneficial for the problem at hand. Specifically, children preferentially sought help from the Active learner on the Similar toy, a novel problem that was related though seemingly more complex than the one the learners had previously solved. However, children did not show any preference when seeking help on the Original toy—the same problem the learners had previously solved—or on the Different toy—a problem quite dissimilar from the learners' problem.

Experiments 2 and 3 further elucidate the aspects of the active-learning process children may be taking into account when assessing a learner's competence, and its relevance to the problem at hand. In Experiment 2, children continued to prefer the Active learner over an Instructed learner on the Similar toy, even when the Active learner was not alone and was offered help. Here, however, children's preference for the Active learner increased from ages 3 to 7, suggesting that in Experiment 1, though there were no significant age effects, the basis for the inference underlying children's preference for the Active learner might have differed from ages 3 to 6.

Specifically, younger children may have relied on the fact that the Active learner was the only one who had been alone and/or for whom no one had thought she needed help, compared to the other learners who had both received assistance, whereas older children may have relied on the actual independence of the Active learner's exploration and discovery. In Experiment 3, from 4 to 8 years of age, children increasingly sought help on all three toys from a Deliberate (active) learner, who intentionally explored and activated the Original toy, compared to an Accidental learner, who did not explore and unintentionally activated the toy. This preference suggests that in Experiments 1 and 2, children's preference for the Active learner was likely based on both her independent exploration and discovery of the Original toy's solution, though younger children's inferences about the learner's problem-solving competence may have been more strongly driven by the discovery (or outcome) than the exploration (or process) that gave rise to it.

Taken together, all three experiments suggest that the cues children use to infer an individual's potential to solve a new problem based on how they learned to solve a problem in the past change across early and middle childhood. When deciding from whom it is best to seek help, younger children may place more weight on whether the individual was previously offered help (regardless of whether or not it was taken) and whether they discovered the solution independently (regardless of whether or not it was intentional). From 3 to 8 years, however, children's reasoning appears to grow more sophisticated: Children can increasingly differentiate between cues that are *correlated with* versus *causally related* to an individual's active-learning competence, preferring an independent explorer and discoverer, even if offered help, and more strongly weighting a goal-directed than an unintentional, lucky discovery. The current experiments cannot determine exactly when or how these shifts occur, but there is evidence that children's tendency to seek the Active (Deliberate) learner's help increased fairly continuously with age, and that at least by age seven they reliably asked this learner for help over the Instructed and Accidental learners. In Experiments 2 and 3, 7- and 8-year-olds' preference was not at ceiling, suggesting that a genuine appreciation of the process of others' learning beyond outcome and an ability to make flexible, context-dependent help-seeking decisions based on this process may continue to develop beyond middle childhood.

These results are in line with the developmental trajectory drawn by previous work on children's understanding of how they and others come to know what they know and related inferences about others' trustworthiness. From 3 to 5 years of age, children become better able to identify the source of their own knowledge (Gopnik & Graf, 1988; Taylor et al., 1994) and to articulate when others' learned from exploration versus instruction (Sobel & Letourneau, 2018); they also become

more likely to seek help from an independently knowledgeable individual versus someone who was assisted (Einav & Robinson, 2011). Children's preference for the Active learner in our experiments increased even beyond age five. This suggests that determining the relative helpfulness of someone who directly versus indirectly discovers abstract information might be harder for children than comparing individuals with knowledge gained via direct visual access versus hearsay (Hu et al., 2015), or who demonstrate prior knowledge versus current ignorance about category labels (Einav & Robinson, 2011). Indeed, previous work on selective trust reveals that children's ability to differentiate and selectively learn from informants based on more nuanced differences in reliability—such as error magnitude, rather than knowledge versus ignorance—develops from 5 to 7 years (Einav & Robinson, 2010). This protracted development in the sophistication of children's help seeking is also in line with evidence suggesting that children's ability to understand and describe learning as a process, and to identify whether knowledge was likely acquired directly versus indirectly improves from 5 to 11 years of age (Lockhart et al., 2016; Sobel & Letourneau, 2015).

Across all three studies, children's preference for the Active (deliberate) learner was strongest when the task they had to solve was analogous to the one that the learner had solved. Indeed, all learners had eventually demonstrated an ability to activate the Original toy, so any learner could help children activate it again. With respect to the Similar and Different toys, it is possible children thought that the problem-solving abilities the Active learner had demonstrated would likely apply to a near-transfer problem (the Similar toy), but not necessarily to a far(ther)-transfer problem (the Different toy).

We argue that in the context of our experiments, children's conservatism seems justified. First children had only witnessed the Active learner figure out one single toy, and they were not given any evidence for the incompetence of the Instructed or Passive learner. Compared to our study, in prior research children were always given more data in which to ground their inference about the relative competence of the potential helpers (i.e., two or more examples, Cluver et al., 2013; Kushnir et al., 2013). Also, in earlier studies the contrast in competence among learners established in the familiarization phase was much stronger than in our experiments. For example, one helper was presented as clearly knowledgeable and successful, while the other clearly failed, claimed ignorance and/or was less socially engaged (e.g., Cluver et al., 2013; Kushnir et al., 2013). The contrast we introduced in our experiments (i.e., the learners were all initially ignorant, but ultimately successful) was arguably more subtle. From this perspective, our task was quite a strict test of the inferences children draw about relative competence and helpfulness from different learning histories. Second, we explicitly stated that the Similar toy was related to the Original toy, while the Different

toy was unrelated, and we offered clear visual cues in support of the toys' (dis)similarity. This information may have increased children's uncertainty about the scope of the Active learner's abilities and led to a more narrow inference about the space of problems she had a greater potential to help solve. Defining the boundaries of a class of problems and knowing when competence in one task is likely to transfer to another is an incredibly important skill for learning to learn (e.g., Brown & Kane, 1988). Future work could explore whether having the Active learner successfully discover the solution to two or more toys, as well as varying the similarity across these demonstrated problems, affects children's willingness to generalize the problem-solving competence they infer.

When deciding whether to seek a learner's help on a novel problem, we find that children were sensitive to whether a learner's previous discoveries had been independent and direct versus assisted and secondhand, as well as whether they had been deliberate versus accidental. Are children sensitive to other cues relevant to an individual's learning and problem-solving abilities—such as the efficiency or the systematicity of an individual's exploratory actions? For example, if children could observe the actual actions the learner took on the toy, would they seek help on a novel problem from an individual who had explored strategically, over one whose actions were apparently random, though equally successful? Might they also prefer to seek help from a previously efficient than an inefficient explorer? Children as young as age 4 can evaluate the informativeness of others' questions (De Simone & Ruggeri, 2020; Ruggeri et al., 2017), 4- and 5 year olds can infer relative competence based on the relative time individuals took to complete the same task (Leonard et al., 2019), and, from ages 5 to 6, children become increasingly sensitive to how informative the evidence a teacher provides is for a given concept (Rhodes et al., 2010). In line with these previous findings and the results of Experiment 3, it is quite possible that children, and particularly 5- to 8 year olds, would also be able to take into account the informativeness and efficiency of others' past exploratory interventions when choosing from whom to seek help. Children's ability to infer others' relative active-learning competence from even more nuanced cues, and how this ability guides their social learning is a rich avenue for future research.

Seeking help from effective learners and explorers not only increases the probability of uncovering new information about the world, but also maximizes opportunities to acquire problem-solving skills that would make you a better independent learner yourself. In this sense, by partnering up with more advanced learners, children not only can learn from the outcome of these learners' discoveries, but also from the process and strategies that led to them. Indeed, good learners can serve as a model of how to perform goal-directed actions that generate useful evidence, how to ask questions that elicit informative

answers from others, or when it is worth exerting effort on a certain task (Frazier et al., 2011; Leonard et al., 2019; Markant & Gureckis, 2014; Mills et al., 2010; Schulz & Bonawitz, 2007). Recent research indicates that young children are sensitive to the quality of others' learning strategies, possibly even before they are able to implement the most efficient strategies on their own. For instance, children as young as 4 years are able to *identify* which agents ask the most informative questions, despite not being able to *generate* such questions until around age seven (De Simone & Ruggeri, 2020; Ruggeri et al., 2017). Children are also sensitive to when others' effort pays off, and understand that the amount of effort needed is a product of both individual skill and task difficulty. For example, they are more likely to imitate an adult's level of effort when the adult's actions are eventually successful (Leonard et al., 2020), and try harder on tasks that were easy for an adult to achieve, compared to tasks that were difficult or impossible for the adult (Lucca et al., 2020). Given these findings, it is quite possible that identifying and seeking help from competent learners not only enables children to achieve their immediate goals, but also serves as a stepping stone to enrich and advance their own toolbox of strategies for epistemic inquiry and future goal achievement.

Children's consistent lack of preference on the Original toy, though, suggests that in our experiments they may have just wanted to learn the activation sequence for each toy, rather than the skills and strategies required to discover that sequence. Indeed, if children had wanted to learn about the *discovery process* more generally, their best bet would have been to team up with the Active learner on the Original toy—as she was the only learner who had demonstrated the ability to figure out this toy through her own (deliberate) exploration. However, children did seek the Active learner's help on the Similar toy. Even though it is possible their only intention was to learn the activation sequence (rather than the process that led to its discovery), if we had actually given children the opportunity to work together with the Active learner to figure out the Similar toy, they eventually would have been privy to her exploratory actions and strategies. In other words, they still would have had the opportunity to acquire inquiry, exploration, and problem-solving skills even if that was not their reason for partnering up with the Active learner in the first place. In this sense, children's ability to identify individuals who are likely able to offer useful assistance in solving a problem may be the first step in recognizing that, in addition to learning solutions, they can also learn strategies from them. Also note that in our experiments the instructions stated that children's goal was to “figure out how to make each toy work”, not to figure out *how to figure out* how to operate the toys. If the goal of the task had been to acquire the skills for figuring out the toys (for instance, because children would need to figure out other toys in the future without assistance), children

might have shown a stronger preference for the Active learner, possibly on all toys.

In our studies, we used help seeking as our dependent measure because we were not only interested in children's inferences about the agents' learning histories, but also in how children used this information to guide their own social learning behavior. Even if the social learning behavior itself was not our focus of interest, help seeking can offer a generative measure for indexing children's sensitivity to differences across individuals. It is action based rather than language based, enabling the same paradigm to be appropriate for a wide age range of participants, with varying vocabulary, linguistic comprehension, and introspective abilities. Moreover, using knowledge for action is motivating, and there is evidence that action- and decision-based measures with relevance to children's own goals may allow researchers to capture earlier competencies and understanding, compared to third-party questions or evaluations (Buttelmann et al., 2009; Setoh et al., 2016; Southgate et al., 2010; Walker et al., 2016; Walker & Gopnik, 2017).

Despite these strengths, there are bounds to what we can conclude from children's help seeking. First, help-seeking paradigms provide only an indirect measure of children's reasoning, as the decision to ask a specific person for help is the consequence of this reasoning, and not a direct read-out (although one might argue that no behavioral dependent variable in psychological research is a direct read-out, after all). In our experiments, the features along which the learners differed and the selective nature of children's preference for the Active learner provide some insight into what children were thinking (i.e., the competence inferred was bounded, and independent, direct exploration and intentional discovery were critical cues). This selectivity also implies that children did not just like the Active learner better or form a positive association with her and/or a negative association with the other learners (otherwise, they should have sought the Active learner's help on all three toys). We are, however, unable to assess exactly what children inferred about the learners, nor why. For example, did children think that the Active learner *knew more* about the Similar toy, that she was *better at discovering* how this type of toy works, or some combination of both? In other words, did they think her exploration conferred or demonstrated deeper mechanical and casual knowledge, or more procedural knowledge and skills for how to go about solving toys? To begin disentangling these possibilities, we could see if children seek help from the Active learner when the task is to fix a broken toy, rather than figure out how to operate a toy. If children think the Active learner has acquired a deeper understanding of the Original toy's mechanism via her learning history, then they should preferentially seek her help to fix both the Original toy and the Similar toy. If, however, they see her exploration only as evidence of procedural knowledge and skills then they should be less likely to seek the Active learner's help when looking for someone to help fix (vs. help figure out) the

Original and Similar toys. Future work could also combine help seeking with explicit questions about what the learners know or the skills they possess to make more precise what children of different ages infer about these learners based on their learning history.

Second, the act of asking someone for help involves several steps, which can complicate conclusions about why children fail to differentiate among potential helpers. In our experiments, to selectively seek help from the Active learner, children had to (1) appreciate that the learners learned how to operate the Original toy in different ways, (2) infer from these different processes that the Active learner is more likely able to solve the problem at hand than the other learner(s), and (3) actually ask the Active learner for help. If children do not ask the Active learner for help, at which step(s) did they encounter problems? We think it is quite unlikely that younger children in our experiments struggled with step 3, given that previous work contrasting helpers in terms of the success of their actions and relative amount or domain of knowledge showed that children as young as age 2 and 3, respectively, were able to plan help-seeking actions in line with their inferences of relative competence (Cluver et al., 2013; Kushnir et al., 2013). We argue, instead, that it is likely a combination of steps 1 and 2. From ages 3 to 8, children are increasingly able to appreciate that the processes by which the learners learned to solve the Original toy differed (step 1). But, above and beyond this ability to represent the processes as distinct, the inferences children draw, based on these processes, about an individual's ability to solve problems and the related likelihood of providing effective help on a given problem are also changing (step 2). More specifically, we conclude from Experiments 2 and 3 that younger and older children pay attention to different cues and/or weight the same cues differently when determining the competence and helpfulness of the learners. Performance on the memory questions provide some insight to children's abilities to recognize that the learners had learned in different ways. Memory performance did increase with age in both Experiments 2 and 3, and in Experiment 2 children with better memory were more likely to seek the Active learner's help. This suggests that the observed differences between younger and older children's decisions may have been, at least partly, due to their difficulty with step 1. At the same time, memory performance was overall high (77% of children in Experiment 2 and 67% in Experiment 3 answered correctly for both learners), and it did not predict choices in Experiment 3. Also, even the youngest children did fairly well (in Experiment 2, 49% of 3-year olds and 69% of 4-year olds answered correctly for both learners, and in Experiment 3, 48% of 4-year olds and 52% of 5-year olds did so), suggesting that even when younger children could differentiate between learners' learning histories, their reasoning about whom was best to ask for help still differed from older children's (step 2).

The demands of our tasks more broadly (e.g., keeping track of two to three learners and asking for help on

three different toys), however, may have masked younger children's abilities to differentiate the learners and consequently their abilities to draw the appropriate inferences and seek help accordingly. Though the memory questions are helpful, we are cautious to overinterpret these results, as they may offer an underestimation of younger children's abilities to distinguish how the learners learned. First, the memory questions came at the end of the experiment, after children had already made three help-seeking decisions. Second, they are language based, and articulating the different processes by which the agents had learned was challenging even for us (especially in Experiment 1). Providing additional scaffolding to support children's ability to distinguish and keep in mind the different learning processes would allow us to more accurately determine the relative contribution of differentiation versus interpretation to the observed age effects in the current experiments.

CONCLUSION

Children not only need to learn a great deal about the world, but they also need to learn *how to learn* about the world. Social learning can facilitate both the acquisition of new knowledge and the skills critical to effective exploration and learning. Building on prior work, especially studies exploring the source of informants' knowledge (e.g., Einav & Robinson, 2011; Hu et al., 2015), our work suggests that children can use others' learning histories to identify good learners—social partners who are effective explorers and problem solvers, who can support the discovery of solutions to novel problems, but also who can serve as role models from whom to learn how to learn. Even though children have much to learn, they face this daunting task with a growing and deepening appreciation of both the process and the outcome of learning, realizing that *what* you know matters, but what sometimes matters even more is *how* you came to know what you know.

AUTHOR CONTRIBUTIONS

All authors contributed to designing the experiments and writing the manuscript. Costanza De Simone and Azzurra Ruggeri facilitated and oversaw data collection. Sophie Bridgers, with assistance from Azzurra Ruggeri, analyzed and visualized the data. Costanza De Simone conducted the reliability analyses.

ACKNOWLEDGMENTS

We thank J. Leonard, B. Meder, M. Pelz, and M. H. Tessler for helpful conversations and insightful comments. We also thank M. Bretzke, C. Kreuzer, and M. Breda for help with data collection and coding, and A. Hubrich for help with reliability coding. We are grateful to the Berlin Natural History Museum and the children and families who participated in this research. This

work was partially funded by a NSF Graduate Research Fellowship (to S. Bridgers).

CONFLICT OF INTEREST STATEMENT

We have no known conflicts of interest to disclose.

DATA AVAILABILITY STATEMENT

The data and analytic code necessary to reproduce the analyses presented here are publicly accessible on the Open Science Framework. The analyses were not preregistered. The full set of materials necessary to attempt to replicate the findings presented here are available upon request from the first author; the learner videos are available on the Open Science Framework. The Open Science Framework project associated with this paper can be viewed at <https://osf.io/mpvab/>.

ORCID

Sophie Bridgers  <https://orcid.org/0000-0003-4522-2697>

Azzurra Ruggeri  <https://orcid.org/0000-0002-0839-1929>

REFERENCES

- Agresti, A., & Min, Y. (2002). Unconditional small-sample confidence intervals for the odds ratio. *Biostatistics*, 3(3), 379–386. <https://doi.org/10.1093/biostatistics/3.3.379>
- Bandura, A., & Walters, R. H. (1977). *Social learning theory* (Vol. 1). Englewood cliffs Prentice Hall.
- Bonawitz, E., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. E. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition*, 120(3), 322–330. <https://doi.org/10.1016/j.cognition.2010.10.001>
- Boyd, R., Richerson, P. J., & Henrich, J. (2011). The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences of the United States of America*, 108(Suppl. 2), 10918–10925. <https://doi.org/10.1073/pnas.1100290108>
- Bridgers, S., Jara-Ettinger, J., & Gweon, H. (2020). Young children consider the expected utility of others' learning to decide what to teach. *Nature Human Behaviour*, 4(2), 144–152. <https://doi.org/10.1038/s41562-019-0748-6>
- Brown, A. L., & Kane, M. J. (1988). Preschool children can learn to transfer: Learning to learn and learning from example. *Cognitive Psychology*, 20(4), 493–523. [https://doi.org/10.1016/0010-0285\(88\)90014-X](https://doi.org/10.1016/0010-0285(88)90014-X)
- 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. <https://doi.org/10.1016/j.cognition.2010.12.001>
- Butler, L. P., & Markman, E. M. (2012). Preschoolers use intentional and pedagogical cues to guide inductive inferences and exploration. *Child Development*, 83(4), 1416–1428. <https://doi.org/10.1111/j.1467-8624.2012.01775.x>
- Buttelmann, D., Carpenter, M., & Tomasello, M. (2009). Eighteen-month-old infants show false belief understanding in an active helping paradigm. *Cognition*, 112(2), 337–342. <https://doi.org/10.1016/j.cognition.2009.05.006>
- Cluver, A., Heyman, G., & Carver, L. J. (2013). Young children selectively seek help when solving problems. *Journal of Experimental Child Psychology*, 115(3), 570–578. <https://doi.org/10.1016/j.jecp.2012.12.011>
- Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, 13(4), 148–153. <https://doi.org/10.1016/j.tics.2009.01.005>
- De Simone, C., & Ruggeri, A. (2020). *What is a good question asker better at? From unsystematic generalization to overgeneralization to adult-like selectivity across childhood*. OSF. <https://doi.org/10.31219/osf.io/6jupw>
- Denison, S., Reed, C., & Xu, F. (2013). The emergence of probabilistic reasoning in very young infants: Evidence from 4.5- and 6-month-olds. *Developmental Psychology*, 49(2), 243–249. <https://doi.org/10.1037/a0028278>
- Denison, S., & Xu, F. (2010). Integrating physical constraints in statistical inference by 11-month-old infants. *Cognitive Science*, 34(5), 885–908. <https://doi.org/10.1111/j.1551-6709.2010.01111.x>
- Einav, S., & Robinson, E. J. (2010). Children's sensitivity to error magnitude when evaluating informants. *Cognitive Development*, 25(3), 218–232. <https://doi.org/10.1016/j.cogdev.2010.04.002>
- 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. <https://doi.org/10.1177/0956797611416998>
- 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. <https://doi.org/10.1111/j.1467-7687.2011.01106.x>
- Gopnik, A., & Graf, P. (1988). Knowing how you know: Young children's ability to identify and remember the sources of their beliefs. *Child Development*, 59, 1366–1371. <https://doi.org/10.2307/1130499>
- Goupil, L., Romand-Monnier, M., & Kouider, S. (2016). Infants ask for help when they know they don't know. *Proceedings of the National Academy of Sciences of the United States of America*, 113(13), 3492–3496. <https://doi.org/10.1073/pnas.1515129113>
- Gweon, H. (2021). Inferential social learning: Cognitive foundations of human social learning and teaching. *Trends in Cognitive Sciences*, 25(10), 896–910. <https://doi.org/10.1016/j.tics.2021.07.008>
- Gweon, H., & Schulz, L. E. (2008). *Stretching to learn: Ambiguous evidence and variability in preschoolers' exploratory play*. In Proceedings of the 30th annual meeting of the Cognitive Science Society, 570–574.
- Gweon, H., & Schulz, L. E. (2011). 16-month-olds rationally infer causes of failed actions. *Science*, 332(6037), 1524. <https://doi.org/10.1126/science.1204493>
- Gweon, H., & Schulz, L. E. (2019). From exploration to instruction: Children learn from exploration and tailor their demonstrations to observers' goals and competence. *Child Development*, 90(1), e148–e164. <https://doi.org/10.1111/cdev.13059>
- Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences of the United States of America*, 107(20), 9066–9071. <https://doi.org/10.1073/pnas.1003095107>
- Hu, J. C., Whalen, A., Buchsbaum, D., Griffiths, T. L., & Xu, F. (2015). *Can children balance the size of a majority with the quality of their information?* In Proceedings of the 37th annual meeting of the Cognitive Science Society.
- Kline, M. A. (2015). How to learn about teaching: An evolutionary framework for the study of teaching behavior in humans and other animals. *Behavioral and Brain Sciences*, 38, e31. <https://doi.org/10.1017/S0140525X14000090>
- Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, 76(6), 1261–1277. <https://doi.org/10.1111/j.1467-8624.2005.00849.x>
- Kuhn, D., & Ho, V. (1980). Self-directed activity and cognitive development. *Journal of Applied Developmental Psychology*, 1(2), 119–133. [https://doi.org/10.1016/0193-3973\(80\)90003-9](https://doi.org/10.1016/0193-3973(80)90003-9)
- Kushnir, T., & Gopnik, A. (2005). Young children infer causal strength from probabilities and interventions. *Psychological Science*, 16(9), 678–683.

- 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), 446–453. <https://doi.org/10.1037/a0031649>
- Kushnir, T., Wellman, H., & Gelman, S. (2008). The role of preschoolers' social understanding in evaluating the informativeness of causal interventions. *Cognition, 107*(3), 1084–1092. <https://doi.org/10.1016/j.cognition.2007.10.004>
- Kushnir, T., Wellman, H. M., & Gelman, S. A. (2009). A self-agency bias in preschoolers' causal inferences. *Developmental Psychology, 45*(2), 597–603. <https://doi.org/10.1037/a0014727>
- Legare, C. H. (2011). Exploring explanation: Explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children. *Child Development, 83*(1), 173–185. <https://doi.org/10.1111/j.1467-8624.2011.01691.x>
- Leonard, J. A., Bennett-Pierre, G., & Gweon, H. (2019). Who is better? Preschoolers infer relative competence based on efficiency of process and quality of outcome. In A. K. Goel, C. M. Seifert, & C. Freksa (Eds.), *Proceedings of the 41st annual conference of the Cognitive Science Society* (pp. 639–645). Cognitive Science Society.
- Leonard, J. A., Garcia, A., & Schulz, L. E. (2020). How adults' actions, outcomes, and testimony affect preschoolers' persistence. *Child Development, 91*(4), 1254–1271. <https://doi.org/10.1111/cdev.13305>
- Liszowski, U. (2013). Using theory of mind. *Child Development Perspectives, 7*(2), 104–109. <https://doi.org/10.1111/cdep.12025>
- 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. <https://doi.org/10.1111/cdev.12469>
- Lucca, K., Horton, R., & Sommerville, J. A. (2020). Infants rationally decide when and how to deploy effort. *Nature Human Behaviour, 4*(4), 372–379. <https://doi.org/10.1038/s41562-019-0814-0>
- Markant, D. B., & Gureckis, T. M. (2014). Is it better to select or to reject? Learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General, 143*(1), 94–122. <https://doi.org/10.1037/a0032108>
- McElreath, R. (2020). *Statistical rethinking: A bayesian course with examples in R and Stan*. CRC Press.
- 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. <https://doi.org/10.1080/15248372.2010.516419>
- 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. <https://doi.org/10.1016/j.jecp.2011.06.003>
- Morey, R., Hoekstra, R., Rouder, J., Lee, D. M., & Wagenmakers, E. (2016). The fallacy of placing confidence in confidence intervals. *Psychonomic Bulletin & Review, 23*, 103–123. <https://doi.org/10.3758/s13423-015-0947-8>
- Nicenboim, B., Schad, D., & Vasisht, S. (2021). *An introduction to Bayesian data analysis for cognitive science*. Chapman and Hall/CRC Statistics in the Social and Behavioral Sciences Series.
- Olson, K. R., Banaji, M. R., Dweck, C. S., & Spelke, E. S. (2006). Children's biased evaluations of lucky versus unlucky people and their social groups. *Psychological Science, 17*(10), 845–846.
- Pasquini, E. S., Corriveau, K. H., Koenig, M., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology, 43*(5), 1216–1226. <https://doi.org/10.1037/0012-1649.43.5.1216>
- Piaget, J., & Cook, M. (1952). *The origins of intelligence in children* (Vol. 8). International Universities Press.
- Rhodes, M., Gelman, S., & Brickman, D. (2010). Children's attention to sample composition in learning, teaching and discovery. *Developmental Science, 13*(3), 421–429. <https://doi.org/10.1111/j.1467-7687.2009.00896.x>
- Rowles, S. P., & Mills, C. M. (2018). Preschoolers sometimes seek help from socially engaged informants over competent ones. *Cognitive Development, 48*, 19–31. <https://doi.org/10.1016/j.cogdev.2018.06.006>
- 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. <https://doi.org/10.1037/dev0000340>
- Ruggeri, A., Swaboda, N., Sim, Z. L., & Gopnik, A. (2019). Shake it baby, but only when needed: Preschoolers adapt their exploratory strategies to the information structure of the task. *Cognition, 193*, 104013. <https://doi.org/10.1016/j.cognition.2019.104013>
- Schulz, L. E. (2012). The origins of inquiry: Inductive inference and exploration in early childhood. *Trends in Cognitive Sciences, 16*(7), 382–389. <https://doi.org/10.1016/j.tics.2012.06.004>
- Schulz, L. E., & Bonawitz, E. (2007). Serious fun: Preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology, 43*(4), 1045–1050. <https://doi.org/10.1037/0012-1649.43.4.1045>
- Schulz, L. E., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science, 10*(3), 322–332. <https://doi.org/10.1111/j.1467-7687.2007.00587.x>
- Setoh, P., Scott, R. M., & Baillargeon, R. (2016). Two-and-a-half-year-olds succeed at a traditional false-belief task with reduced processing demands. *Proceedings of the National Academy of Sciences of the United States of America, 113*(47), 13360–13365. <https://doi.org/10.1073/pnas.1609203113>
- Shafiq, P., Goodman, N. D., & Frank, M. C. (2012). Learning from others: The consequences of psychological reasoning for human learning. *Perspectives on Psychological Science, 7*(4), 341–351. <https://doi.org/10.1177/1745691612448481>
- Sierksma, J., & Shutts, K. (2020). When helping hurts: Children think groups that receive help are less smart. *Child Development, 91*(3), 715–723. <https://doi.org/10.1111/cdev.13351>
- Sim, Z. L., Mahal, K., & Xu, F. (2017). Learning about causal systems through play. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. J. Davelaar (Eds.), *Proceedings of the 39th annual conference of the Cognitive Science Society* (pp. 1078–1083). Cognitive Science Society.
- Sim, Z. L., & Xu, F. (2017). Learning higher-order generalizations through free play: Evidence from 2- and 3-year-old children. *Developmental Psychology, 53*(4), 642–651.
- Sobel, D. M., & Corriveau, K. H. (2010). Children monitor individuals' expertise for word learning. *Child Development, 81*(2), 669–679. <https://doi.org/10.1111/j.1467-8624.2009.01422.x>
- Sobel, D. M., & Letourneau, S. M. (2015). Children's developing understanding of what and how they learn. *Journal of Experimental Child Psychology, 132*, 221–229. <https://doi.org/10.1016/j.jecp.2015.01.004>
- Sobel, D. M., & Letourneau, S. M. (2018). Preschoolers' understanding of how others learn through action and instruction. *Child Development, 89*(3), 961–970. <https://doi.org/10.1111/cdev.12773>
- Sobel, D. M., & Sommerville, J. (2010). The importance of discovery in children's causal learning from interventions. *Frontiers in Psychology, 1*, 176. <https://doi.org/10.3389/fpsyg.2010.00176>
- Sommerville, J., Woodward, A., & Needham, A. (2005). Action experience alters 3-month-old infants' perception of others' actions. *Cognition, 96*(1), B1–B11. <https://doi.org/10.1016/j.cognition.2004.07.004>
- Southgate, V., Chevallier, C., & Csibra, G. (2010). Seventeen-month-olds appeal to false beliefs to interpret others' referential communication. *Developmental Science, 13*(6), 907–912. <https://doi.org/10.1111/j.1467-7687.2009.00946.x>

- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, *348*(6230), 91–94. <https://doi.org/10.1126/science.aaa3799>
- Taylor, M., Esbensen, B. M., & Bennett, R. T. (1994). Children's understanding of knowledge acquisition: The tendency for children to report that they have always known what they have just learned. *Child Development*, *65*(6), 1581–1604. <https://doi.org/10.1111/j.1467-8624.1994.tb00837.x>
- Tomasello, M. (2009). *Why we cooperate*. MIT Press.
- Tomasello, M., Kruger, A., & Ratner, H. (1993). Cultural learning. *Behavioral and Brain Sciences*, *16*(3), 495–552. <https://doi.org/10.1017/S0140525X0003123X>
- Vredenburg, C., & Kushnir, T. (2016). Young children's help-seeking as active information gathering. *Cognitive Science*, *40*(3), 697–722. <https://doi.org/10.1111/cogs.12245>
- Vygotsky, L. S. (1997). *The collected works of LS Vygotsky: Problems of the theory and history of psychology* (Vol. 3). Springer Science & Business Media.
- Walker, C. M., Bridgers, S., & Gopnik, A. (2016). The early emergence and puzzling decline of relational reasoning: Effects of knowledge and search on inferring abstract concepts. *Cognition*, *156*, 30–40. <https://doi.org/10.1016/j.cognition.2016.07.008>
- Walker, C. M., & Gopnik, A. (2017). Discriminating relational and perceptual judgments: Evidence from human toddlers. *Cognition*, *166*, 23–27. <https://doi.org/10.1016/j.cognition.2017.05.013>
- Wu, Y., & Gweon, H. (2021). Preschool-aged children jointly consider others' emotional expressions and prior knowledge to decide when to explore. *Child Development*, *92*(3), 862–870. <https://doi.org/10.1111/cdev.13585>
- Xu, F., & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(13), 5012–5015. <https://doi.org/10.1073/pnas.0704450105>
- Xu, F., & Kushnir, T. (2013). Infants are rational constructivist learners. *Current Directions in Psychological Science*, *22*(1), 28–32. <https://doi.org/10.1177/0963721412469396>

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Bridgers, S., De Simone, C., Gweon, H., & Ruggeri, A. (2023). Children seek help based on how others learn. *Child Development*, *94*, 1259–1280. <https://doi.org/10.1111/cdev.13926>