nature human behaviour

ARTICLES
https://doi.org/10.1038/s41562-019-0748-6

Young children consider the expected utility of others' learning to decide what to teach

Sophie Bridgers[1*], Julian Jara-Ettinger[2] and Hyowon Gweon[1*]

Direct instruction facilitates learning without the costs of exploration, yet teachers must be selective because not everything can nor needs to be taught. How do we decide what to teach and what to leave for learners to discover? Here we investigate the cognitive underpinnings of the human ability to prioritize what to teach. We present a computational model that decides what to teach by maximizing the learner's expected utility of learning from instruction and from exploration, and we show that children (aged 5–7 years) make decisions that are consistent with the model's predictions (that is, minimizing the learner's costs and maximizing the rewards). Children flexibly considered either the learner's utility or their own, depending on the context, and even considered costs they had not personally experienced, to decide what to teach. These results suggest that utility-based reasoning may play an important role in curating cultural knowledge by supporting selective transmission of high-utility information.

Humans actively explore their surroundings and learn from their own experience. Even young children direct their own learning through exploratory play, and update their beliefs from self-generated evidence[1–4]. However, exploration involves uncertainty; learners may not know the time and effort required to make a discovery, and even after much trial and error, they may fail to discover anything at all. Such uncertainty can make it especially challenging for novice learners to decide what or when to explore, because they often lack the knowledge or experience to estimate the expected costs and rewards of exploration.

Social learning provides an effective solution to this problem. Teaching, in particular, is a powerful way to facilitate learning. When knowledgeable individuals teach what they know, naïve learners can avoid the uncertainties of exploration and acquire useful knowledge, even when self-guided discovery is too costly[5]. For instance, a hunter-gatherer who already knows how to make fire can save others from hours of trial and error by showing them how to make it. However, teaching requires an investment of time, effort and resources[6], making it necessary for teachers to be selective in what they choose to teach. Indiscriminately teaching anything would be inefficient; not all of one's knowledge is useful for others, and some knowledge can be acquired easily without being taught, or only through direct experience. Relying on learners' requests to decide what to teach would also be ineffective; ignorant learners may be unaware of what they need to learn or unable to communicate their requests. Thus, while teaching yields significant benefits for learners, it also presents a decision problem for teachers: what needs to be taught, and what can be left for learners to discover on their own?

As adults, we share the intuition that it is better to teach things that are useful, novel, and interesting for others (i.e., rewarding), or difficult, risky or even impossible to discover by oneself (i.e., costly). Although what humans find useful, interesting, or difficult can differ between cultures, change across generations and develop over a lifetime, the tendency to prioritize information that one considers to be rewarding and costly to acquire may be grounded in basic social cognition. We propose that the human ability to decide what to teach is rooted in early-developing capacities to reason about other minds, not only in terms of what other people know or want (theory of mind)[7], but also what is rewarding for other people to acquire and what is costly to achieve on their own (naïve utility calculus)[8]. We formalize these intuitions as decisions that maximize learners' expected utilities, and demonstrate that even young children can decide what to teach based on an intuitive cost–benefit analysis of learning from instruction versus learning from exploration—arguably the most and least teacher-driven forms of learning.

Unlike observational learning or imitation, instruction involves a knowledgeable teacher who deliberately provides helpful information for the learner, often accompanied by ostensive cues (for example, pointing) that signal pedagogical communication of information. Some theories suggest that young children are prepared to interpret information in these contexts as conveying generic, shared knowledge[9], while others emphasize the role of mutual, recursive mental-state reasoning between the teacher and the learner that constrains the learner's inferences[10,11]. Yet the majority of theoretical and empirical work on human pedagogy has focused on how learners benefit from teaching, with an implicit assumption that the teacher already knows what knowledge or concept needs to be taught. We address a question that complements this previous work: how do teachers decide what to teach in the first place?

Our work is grounded in a broader computational framework that has successfully captured how humans interpret others' behaviours through a naïve utility calculus[8]: people have an intuitive theory of agents as 'utility maximizers', expecting agents to pursue rewarding goals while minimizing the associated costs[12–14]. Critically, however, agents can only maximize their utilities insofar as they have an accurate representation of the costs and rewards of their actions; a learner who is ignorant of, or even mistaken about, the cost–reward structure of the world might incur significant costs only to gain little to no reward. The key insight behind our model is that a 'teacher'—a knowledgeable agent who already understands the cost–reward structure of the world—can help the learner reduce or eliminate the costs and uncertainties of exploration. By selectively teaching what is rewarding and/or costly to learn, the teacher can ensure that the

[1]Department of Psychology, Stanford University, Stanford, CA, USA. [2]Department of Psychology, Yale University, New Haven, CT, USA.
*e-mail: sbridge@stanford.edu; gweon@stanford.edu

learner reaps the rewards of learning while also bypassing potentially extreme costs. Importantly, such prioritization of high-utility information still allows the learner to explore and make additional discoveries, even if the teacher cannot teach everything.

Following the focus of our computational model, our experiments are designed to test young children's ability to decide what it is best to teach and what can be left for the learner to discover on their own. Anticipating what information would be rewarding for others and how costly it would be for them to acquire it on their own involves a host of cognitive capacities that go far beyond a general motivation to help or teach; one must reason about the process by which a naïve learner might explore their environment, and understand what information would be more rewarding for the learner, or more costly for the learner to discover without being taught. Thus, one might assume that such sophisticated decision-making abilities would take years of experience or even explicit training to develop. However, there are reasons to believe that the prerequisite representational and inferential capacities for such reasoning are present in childhood.

First, children, as learners themselves, selectively engage in active exploration, especially when additional discovery is likely[4,11,15,16], and rationally decide whether to explore or seek help[17,18]. This work suggests that an intuitive understanding of the relative benefits of social learning and self-guided exploration may be present early in life. Second, by the beginning of formal schooling, children already have an abstract understanding of what constitutes informative teaching (that is, the information communicated should be accurate and sufficient, yet not superfluous), and readily evaluate others' teaching accordingly[19–23]. As teachers themselves, children can also tailor the content of their communication based on the learner's goal or knowledge state[24–26], readily selecting appropriate evidence to communicate a concept[27], correct another's false belief[28] or disambiguate a causal system[22,29]. In particular, when evidence is physically costly to generate (for example, demonstration requires multiple actions or walking over a distance), 5- to 6-year-olds resist providing unnecessary information[22,29], suggesting that children are sensitive not only to the value of information but also the cost of generating it.

Yet a hallmark of helpful, effective teaching is its potential to make learning less costly from the learner's perspective; a tendency to minimize the overall cost of generating information could emerge solely from the teacher's motivation to minimize their own costs, rather than a regard for the learner. Given that even infants and young children expect others to act in ways that maximize their expected utilities[30–34], an important open question is whether children can actively choose to maximize others' utilities. Consistent with this possibility, one study has reported that children prioritize teaching what was taught by an adult (over what children discovered on their own) when the taught information is causally opaque and thus difficult to discover[35]. However, children in this study could have privileged this information because it was provided by a knowledgeable adult or because the taught information was novel or surprising to them. To test whether young children have a genuine understanding of the utilities of others' learning, it is critical to distinguish between the role of reward (that is, the pleasure or benefit of knowing something) and the role of cost (that is, the expected time or effort required to acquire the knowledge) in children's decisions.

Combining developmental and computational approaches, the current work provides a unifying utility-based framework for how humans decide what to teach. While previous utility-based models of social reasoning have focused on formalizing people's inferences about others' actions or mental states[8,13,14], here we are concerned with how these inferences inform people's decisions about what it is best to teach. Using a family of computational models, we generate precise, quantitative predictions about how a rational teacher decides what to teach depending on how this choice modulates others' expected utilities. We then ask whether the ability to consider the expected costs and rewards of learning (from exploration and from instruction) is present even early in childhood, by comparing children's decisions with the model predictions. We chose 5–7 years as our target age range, because children of this age consider others' mental states to decide what to communicate as teachers[22,28,29], and can infer and integrate expected costs and rewards when reasoning about others' goal-directed actions[8].

Consider a scenario in which an agent (the child participant) has to decide which of two causal devices a teacher (the experimenter) should teach to a naïve learner (note that children were not asked to teach the learner themselves; see Methods). Each device has a distinct causal mechanism that generates an effect, and the child and the experimenter already know how to activate both devices. The learner will learn how to activate the device that is chosen by the participant, but because only one device can be taught, they will also have to figure out how to activate the other device through their own exploration. Formally, this decision can be described as a choice between two possible teaching plans (that is, teach device X and let the learner explore device Y, or vice versa) that differ in their expected utilities from the learner's perspective. For each teaching plan, the learner's expected utility can be described as the difference between their expected rewards and costs of learning one device through instruction and learning the other through exploration (Fig. 1a). By choosing to teach the device that generates a more rewarding effect or has a causal mechanism that is more costly to discover, the participant can maximize the learner's expected utility.

In our experiments, the devices were causal toys that varied in discovery reward (high or low) and discovery cost (low, medium, high or extra-high). High-reward toys had an orb that spun around and lit up in different colours when activated. Low-reward toys played music. Low-cost toys had one big button (enabling almost immediate discovery). Medium- and high-cost toys had one big button and six small buttons. The medium-cost toy required pressing one particular small button to activate it and thus was less difficult to figure out than the high-cost toys, which required pressing the big button and a particular small button at the same time. Activating the extra-high-cost toy also required simultaneously pressing the one big button and one particular small button, but it had 11 small buttons, entailing much trial and error before discovery.

Our main and alternative models formalize different hypotheses about which variables children consider when deciding which toy the experimenter should teach. The main (full) model assumes children will consider both the costs and rewards of the toy that would be taught (that is, the chosen toy's activation cost and activation reward), as well as the costs and rewards of the toy the learner would explore on their own (that is, the unchosen toy's discovery cost and discovery reward), and select the teaching plan with the higher net utility as the best one to execute.

Four alternative models consider subsets of these variables and represent possible developmental changes in children's reasoning. While previous work has shown an early-developing understanding of others' costs and rewards[31,32,36], as well as the ability to reason about uncertain future events[37–39], it is possible that children do not consider all of the variables represented in the full model to decide what to teach. Children might consider some variables to be less relevant, or lack the representational or processing capacities to integrate them into a single decision. For instance, since it is easy to activate the toys and obtain their rewards once one knows how they work, children might just focus on the learner's discovery costs (the costs-only model). Alternatively, since children already know how the toys work, they might fail to consider the discovery costs for a naïve learner and only consider the rewards (the rewards-only model). It is also possible that children consider both costs and rewards but fail to represent the utilities of both instruction and

a

Activation reward (R_A) | Activation cost (C_A) | Discovery reward (R_D) | Discovery cost (C_D)

$$U(\text{teach}_X) = R_A(\text{toy}_X) - C_A(\text{toy}_X) + P_E(R_D(\text{toy}_Y) - C_D(\text{toy}_Y))$$

$$U(\text{teach}_Y) = R_A(\text{toy}_Y) - C_A(\text{toy}_Y) + P_E(R_D(\text{toy}_X) - C_D(\text{toy}_X))$$

Chosen toy (instructed) | Unchosen toy (explored)

b

Model	Activation reward	Activation cost	Discovery reward	Discovery cost
Full model	✓	✓	✓	✓
Costs only	X	✓	X	✓
Rewards only	✓	X	✓	X
Instruction only	✓	✓	X	X
Exploration only	X	X	✓	✓

Fig. 1 | Full model equation and table of model space. a, Equations of the costs and rewards considered in each teaching plan in the full model. Activation reward and discovery reward are assumed to be equal. The discovery cost is always greater than the activation cost, except for the low-cost toy whose activation and discovery costs are 1. The cost function is assumed to be the physical cost of pressing buttons and is based on the toys' physical properties. The toys' rewards were sampled from beta distributions and grounded in children's relative preferences for the lights and the music (see Methods for more details). P_E is the probability that the learner will explore the unchosen toy. Alternative models consider different subsets of these variables, as depicted in **b**. **b**, Various components used by the full model and four alternative (lesioned) models to calculate the utility of each teaching plan. Each row is a model and each column is a component participants might consider when deciding what to teach. The full model includes all components, whereas each alternative model includes a subset.

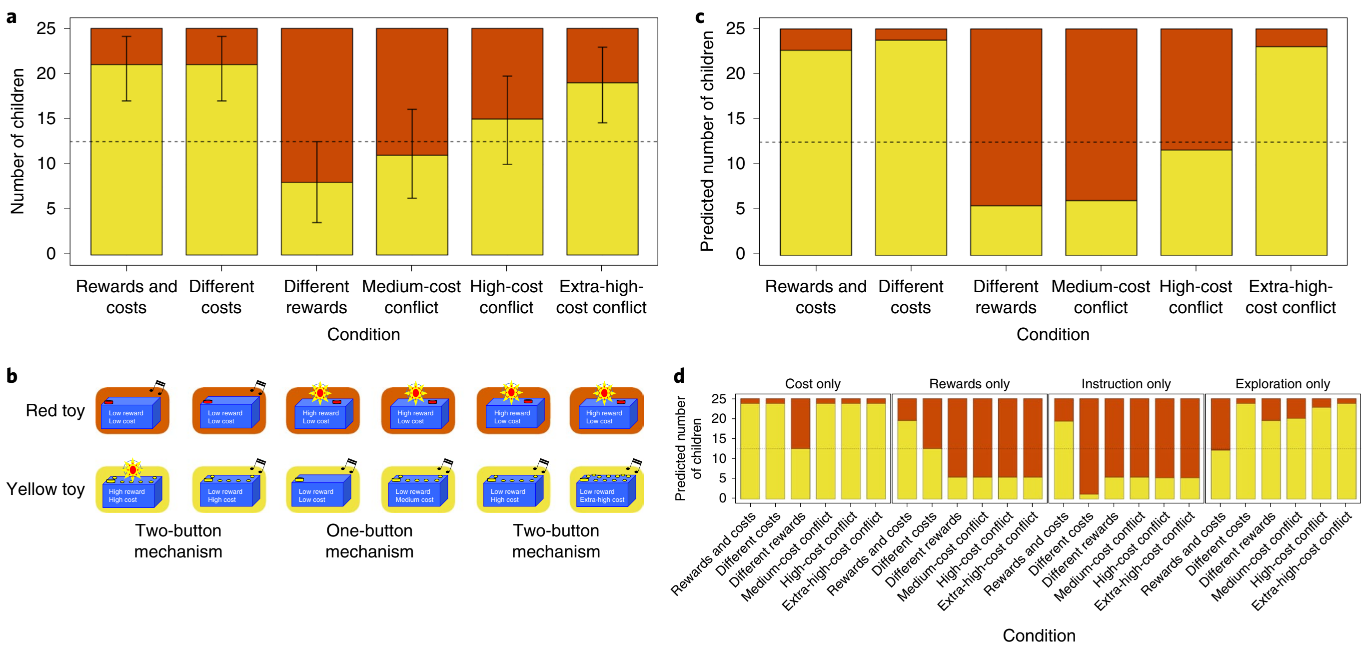

Fig. 2 | Experiment 1: stimuli, behavioural results and model predictions. a, Choices made by children in experiment 1 ($n = 25$ per condition). Red represents the number of children who selected the red toy, as shown in the top row of **b**, while yellow represents the number of children who selected the yellow toy, as shown in the bottom row of **b**. Error bars are bootstrapped 95% CIs. **b**, Schematic of the toys used in each condition of experiment 1 and as shown in **a**, **c**, and **d**. **c**, The full model's predictions of children's choices. **d**, Predictions made by each alternative model.

exploration; they may only concern themselves with the immediate consequence of what would be taught (the instruction-only model), or just focus on the trial and error the learner will experience trying to activate the unchosen toy (the exploration-only model). Figure 1b summarizes the different considerations of the main and alternative models (see Methods and Supplementary Note 1 for more details).

In experiment 1, we tested the model predictions across six between-subjects conditions that presented children with different pairs of toys (Fig. 2b). The choice was over-determined in the first condition: (1) rewards and costs: one toy was higher in both cost and reward than the other toy (high-reward/high-cost toy versus low-reward/low-cost toy). In the next two conditions, either the cost or the reward varied across toys while the other was matched: (2) different costs: the toys were matched in reward but varied in cost (low-reward/high-cost toy versus low-reward/low-cost toy); and (3) different rewards: the toys were matched in cost but varied in reward (high-reward/low-cost toy versus low-reward/low-cost toy). In the remaining three conditions, a high-reward/low-cost toy was contrasted with a low-reward toy that varied in its discovery cost: (4) medium-cost conflict; (5) high-cost conflict; and (6) extra-high-cost conflict. These three conditions forced children to trade-off costs and rewards because one toy was higher in rewards while the other was higher in discovery costs. Formal model comparison will allow us to determine how well our full model captures the pattern of empirical data relative to the simpler alternatives. We also present additional experiments showing that children flexibly consider the learner's utility or their own, depending on the context (experiment 2), and that they can infer the expected costs of exploration, even without having personally experienced these costs (experiment 3).

Results

Experiment 1. The full model predicts that children will prefer the high-cost/high-reward toy in the rewards and costs condition, because this choice maximizes the learner's utility by both increasing their rewards (ensures they will enjoy the high-reward effect) and decreasing their costs (they will not have to go through the trouble of figuring out the high-cost toy). In the different costs and different rewards conditions, it predicts a preference for the high-cost toy and high-reward toy, respectively. However, note that the quantitative predictions are not symmetrical across these two conditions; the model assumes that the relative difference in the toys' rewards (reward from seeing the lights versus hearing the music) in the different rewards condition is more variable and smaller than the difference in the toys' discovery costs in the different costs condition. This assumption reflects children's intuitions about which effect is 'cooler' (see Supplementary Methods 2 and Supplementary Results 2). In our conflict conditions, the model predicts that children's tendency to select the low-reward toy over the high-reward toy will increase as a function of the magnitude of the relative difference in the toys' discovery costs.

In experiment 1, children's ($n=25$ per condition; mean $(M)\pm$ s.d. $=6.35\pm0.97$ years of age; 53% female) choices were highly consistent with the full model's predictions. In the conditions where rewards and costs were not in conflict, children showed a tendency to choose toys that were either higher in both rewards and costs (rewards and costs condition: 84% (21/25); 95% confidence interval (CI) = 0.68–0.96), higher in costs (different costs condition: 84% (21/25); 95% CI = 0.68–0.96) or higher in rewards (different rewards condition: 68% (17/25); 95% CI = 0.50–0.86). All reported CI's were bootstrapped using 10,000 samples.

In the conditions involving a trade-off between costs and rewards, children's tendency to select the low-reward toy over the high-reward toy increased as the expected discovery cost of the low-reward toy increased. While only 44% (11/25; 95% CI = 0.25–0.64) selected the low-reward toy in the medium-cost conflict condition, this increased to 60% (15/25; 95% CI = 0.40–0.79) in the high-cost conflict condition and 76% (19/25; 95% CI = 0.58–0.92) in the extra-high-cost conflict condition. Children's preference for the toy higher in cost increased linearly across these conditions (logistic regression: chose high-reward toy ~ conflict condition; $\beta=-0.226$; $t(74)=-2.348$; $P=0.0216$). See Fig. 2a,c for behavioural data and full model predictions (for more details, see Supplementary Results 1).

To evaluate how well each model captures children's responses, we calculated the likelihood of each model generating the pattern of results across all six conditions. Likelihood ratios between the full model and each alternative model represent how well the full model explains the data relative to each alternative. All comparisons favoured the full model by at least five orders of magnitude (all likelihood ratios $>1\times10^5$), providing compelling support for the full model over the four alternatives (see Fig. 3). For further details on how each alternative model failed to capture the behavioural data, see Supplementary Note 1.

The tight correspondence between the behavioural results and the full model suggests that children in our task coordinated multiple considerations for a learner to decide what to teach. Children did not consider just the expected rewards or just the expected costs for the learner, nor did they reason about just the utility of instruction or just the utility of exploration. Rather, they considered the relative expected costs and rewards for a learner under different teaching plans (that is, learning one toy from instruction and the other from exploration, or vice versa) and chose the plan that maximized the learner's expected utility.

Experiment 2. Our experimental task was designed to minimize the influence of children's own utilities on their decisions; children did not demonstrate the toys themselves to the learner, but instead made a decision about which toy the experimenter should teach. Note also that at the time of this decision, the cost of generating an effect with the medium-cost, high-cost and extra-high-cost toys was low, similar to that of the alternative low-cost toy, and was always lower for the children than for the learner (that is, activation cost < discovery cost, especially for the high- and extra-high-cost toys). Thus, the fact that children prioritized toys that were relatively higher in cost suggests that they considered the learner's expected discovery costs.

However, it is possible that children preferred teaching the higher-cost toy not because it would be harder for the learner to figure out but because their own previous success in figuring out how it worked made it more attractive than the low-cost toy. Alternatively, the existence of multiple buttons (despite being inert) may have made the higher-cost toy more fun to play with or more visually appealing than the low-cost toy with a single button. If children's choices are based on their own rewards, children might show a preference for the relatively higher-cost toy, even when there is no learner to consider, such as when they are choosing a toy with which they themselves would like to play. However, if children's decisions about what it is best to teach are based on a genuine consideration of the learner's expected costs of exploration, they should prioritise high-cost toys when they are choosing a toy for the learner but not when choosing a toy for themselves. We tested this prediction in experiment 2.

Children ($n=25$ per condition; $M\pm$ s.d. $=6.19\pm0.77$ years of age; 54% female) were asked to choose a toy to teach a learner (teach condition), or for themselves to play with (play condition), between a pair of toys that pit the high-cost mechanism against the (intended) high-reward effect (high-reward/low-cost toy versus low-reward/high-cost toy). The teach condition was a replication of experiment 1's high-cost conflict condition, whereas the play condition was identical to the teach condition except for the final prompt: children were asked with which toy they would like to play.

As predicted, children's choice of toys differed across conditions: they were more likely to choose the low-reward/high-cost toy in the teach condition than in the play condition (72% (18/25) versus 20% (5/25); two-tailed Fisher's exact test, $P<0.001$; odds ratio (OR) = 9.71; see Fig. 4). In fact, the majority of children in the play condition chose the high-reward/low-cost toy (80% (20/25); two-tailed binomial test, $P=0.004$; 95% CI = 0.63–0.95), reflecting children's strong preference for the lights effect. In contrast, the majority of children in the teach condition chose the low-reward/high-cost toy (72% (18/25); two-tailed binomial test, $P=0.043$; 95% CI = 0.53–0.89; no significant difference from experiment 1's high-cost conflict condition: two-tailed Fisher's exact test, $P=0.55$; OR = 1.70).

These results further support the interpretation of children's choices in experiment 1 in terms of the full model: their decisions to teach the higher-cost toys probably reflected a concern for the learner's expected cost of exploration rather than a personal preference for these toys. Furthermore, children's decisions were flexible depending on the context: they prioritized the learner's utilities over their own when deciding what to teach, but without a learner to consider, children readily considered their own preferences.

Experiment 3. Together, experiments 1 and 2 suggest that children were sensitive to what would be difficult or time-consuming for others to learn from exploration and understood that teaching can effectively reduce or eliminate such costs. However, it remains unclear whether children's decisions are based on their own past experience of struggling to explore the toys with high discovery costs, or the ability to infer the cost that others would incur to make a meaningful discovery even in the absence of direct previous experience with these costs.

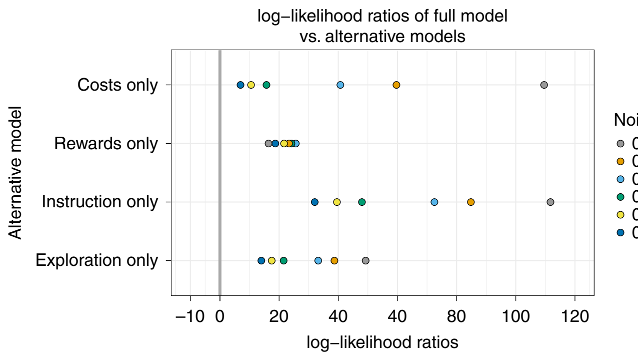

Fig. 3 | Experiment 1: model comparison by noise parameter. Log-likelihood ratios for the full model versus the alternative models (log[L(data|full model)/L(data|alternative model)], where L is the likelihood function) as a function of different noise parameters. Values above 0 (meaning that the likelihood ratio is >1) reflect that the full model explains the data better than the alternative model. Values below 0 (meaning that the likelihood ratio is <1) indicate that the alternative model explains the data better than the full model. The full model better explains the data compared with the four alternative (lesioned) models for all noise-parameter values. For the model comparison reported in the main text, we used a small degree of noise ($\alpha = 0.1$), as shown in light blue, and all comparisons favoured the full model by at least five orders of magnitude (all likelihood ratios > 1×10^5).

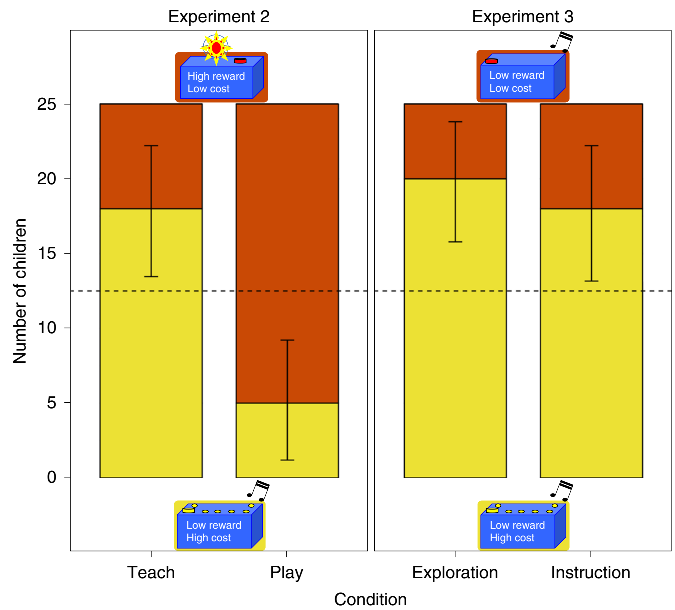

Fig. 4 | Experiments 2–3: behavioural results. Bar plots show children's responses in experiment 2's teach and play conditions (left; $n = 25$ per condition) and experiment 3's exploration and instruction conditions (right; $n = 25$ per condition). Yellow represents the number of children selecting the yellow toys as best to teach (or play with), while red represents the number of children selecting the red toys. Error bars are bootstrapped 95% CIs. In experiment 2, children more often chose the low-reward/high-cost toy in the teach condition than in the play condition (two-tailed Fisher's exact test, $P < 0.001$; OR = 9.71). Children in the play condition chose the high-reward/low-cost toy (two-tailed binomial test, $P = 0.004$; 95% CI = 0.63–0.95), while children in the teach condition chose the low-reward/high-cost toy (two-tailed binomial test, $P = 0.043$; 95% CI = 0.53–0.89). In experiment 3, children chose the high-cost toy in both conditions (two-tailed binomial tests; exploration: $P = 0.004$; 95% CI = 0.63–0.95; instruction: $P = 0.043$; 95% CI = 0.53–0.89; two-tailed Fisher's exact test across conditions, $P = 0.742$; OR = 0.65).

The ability to draw inferences about the expected costs of discovery is critical for continued transmission of knowledge that is difficult to acquire without being taught; without this ability, a learner who acquires high-cost information without personally incurring the costs (for example, because they learned it through direct instruction) might fail to prioritize teaching this information to others. Is children's ability to consider others' costs limited to what they have directly experienced in the past, or can they infer these costs without experience? Given recent work on 3- to 5-year-olds' abilities to estimate the relative difficulty of novel tasks[40], we predicted that children would be able to consider the expected discovery costs for a learner even in the absence of experiencing these costs themselves.

In experiment 3, children ($n = 25$ per condition; $M \pm$ s.d. = 6.20 ± 0.79 years of age; 60% female) made a teaching decision between two toys that differed only in costs (low-reward/high-cost toy versus low-reward/low-cost toy). In the exploration condition, children first explored the toys to learn how they worked (replicating experiment 1's different costs condition). In the instruction condition, children did not interact with the toys but watched the experimenter demonstrate how to activate them; thus, children learned from instruction without the cost of exploration. If children's decisions of what it is best to teach in our experiments critically depend on self-experience, they should choose the high-cost toy in the exploration condition but not in the instruction condition. However, if children are able to infer the discovery costs for a learner unfamiliar with the toys, they should select the high-cost toy in both conditions.

Children chose the high-cost toy in the exploration condition, replicating the results from experiment 1 (80% (20/25); two-tailed binomial test, $P = 0.004$; 95% CI = 0.63–0.95). Critically, children also showed this preference in the instruction condition (72% (18/25); two-tailed binomial test, $P = 0.043$; 95% CI = 0.53–0.89). We found no credible evidence of a difference across conditions (two-tailed Fisher's exact test, $P = 0.742$; OR = 0.65; see Fig. 4), nor between either condition and experiment 1's different costs condition (two-tailed Fisher's exact tests; exploration: $P = 1.00$; OR = 0.77; instruction: $P = 0.496$; OR = 0.50). These results suggest that children can consider the costs of others' learning and make decisions that maximize others' utilities, even when they learned from direct instruction and never personally experienced the cost of exploration. This experiment also provides further support for the idea that when children selected the toys higher in discovery costs as the best ones to teach, they did so because they anticipated that the learner would find them more difficult to discover on their own, and not simply because children received help with these toys via the experimenter's prompts (see Methods).

Discussion

Humans are incredibly sophisticated social learners. However, part of what allows us to learn so effectively from others is that we are also sophisticated social teachers. Without helpful teachers who prioritize the communication of knowledge that is important yet costly to acquire, human social learning would not be nearly as successful. The goal of the current work was to understand the general principles and social and cognitive capacities that underlie our commonsense intuitions about what it is best to teach.

Complementing previous work that focused on the actual process of teaching a given concept, we presented a computational-level account[41] that characterizes how human teachers decide which concepts or what knowledge to teach, and empirically tested its predictions with young children. Our account is grounded in the theoretical proposal that humans expect other agents to act in ways that maximize their utilities[8,13,14], as well as recent empirical work on early-developing abilities to reason about the utilities of others' goal-directed actions[31,32,34]. The full model formalizes our main hypothesis that children can consider the costs and rewards of learning from instruction and learning from exploration to decide what it is best to teach. Children's decisions in experiment 1 were most consistent with the full model, compared with simpler alternatives, and were graded with respect to the learner's utility. These results suggest that children have an abstract understanding of how teaching and exploration differentially influence the utilities of learning and can integrate these variables into a single teaching decision.

Additionally, children considered the learner's costs of exploration selectively when deciding what to teach, but not when deciding what to play with (experiment 2), even when they had not personally incurred these costs (experiment 3). Since even adults often fall victim to a 'curse of knowledge', both in everyday reasoning[42,43] and teaching[44], it is impressive that children chose to teach what would be more costly to learn from a naïve learner's perspective even though it was no longer costly for them (that is, both toys were easy to activate once learned). Given that even infants expect other agents to avoid costly options and pursue goals that are lower in costs[31,32,45], one might have expected children to have chosen the toy that was less costly to discover as the best one to teach. Yet, children made the opposite choice, suggesting that they understood the goal of teaching (that is, maximizing the learner's utilities) and could distinguish the cost of discovering how the toy works from the cost of activating it.

Our experiments were designed to maximize the chances of identifying such competence in young children. First, we situated the task in a dyadic, pedagogical context with which US children are highly familiar, where the learner's goals (the learner wants to learn about both toys) and the teacher's knowledge and motivation (the experimenter knows about both toys and wants to help the learner) are made clear. Second, by asking children to choose which toy the experimenter should teach (rather than asking children to directly teach the learner), we were able to isolate children's ability to make a prosocial decision for the learner without the influence of their own costs and rewards as teachers. Third, the simple forced-choice task of teaching one of two toys (and leaving the other to be explored) was a distillation of the real-world constraints that not everything can be taught and that selecting to teach certain concepts may come at the expense of teaching others. Note that teaching either toy was helpful; by asking children to choose just one, we could identify which teaching plan they considered more helpful. Beyond making proactive decisions to protect others from failing[46], children also anticipate the expected costs of discovery, and prioritize what to teach in ways that help others to learn more effectively and efficiently.

The current work is an example of how developmental experiments and theory-driven computational modelling can mutually inform one another. Our full model assumes that children are able to reason about rewards and costs simultaneously, and trade-off between them. While integration of rewards and costs is a basic assumption in utility-based reasoning, one might question whether such integration is necessary to explain children's decisions. For example, children in our study might have relied on a set of rules or heuristics that approximate reasoning about utilities rather than simultaneously considering both rewards and costs (see Supplementary Note 1). Given that children of this age are unlikely to have acquired explicit rules about what to prioritize when teaching others, even the use of such heuristics that systematically consider either costs or rewards (and especially the tendency to preferentially teach high-cost options) would still be an impressive feat. The conflict conditions in experiment 1, however, provide some evidence against such heuristics-based accounts. In these conditions, the difference in expected rewards is held constant while the difference in expected discovery costs between the two toys increases. If children were simply considering either the difference in rewards or the difference in costs, their decisions would not differ across these conditions. Children's tendency to select the higher-cost toy gradually increased with the magnitude of the difference in discovery costs, suggesting a trade-off between the toys' discovery costs and rewards. Collectively, our work shows how smart, considerate teaching decisions can naturally emerge from early-developing social and cognitive capacities (that is, an understanding of what others want or know[47], as well as what is costly or rewarding for others[8]) without the need to posit explicit norms about what constitutes helpful teaching. A similar approach has been used to derive Gricean maxims from utility-theoretic models of pragmatic communication[48].

Further work is still needed to provide a more rigorous test of the hypothesis that children's teaching decisions vary flexibly as a function of the relative differences in the toys' costs and rewards. This endeavour would benefit from finding more precise ways of estimating the shape of children's subjective cost and reward functions. We made reasonable simplifying assumptions to estimate these functions (see Supplementary Note 1), and our full model captured the overall pattern of children's choices at the group level. However, it cannot capture potential individual differences in the exact shape of these functions. Developing models that can predict children's choices with better quantitative fit, and linking group-level predictions to individuals' behaviour, remain important challenges for future work.

We aimed to formalize the basic ingredients that contribute to human teaching decisions within the broader context of utility-based reasoning. From this perspective, our full model is by no means a complete model of human teaching. Yet a key strength is that it easily lends itself to extensions that incorporate additional factors that might influence teaching decisions in more complex, real-world contexts. For example, real-world teaching often involves flexibly deciding how much time or effort one ought to invest in teaching others depending on the context. Although we deliberately designed the current task to isolate children's consideration of others' utilities from their own, previous work on children's resistance to providing redundant information suggests that consideration of one's own utilities may indeed influence children's behaviours as teachers[22,29]. Would children be less likely to teach the toys high in discovery cost if these toys were also more costly to demonstrate? Future work might explore this question by extending the model to incorporate the teacher's costs, and asking children to choose and directly teach one of two toys that vary in activation costs.

Our models also assumed that exploration is costly and that the discovery reward of a toy is the same as its activation reward. Correspondingly, the actual process of repeatedly pressing inert buttons on the (extra-)high-cost toys was not particularly exciting for the children. However, a helpful teacher might also understand that the process of exploration and discovery itself can be beneficial for learning[49] and that pursuing a challenging goal may be intrinsically rewarding[50]. Anecdotally, two children in experiment 1 chose to teach the low-cost toy and explicitly mentioned the value of letting the learner explore the more challenging toy on their own (see Supplementary Results 1). Although we can incorporate the added value of exploration and discovery into our model by sampling different rewards for activation and discovery, measuring and manipulating these variables remains a challenge for empirical work. How and when children acquire such a nuanced understanding of the costs and rewards of learning is an intriguing topic for research, with deep implications for promoting motivation in classrooms and informal learning contexts[51].

Children's success in our study raises questions about what really develops and why these children might still fare poorly as teachers in the real world. One possibility is that children's ability to consider more abstract costs and rewards develops over time. Our task involved concrete, perceptual rewards and physical costs, which lend themselves well to formalization, and are easy for even young children to understand. However, adult teachers also understand the long-term expected utility of generalizable information and skills (for example, natural pedagogy[9]) and discovery costs that cannot be gleaned from visual complexity, such as causal complexity[52] or the physical and mental demands of novel tasks[40]. Children's ability to consider others' utility functions that differ from their own may also continue to develop. In our task, because children were presented with a generic learner whose specific preferences and competencies were unknown, children might have reasonably assumed that what was rewarding or costly for them in the past would also be rewarding or costly for the learner. Making teaching decisions for others who have quite different utility functions than they do (for example, different preferences or competencies[31,36,53]) may be more challenging and improve in tandem with children's theory of mind and executive control[25,28,54].

The ability to prioritize teaching knowledge that yields a higher utility for the learner, when implemented by multiple individuals over generations, may be a powerful force that shapes our collective knowledge. It would allow human societies to incorporate new discoveries, preserve timeless wisdom, and filter out obsolete skills and knowledge, thereby curating a collection of cultural knowledge deemed valuable by its members. Our computational account provides a possible explanation for how such curation might occur. Furthermore, over time, the tendency to teach high-utility knowledge may give rise to different cultural norms about what it is best to teach. The cognitive mechanisms that subserve the formation of such norms might also allow them to diversify and evolve over time; as the intellectual and technological repertoire of a society grows and changes, its members can reassess the utility of existing knowledge and flexibly update the content of these norms to reflect the changes in the cost–reward structure of their environment.

Relatedly, human societies show significant variability in the degree to which they rely on direct pedagogy as the primary means of knowledge transfer (as opposed to observation or exploration)[55,56]. Our work suggests that such variability might reflect the relative utility of direct instruction, compared to other ways of learning, given the nature of the knowledge and skills to be acquired; teaching may not be necessary if learning can occur through other means that do not require a 'teacher' who is willing to increase others' utilities at the expense of their own[6,57]. Compared with more abstract knowledge (such as algebra), practical skills (such as cooking) or social conventions (such as who sits where during mealtime) may be easier to acquire through observation and imitation. In such cases, knowledgeable individuals might slow their actions or tolerate onlookers as a way to facilitate others' learning. But even the tendency to slow down or allow onlookers may be higher for goals that are more costly to achieve for the observer[5,55]. It is thus possible that selective communication via an intuitive cost–benefit analysis is a culturally universal yet distinctively human ability that extends beyond pedagogical contexts. However, our work leaves open the question of whether prioritization of high-utility information is observed in non-pedagogical social learning contexts or in cultures where direct instruction is rare[16,55].

Theories of human learning have often contrasted the value of active exploration and discovery[1–4,49,58] and the importance of social learning[9,59–63]. These views have fuelled theoretical debates about the relative efficacy of instruction versus active learning[11,64], and have raised practical questions in real-world classrooms about whether to teach students or encourage them to figure things out on their own. The current work provides a different perspective: rather than replacing or hindering exploration, teachers can optimize learning by ensuring that learners acquire knowledge that is important yet difficult to learn from exploration, while guiding learners to explore things that can or should be acquired on their own. Overall, we provide the theoretical, formal and empirical groundwork for studying the development of human intuitions about what to teach. Even in their early years as learners, human children understand how to selectively share what they know to help others learn.

Methods

Implementation of the main (full) model. Given two devices (toy X and toy Y), our main (full) model compares the learner's expected utilities under the following two teaching plans: teach_X (and let the learner explore Y) and teach_Y (and let the learner explore X). The learner's expected utility under a given teaching plan is defined as the linear sum of the following components: after being taught a toy, the learner incurs a small activation cost (C_A) to obtain the activation reward (R_A) of this chosen toy; in exploring the other toy (the unchosen toy), the learner incurs a discovery cost (C_D), and obtains the discovery reward (R_D) upon successful discovery. Making a decision about what to teach (and consequently what the learner will explore) is formalized as choosing the plan that maximizes the learner's expected utilities (see Fig. 1a).

The cost function is assumed to be the physical cost of pressing buttons, and is based on the toys' physical properties. The activation cost (C_A) is the number of button presses required to activate the toy: 1 for low-cost toys and the medium-cost toy (all of which have a single-button mechanism), and 2 for the high-cost toys and the extra-high-cost toy (all of which require two specific buttons to be pressed together). These values are both small, reflecting the fact that once an agent knows how the two toys work, both are easy to activate. The discovery cost (C_D) of a toy is the expected number of button presses required for a naïve learner to explore and discover how to activate it, assuming that the learner would first try simpler actions (pressing one button at a time) and then try more complex actions (pressing combinations of two buttons) once the simpler hypothesis space has been depleted (low-cost toys: $C_D = 1$; medium-cost toy: $C_D = E(C_D) = 4$; high-cost toys: $C_D = E(C_D) = 29$; extra-high-cost toy: $C_D = E(C_D) = 79$; see Supplementary Note 1 for more details).

The activation reward (R_A) and discovery reward (R_D) of a toy are assumed to be equal, but variable across participants, to allow for the possibility that different participants assign different subjective rewards to each effect. The toys' rewards were sampled from beta distributions (ranging from 1–87, to allow the maximum reward of a toy to be as high as three times the expected discovery cost of the high-cost toys, which is also higher than the expected discovery cost of the extra-high-cost toy). We fit these distributions such that there was an 80% probability that a sample from the lights (high-reward) distribution would be higher in value than a sample from the music (low-reward) distribution, matching empirical data on children's relative preference for these causal effects (80% of children (20/25) preferred the lights; see Supplementary Methods 2, Supplementary Results 2 and Supplementary Note 1 for more details).

It is possible that children's reasoning about rewards and costs is more complex than how we have implemented these functions. For example, cost may not increase linearly with each button press, and reward might be the highest on first observation and then degrade with each subsequent observation due to decreasing novelty. In the context of our experiment, model predictions depend on the relative reward and cost rather than their absolute values; identifying the precise shape of children's cost and reward functions is beyond the scope of the current work. Thus, we adopt the simplest cost function that is consistent with how children actually explored the medium-cost, high-cost and extra-high-cost toys (that is, most children first pressed each button individually before attempting combinations of buttons), and the simplest reward function where rewards remain constant across time. Some children might also assume an extra value to the process of exploration and discovery (for example, $R_A < R_D$), or consider the utility of the teacher. See Discussion for how the current model could be extended to incorporate these possibilities.

Additionally, we considered the probability (P_E) that the learner would explore the unchosen toy. This exploration certainty parameter (higher means more certain) reflects the idea that teaching ensures immediate learning and guarantees a reward, while there is more uncertainty about the learner's future exploratory behaviours and whether or not they will successfully activate the unchosen toy. For example, they may never explore this toy, or even if they do, they might not figure it out, which is of particular concern for the high-cost and extra-high-cost toys.

Our full model considers all of these components to compute the learner's expected utilities under each teaching plan (see Fig. 1), and chooses the plan with the higher overall expected utility for the learner. To generate the model predictions reported in the main text (Fig. 2c,d), we used $P_E = 0.5$, as well as a small degree of noise ($\alpha = 0.1$) in children's choices. However, model fits are robust to: (1) different distributions of the reward space; (2) different ranges of reward values; and (3) different values of P_E and α (see Supplementary Note 1 for details). Model code and full predictions can be found at https://osf.io/wunbq/.

Participants. We recruited 250 5-, 6- and 7-year-olds ($M \pm$ s.d. = 6.29 ± 0.90 years of age; 55% female; $n = 25$ per condition for all experiments) for experiments 1–3 from a local children's museum in the San Francisco Bay Area. Experiment 1 involved 150 participants aged 5, 6 and 7 years ($M \pm$ s.d. = 6.35 ± 0.97 years; 53% female). An additional nine children were excluded from the analysis due to difficulty understanding English (four), a missing date of birth (two), an inability to learn the mechanism of the extra-high-cost toy (two) or parental interference (one). Experiment 2 involved 50 participants aged 5, 6 and 7 years ($M \pm$ s.d. = 6.19 ± 0.77 years; 54% female). An additional three children were excluded from the analysis due to a missing date of birth (one), parental interference (one) or experimenter error (one). Experiment 3 involved 50 participants aged 5, 6 and 7 years ($M \pm$ s.d. = 6.20 ± 0.79 years; 60% female). An additional five children were excluded from the analysis due to experimenter error (two) or not completing the procedure (three).

No statistical methods were used to predetermine sample sizes, but our sample sizes were pre-set and similar to those reported in previous publications on children's teaching[22,29,35]. All procedures were approved by Stanford University's Institutional Review Board for Human Subjects Research. Informed consent was obtained from the legal guardians of all child participants.

For experiment 1, participants were randomly assigned to the different rewards and high-cost conflict conditions, followed by the rewards and costs and different costs conditions. The medium-cost and extra-high-cost conflict conditions were added during the review process (participants were randomly assigned to either condition). See https://osf.io/5dmga/ for preregistration, and Supplementary Results 1 for further statistical analyses of the behavioural results. In experiment 2, we first collected data for the play condition, followed by the teach condition (a replication of the high-cost conflict condition in experiment 1). In experiment 3, we first collected data for the instruction condition, followed by the exploration condition (a replication of the different costs condition in experiment 1). Data collection and analysis were not performed blind to the conditions of the experiments. Children's responses were recorded offline from the video by S.B. and a second researcher blind to the experimental hypotheses. See Supplementary Methods 1 for details.

Materials. We constructed seven rectangular toys (25 cm × 15 cm × 15 cm) from foam board, felt, push-buttons and electrical circuits. Low-cost toys had one (large) button that activated the toys' effects. The medium-cost toy had one large and six identical small buttons; pressing one of the small buttons activated the toy. The high-cost toys also had one large and six identical small buttons, but to activate these toys, the large button and one particular small button had to be pressed at the same time. The extra-high-cost toy had one large button and 11 identical small buttons, and just like the high-cost toys, to activate this toy, the large button and one particular small button had to be pressed at the same time. When activated, low-reward toys played music, whereas high-reward toys had a plastic orb that lit up different colours and spun around (see Fig. 2b for a schematic of the toys used in each condition of experiment 1, and Fig. 4 for a schematic of the toys used in experiments 2 and 3).

Procedure. Children were tested individually in a quiet room separate from the main exhibits of the museum. Experiment 1 began with the discovery phase. The experimenter produced two toys (determined by condition), explained that she did not know how the toys worked, and asked children for help figuring them out. Children explored both toys until they activated them. Children quickly figured out low-cost toys, immediately pressing the single button (experiment 1: $M_{\text{discovery time}} \pm$ s.d. = 0.28 ± 1.02 s). Figuring out the medium-cost toy took longer (experiment 1: $M_{\text{discovery time}} \pm$ s.d. = 9.33 ± 6.37 s). Figuring out the high-cost toys and extra-high-cost toy took much longer (experiment 1: high-cost toys: $M_{\text{discovery time}} \pm$ s.d. = 81.22 ± 49.42 s; extra-high-cost toy: $M_{\text{discovery time}} \pm$ s.d. = 94.52 ± 61.79 s) and involved pressing many buttons. If children stopped exploring the high-cost or extra-high-cost toys before successfully activating them, the experimenter prompted them to continue, and provided suggestions of what they could try next (for example, "I wonder what would happen if you pressed two buttons at the same time."). These prompts were delivered as though they were spontaneous, and with uncertainty, to keep up the charade that the experimenter did not know how these toys worked. Critically, the experimenter never explicitly told or showed children how to activate the toys. Once children figured out both toys, they were asked to demonstrate each toy twice (with the order counterbalanced) to ensure that they knew how to activate the toys and had the chance to experience the ease of activating them after discovery (i.e., the minimal activation cost).

Next came the choice phase. The experimenter explained that her friend would play with the toys later all by herself but knew nothing about them, so beforehand, the experimenter would help by teaching her how one of the toys worked. The experimenter then asked: "Which toy should I teach her?" Children selected a toy and explained their choice. Children also answered two follow-up questions: (1) which toy was harder to figure out?; and (2) which effect was cooler?

In experiment 2, the teach condition was identical to experiment 1's high-cost conflict condition. The only difference was in the choice phase of the play condition; the experimenter said she needed to work on something and asked the children to choose one toy to play with while she was gone. In experiment 3, the exploration condition was identical to experiment 1's different costs condition. The only difference was in the instruction condition; a demonstration phase replaced the discovery phase. Children never interacted with the toys; instead, the experimenter taught children how the toys worked: she pressed the appropriate buttons to generate the effect, and added, "this is the only way to make the toy go" (see Supplementary Methods 1).

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

The data and analysis scripts that support the findings of this study are available at https://osf.io/wunbq/.

Code availability

Model code and full predictions can be found at https://osf.io/wunbq/.

Received: 7 September 2018; Accepted: 29 August 2019;
Published online: 14 October 2019

References

1. Piaget, J. *The Child's Conception of the World* (Harcourt Brace Jovanovich, 1929).
2. Singer, D. G., Golinkoff, R. M. & Hirsh-Pasek, K. *Play = Learning: How Play Motivates and Enhances Children's Cognitive and Social-Emotional Growth* (Oxford Univ. Press, 2006).
3. Schulz, L. The origins of inquiry: inductive inference and exploration in early childhood. *Trends Cogn. Sci.* **16**, 382–389 (2012).
4. Stahl, A. E. & Feigenson, L. Observing the unexpected enhances infants' learning and exploration. *Science* **348**, 91–94 (2015).
5. Kline, M. A. How to learn about teaching: an evolutionary framework for the study of teaching behavior in humans and other animals. *Behav. Brain Sci.* **38**, e31 (2015).
6. Caro, T. & Hauser, M. Teaching in nonhuman animals. *Q. Rev. Biol.* **67**, 151–174 (1992).
7. Wellman, H. & Cross, D. Theory of mind and conceptual change. *Child Dev.* **72**, 702–707 (2001).
8. Jara-Ettinger, J., Gweon, H., Schulz, L. E. & Tenenbaum, J. B. The nave utility calculus: computational principles underlying commonsense psychology. *Trends Cogn. Sci.* **20**, 589–604 (2016).
9. Csibra, G. & Gergely, G. Natural pedagogy. *Trends Cogn. Sci.* **13**, 148–153 (2009).
10. Shafto, P., Goodman, N. D. & Griffiths, T. L. A rational account of pedagogical reasoning: teaching by, and learning from, examples. *Cogn. Psychol.* **71**, 55–89 (2014).
11. Bonawitz, E. et al. The double-edged sword of pedagogy: instruction limits spontaneous exploration and discovery. *Cognition* **120**, 322–330 (2011).
12. Baker, C. L., Jara-Ettinger, J., Saxe, R. & Tenenbaum, J. B. Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nat. Hum. Behav.* **1**, 0064 (2017).
13. Lucas, C. G. et al. The child as econometrician: a rational model of preference understanding in children. *PLoS One* **9**, e92160 (2014).
14. Jern, A., Lucas, C. G. & Kemp, C. People learn other people's preferences through inverse decision-making. *Cognition* **168**, 46–64 (2017).
15. Schulz, L. & Bonawitz, E. Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Dev. Psychol.* **43**, 1045–1050 (2007).
16. Shneidman, L., Gweon, H., Schulz, L. E. & Woodward, A. L. Learning from others and spontaneous exploration: a cross-cultural investigation. *Child Dev.* **87**, 723–735 (2016).
17. Gweon, H. & Schulz, L. 16-month-olds rationally infer causes of failed actions. *Science* **332**, 1524 (2011).
18. Was, A. M. & Warneken, F. Proactive help-seeking: preschoolers know when they need help, but do not always ask for it. *Cogn. Dev.* **43**, 91–105 (2017).
19. Sobel, D. M. & Kushnir, T. Knowledge matters: how children evaluate the reliability of testimony as a process of rational inference. *Psychol. Rev.* **120**, 779–797 (2013).
20. Koenig, M. A. & Harris, P. L. Preschoolers mistrust ignorant and inaccurate speakers. *Child Dev.* **76**, 1261–1277 (2005).
21. Gweon, H. & Asaba, M. Order matters: children's evaluation of underinformative teachers depends on context. *Child Dev.* **89**, e278–e292 (2018).
22. Gweon, H., Shafto, P. & Schulz, L. Development of children's sensitivity to over-informativeness in learning and teaching. *Dev. Psychol.* **54**, 2113–2125 (2018).
23. Sobel, D. M. & Letourneau, S. M. Children's developing knowledge of and reflection about teaching. *J. Exp. Child Psychol.* **143**, 111–122 (2016).
24. Shatz, M. & Gelman, R. The development of communication skills: modifications in the speech of young children as a function of listener. *Monogr. Soc. Res. Child Dev.* **38**, 1–38 (1973).
25. Ziv, M., Solomon, A., Strauss, S. & Frye, D. Relations between the development of teaching and theory of mind in early childhood. *J. Cogn. Dev.* **17**, 264–284 (2016).
26. Baer, C. & Friedman, O. Fitting the message to the listener: children selectively mention general and specific facts. *Child Dev.* **89**, 461–475 (2018).

27. Rhodes, M., Gelman, S. & Brickman, D. Children's attention to sample composition in learning, teaching and discovery. *Dev. Sci.* **13**, 421–429 (2010).
28. Bass, I. et al. Children's developing theory of mind and pedagogical evidence selection. *Dev. Psychol.* **55**, 286–302 (2019).
29. Gweon, H. & Schulz, L. From exploration to instruction: children learn from exploration and tailor their demonstrations to observers' goals and competence. *Child Dev.* **90**, e148–e164 (2019).
30. Gergely, G., Nádasdy, Z., Csibra, G. & Bró, S. Taking the intentional stance at 12 months of age. *Cognition* **56**, 165–193 (1995).
31. Jara-Ettinger, J., Gweon, H., Tenenbaum, J. B. & Schulz, L. E. Children's understanding of the costs and rewards underlying rational action. *Cognition* **140**, 14–23 (2015).
32. Liu, S., Ullman, T. D., Tenenbaum, J. B. & Spelke, E. S. Ten-month-old infants infer the value of goals from the costs of actions. *Science* **358**, 1038–1041 (2017).
33. Pesowski, M. L., Denison, S. & Friedman, O. Young children infer preferences from a single action, but not if it is constrained. *Cognition* **155**, 168–175 (2016).
34. Liu, S., Gonzalez, G. & Warneken, F. Worth the wait: children trade off delay and reward in self- and other-benefiting decisions. *Dev. Sci.* **6**, e12702–e12708 (2018).
35. Ronfard, S., Was, A. M. & Harris, P. L. Children teach methods they could not discover for themselves. *J. Exp. Child Psychol.* **142**, 107–117 (2016).
36. Repacholi, B. & Gopnik, A. Early reasoning about desires: evidence from 14- and 18-month-olds. *Dev. Psychol.* **33**, 12–20 (1997).
37. Téglás, E. et al. Pure reasoning in 12-month-old infants as probabilistic inference. *Science* **332**, 1054–1059 (2011).
38. Lagattuta, K. H. Linking past, present, and future: children's ability to connect mental states and emotions across time. *Child Dev. Perspect.* **8**, 90–95 (2014).
39. Atance, C. M. & O'Neill, D. K. Episodic future thinking. *Trends Cogn. Sci.* **5**, 533–539 (2001).
40. Gweon, H., Asaba, M. & Bennett-Pierre, G. Reverse-engineering the process: adults and preschoolers' ability to infer the difficulty of novel tasks. In *Proc. 39th Annual Conference of the Cognitive Science Society* 458–463 (Cognitive Science Society, 2017).
41. Marr, D. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (MIT Press, 1982).
42. Camerer, C., Loewenstein, G. & Weber, M. The curse of knowledge in economic settings: an experimental analysis. *J. Polit. Econ.* **97**, 1232–1254 (1989).
43. Birch, S. & Bloom, P. The curse of knowledge in reasoning about false beliefs. *Psychol. Sci.* **18**, 382–386 (2007).
44. Aboody, R., Velez-Ginorio, J., Santos, L. R. & Jara-Ettinger, J. When teaching breaks down: teachers rationally select what information to share, but misrepresent learners' hypothesis spaces. In *Proc. 40th Annual Meeting of the Cognitive Science Society* 70–75 (2018).
45. Gergely, G. & Csibra, G. Teleological reasoning in infancy: the naïve theory of rational action. *Trends Cogn. Sci.* **7**, 287–292 (2003).
46. Martin, A. & Olson, K. R. When kids know better: paternalistic helping in 3-year-old children. *Dev. Psychol.* **49**, 2071–2081 (2013).
47. Wellman, H., Cross, D. & Watson, J. Meta-analysis of theory-of-mind development: the truth about false belief. *Child Dev.* **72**, 655–684 (2001).
48. Goodman, N. D. & Frank, M. C. Pragmatic language interpretation as probabilistic inference. *Trends Cogn. Sci.* **20**, 818–829 (2016).
49. Gureckis, T. M. & Markant, D. B. Self-directed learning: a cognitive and computational perspective. *Perspect. Psychol. Sci.* **7**, 464–481 (2012).
50. Inzlicht, M., Shenhav, A. & Olivola, C. Y. The effort paradox: effort is both costly and valued. *Trends Cogn. Sci.* **22**, 337–349 (2018).
51. Dweck, C. S. *Mindset: The New Psychology of Success* (Random House, 2006).
52. Ahl, R. E. & Keil, F. C. Diverse effects, complex causes: children use information about machines' functional diversity to infer internal complexity. *Child Dev.* **88**, 828–845 (2017).
53. Magid, R. W., Depascale, M. & Schulz, L. E. Four-and 5-year-olds infer differences in relative ability and appropriately allocate roles to achieve cooperative, competitive, and prosocial goals. *Open Mind* **2**, 72–85 (2018).
54. Benson, J. E., Sabbagh, M. A., Carlson, S. M. & Zelazo, P. D. Individual differences in executive functioning predict preschoolers' improvement from theory-of-mind training. *Behav. Brain Sci.* **49**, 1615–1627 (2013).
55. Boyette, A. H. & Hewlett, B. S. Teaching in hunter-gatherers. *Rev. Philos. Psychol.* **9**, 771–797 (2018).
56. Legare, C. H. Cumulative cultural learning: development and diversity. *Proc. Natl Acad. Sci. USA* **114**, 7877–7883 (2017).
57. Tomasello, M. Born (and bred) to help. in *Why We Cooperate* 1–48 (MIT Press, 2009).
58. Bruner, J. S., Jolly, A. & Sylva, K. *Play: Its Role in Development and Evolution* (Basic Books, 1976).
59. Vygotsky, L. S. *Mind in Society: The Development of Higher Psychological Processes* (Harvard Univ. Press, 1980).
60. Boyd, R., Richerson, P. J. & Henrich, J. The cultural niche: why social learning is essential for human adaptation. *Proc. Natl Acad. Sci. USA* **108**, 10918–10925 (2011).
61. Tomasello, M., Carpenter, M., Call, J., Behne, T. & Moll, H. Understanding and sharing intentions: the origins of cultural cognition. *Behav. Brain Sci.* **28**, 691–735 (2005).
62. Legare, C. H. & Nielsen, M. Imitation and innovation: the dual engines of cultural learning. *Trends Cogn. Sci.* **19**, 688–699 (2015).
63. Heyes, C. Who knows? Metacognitive social learning strategies. *Trends Cogn. Sci.* **20**, 204–213 (2016).
64. Klahr, D. & Nigam, M. The equivalence of learning paths in early science instruction: effect of direct instruction and discovery learning. *Psychol. Sci.* **15**, 661–667 (2004).

Acknowledgements

We thank C. Dweck, M. C. Frank, E. Markman, M. H. Tessler, M. Asaba, K. Weisman and N. Vélez for helpful conversations and insightful comments. We thank G. Bennett-Pierre, A. Singh, F. Kramer, A. Garron and N. Chandaria for help with data collection and coding. We are grateful to the Palo Alto Junior Museum and Zoo, the Tech Museum of Innovation in San Jose and the children and families who participated in this research. This work was funded by a John Templeton Foundation Varieties of Understanding grant (to H.G.), a James S. McDonnell Scholar Award (to H.G.) and an NSF Graduate Research Fellowship (to S.B.). In addition, this material is based upon work supported by the Center for Brains, Minds, and Machines (CBMM), funded by NSF-STC award CCF-1231216. The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Author contributions

S.B. and H.G. conceived of and designed the experiments. S.B. collected and analysed the data. J.J.-E. designed, implemented and conducted the formal model comparisons, with assistance from S.B. and H.G. S.B., H.G. and J.J.-E. interpreted the results and wrote and edited the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at https://doi.org/10.1038/s41562-019-0748-6.

Correspondence and requests for materials should be addressed to S.B. or H.G.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2019

nature research

Corresponding author(s): Sophie Bridgers and Hyowon Gweon

Last updated by author(s): Aug 21, 2019

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see Authors & Referees and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed	
☐	☒	The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☒	☐	A description of all covariates tested
☒	☐	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☐	☒	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒	Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	No software was used.
Data analysis	To visualize and analyze children’s behavior we used R version 3.5.2 and the following packages: tidyverse version 1.2.1, tidyboot 0.1.1, cowplot 0.9.4, and lubridate 1.7.4 To implement our models and generate predictions, we used Python version 2.7.10. To visualize our model predictions and conduct formal model comparison, we used R version 3.5.2 and the following packages: tidyverse version 1.2.1, boot version 1.3-20, and glue 1.3.0

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

All data and analysis scripts will be made available on the Open Science Framework upon publication at the following link: https://osf.io/wunbq/

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☐ Life sciences ☒ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	Quantitative experimental
Research sample	We recruited 250 5-, 6-, and 7-year-olds (M(SD) = 6.29(0.90)yrs; 55% female) for Experiments 1-3 from a local children's museum; we recruited a separate group of 25 5-, 6-, and 7-year-olds (M(SD) = 6.15(0.93) yrs; 60% female) for the Norming Experiment from local children's museums. A range of ethnicities proportional to the local population were included. We selected 5 to 7 years as our target age range for two reasons. First, children this age understand teaching as a process that causes knowledge change (Sobel & Letourneau, 2016), and consider others' mental states to decide what to communicate as teachers (Gweon, Pelton, Konopka, & Schulz, 2014). Second, they can infer and integrate the expected costs and rewards of others' actions in their reasoning (see Jara-Ettinger et al., 2015). Thus, we expected that children this age might be capable of simulating another person's utility to make decisions on their behalf.
Sampling strategy	We randomly sampled and assigned children to one of two conditions in Experiment 1 (Different Rewards and High Cost Conflict or Different Costs and Rewards & Costs). The Medium Cost Conflict and the Extra-High Cost Conflict conditions in Experiment 1 were run during the review process and children were randomly assigned to one of these two conditions. The sample size per condition (n = 25) was consistent with previous conditions and preregistered along with hypotheses and analyses. The preregistration for these conditions can be found at: https://osf.io/5dmga/ In Experiment 2, we first collected data for the Play condition, followed by the Teach condition (a replication of the High Cost Conflict condition in Experiment 1). In Experiment 3, we first collected data for the Instruction condition, followed by the Exploration condition (a replication of the Different Costs condition in Experiment 1). We pre-set our sample size at 25 children per condition for all experiments. This was based on closely related prior work on children's ability to teach (e.g., Ronfard & Corriveau 2016, Gweon & Schulz 2019, Gweon Shafto & Schulz, 2018); our sample size was similar to those used in these studies.
Data collection	Data was collected in a room separate from the main exhibits at the museum by either the first author or an undergraduate research assistant. Although the experimenters were not blind to the study's hypothesis, they were thoroughly trained to ensure that the experiment was run without any indication of how children should respond. Children's responses were recorded by video when parent permission was given. We had parental permission to video record all but 9 children whose responses were recorded using pen and paper by the experimenter. Sometimes members of the child participant's family were in the room while the experiment was being conducted.
Timing	Experiment 1: 4/15/2015 - 1/15/2016 and 1/12/2019 - 1/27/2019 (preregistered conditions run during the review process); Experiment 2: 5/3/2015-3/25/2017; Experiment 3: 5/8/2016-3/25/2017; Norming Experiment: 4/23/2017 - 6/29/2017. Any breaks are due to school holidays, museum holidays, and travel.
Data exclusions	Seventeen participants were excluded from analysis across our three experiments (see SI). In Experiment 1, 9 children were excluded from analysis due to difficulty understanding English (4), missing date of birth (2), an inability to learn the mechanism of the extra-high-cost toy (2), or parental interference (1). In Experiment 2, 3 children were excluded from analysis due to missing date of birth (1), parental interference (1), or experimenter error (1). In Experiment 3, 5 children were excluded from analysis due to experimenter error (2) or not completing the procedure (3). No children were excluded from the Norming Experiment. Exclusion criteria were pre-established and consistent with prior research in the lab.
Non-participation	Three participants began the experiment but did not complete it due to being too shy to respond to the experimenter's questions.
Randomization	Allocation to experimental conditions was random. We recruited until we had 25 participants per condition per experiment that we could include in data analyses.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology
☒	☐ Animals and other organisms
☐	☒ Human research participants
☒	☐ Clinical data

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Human research participants

Policy information about studies involving human research participants

Population characteristics	See above.
Recruitment	The main experimenter approached families in the museum and asked if they had children between the ages of 5 and 7 years and if they would be interested in participating in a study exploring children's teaching. The museum has no admission fee, and attracts visitors from a wide range of SES, ethnic, and cultural backgrounds. Although it is possible that parents who willingly agreed to have their child participate in research may be somewhat more interested in science, we have no reason to expect that such a tendency would bias the results of the current study.
Ethics oversight	Stanford University's Human Subjects Research and IRB

Note that full information on the approval of the study protocol must also be provided in the manuscript.