** Accepted for publication in Nature Human Behaviour(June 2019).

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

Sophie Bridgers^{1*}, Julian Jara-Ettinger², and Hyowon Gweon^{1*}

¹Department of Psychology, Stanford University,
450 Serra Mall, Bldg. 420, Stanford, CA, USA

²Department of Psychology, Yale University, New Haven, CT, USA

Corresponding Authors:

Sophie Bridgers & Hyowon Gweon

Department of Psychology, Stanford University

450 Serra Mall, Bldg. 420, Stanford, CA, USA 94305

E-mail: sbridge@stanford.edu, gweon@stanford.edu

Abstract

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 prioritise what to teach. We present a computational model that decides what to teach by maximising the learner's expected utility of learning from instruction and from exploration, and show that children (age 5-7) make decisions that are consistent with the model's predictions (i.e., minimising the learner's costs and maximising 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^{1,2,3,4}. 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 costly⁵. 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 resources⁶, 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 upon 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 prioritise 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 Mind⁷) but also what is rewarding for other people to acquire and what is costly to achieve on their own (Naïve Utility Calculus⁸). We formalise these intuitions as decisions that maximise learners’ expected utilities and demonstrate that even young children can rationally decide what to teach based on an intuitive cost-benefit analysis of learning from instruction vs. 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 (e.g., 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 knowledge⁹, while others emphasise 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 prior work: How 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 Calculus⁸: People have an intuitive theory of agents as “utility-maximisers”, expecting agents to pursue rewarding goals while minimising the associated costs^{12,13,14}. Critically, however, agents can only maximise 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 rewards. The key insight behind our model is that a “teacher”—a

55 knowledgeable agent who already understands the cost-reward structure of the world—can
56 help the learner reduce or eliminate the costs and uncertainties of exploration. By
57 selectively teaching what is rewarding and/or costly to learn, the teacher can ensure that
58 the learner reaps the rewards of learning while also bypassing potentially extreme costs.
59 Importantly, such prioritisation of high-utility information still allows the learner to
60 explore and make additional discoveries even if the teacher cannot teach everything.

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

72 First, children, as learners themselves, selectively engage in active exploration
73 especially when additional discovery is likely^{15,11,16,4}, and rationally decide whether to
74 explore or seek help^{17,18}. This work suggests that an intuitive understanding of the relative
75 benefits of social learning and self-guided exploration may be present early in life. Second,
76 by the beginning of formal schooling, children already have an abstract understanding of
77 what constitutes informative teaching (i.e., the information communicated should be
78 accurate, sufficient, yet not superfluous), and readily evaluate others’ teaching
79 accordingly^{19,20,21,22,23}. As teachers themselves, children can also tailor the content of their
80 communication based on the learner’s goal or knowledge state^{24,25,26}, readily selecting
81 appropriate evidence to communicate a concept²⁷, correct others’ false beliefs²⁸, or

disambiguate a causal system^{29,22}. In particular, when evidence is physically costly to generate (e.g., demonstration requires multiple actions or walking over a distance), 5- to 6-year-olds resist providing unnecessary information^{29,22}, 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 minimise the overall cost of generating information could emerge solely from the teacher’s motivation to minimise their own costs rather than a regard for the learner. Given that even infants and young children expect others to act in ways that maximise their expected utilities^{30,31,32,33,34}, an important open question is whether children can actively choose to maximise others’ utilities. Consistent with this possibility, one study reports that children prioritise 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 discover³⁵. However, children in this study could have privileged this information simply 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 the role of reward (i.e., pleasure or benefit of knowing something) and the role of cost (i.e., 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 prior utility-based models of social reasoning have focused on formalising 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 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 to the model predictions. We chose 5 to 7 years as our target age range, because children this age consider others’ mental states to decide what to communicate as teachers^{29,22,28} and can infer and integrate expected costs and rewards when reasoning about others’ goal-directed actions⁸.

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, she will also have to figure out how to activate the other device through her own exploration. Formally, this decision can be described as a choice between two possible teaching plans (i.e., teach device X and let the learner explore device Y or vice versa) which 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 her expected rewards and costs of learning one device through instruction and learning the other through exploration (Figure 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 maximise 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). The medium-cost and high-cost toys had one big button and six small buttons; medium-cost toys required pressing one particular small button to activate, and thus were relatively less difficult to figure out than the high-cost toys that required pressing the big button and a

particular small button at the same time. Activating the extra-high-cost toy also required simultaneously pressing 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 formalise 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 (Chosen Toy's Activation Cost and Activation Reward), as well as the costs and rewards of the toy the learner would explore on her own (Unchosen Toy's Discovery Cost and Discovery Reward), and selects the teaching plan with the higher net utility as best to execute.

Four alternative models consider subsets of these variables and represent possible developmental changes in children's reasoning. While prior work has demonstrated an early-developing understanding of others' costs and rewards^{36,31,32}, as well as the ability to reason about uncertain future events^{37,38,39}, 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 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 (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 (Rewards-only model). It is also possible that children do consider both costs and rewards but fail to represent the utilities of both instruction and exploration; they may only concern themselves with the immediate consequence of what would be taught (Instruction-only model), or just focus on the trial-and-error the learner will experience trying to activate the Unchosen toy (Exploration-only model). Figure 1b summarises the different considerations of the main and alternative models (see Methods and Supplementary Note 1 for more details).

In Experiment 1, we test the model predictions across six between-subjects conditions that present children with different pairs of toys (Figure 2b). The choice was over-determined in the first condition: (1) Rewards & Costs: One toy is higher in both cost and reward than the other toy (high-reward/high-cost toy vs. 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 are matched in reward but vary in cost (low-reward/high-cost toy vs. low-reward/low-cost toy), (3) Different Rewards: The toys are matched in costs but vary in reward (high-reward/low-cost toy vs. 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 force children to trade-off costs and rewards because one toy is higher in rewards while the other is 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 children will prefer the high-cost/high-reward toy in the Rewards & Costs condition, because this choice maximises the learner’s utility by both increasing her rewards (ensures she will enjoy the high-reward effect) and decreasing her costs (she 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. Note however that the quantitative predictions are

not symmetric across these two conditions; reflecting children’s intuitions about which effect is “cooler” (see Supplementary Methods 2 and Results 2), the model assumes that the relative difference in the toys’ rewards (reward from seeing the lights vs. 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. 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 ($n = 25/\text{condition}$; $M(SD) = 6.35(0.97)\text{yrs}$; 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 & 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 CI’s reported are 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, it was 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 Figure 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 $> 1e5$), providing compelling support for the Full model over the four alternatives (see Figure 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 and rationally decided what to teach. Children did not just consider the expected rewards or just the expected costs for the learner; nor did they reason about just the utility of instruction or the utility of exploration. Rather, they considered the relative expected costs and rewards for a learner under different teaching plans (i.e., learning one toy from instruction and the other from exploration or vice versa) and chose the plan that maximised the learner’s expected utility.

Experiment 2

Our experimental task was designed to minimise the influence of children’s own utilities on their decisions; children did not demonstrate the toys themselves to the learner and 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 always lower for the children than for the learner (i.e., Activation Cost $\ll$ Discovery Cost, especially for the high- and extra-high-cost toys). Thus, the fact that children prioritised 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 prior success in figuring out how it works made it more attractive than the low-cost toy. Alternatively,

the existence of multiple buttons (even though 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 is best to teach are based on a genuine consideration of the learner’s expected costs of exploration, then they would prioritise high-cost toys when they are choosing a toy for the learner but not when choosing a toy for themselves. We test this prediction in Experiment 2.

Children ($n = 25/condition$; $M(SD) = 6.19(0.77)$ yrs; 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 vs. low-reward/high-cost toy). The Teach condition was a replication of Experiment 1’s High Cost Conflict condition; 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 than in the Play condition (72% (18/25) vs. 20% (5/25), two-tailed Fisher’s exact test, $p < 0.001$, Odds ratio (OR) = 9.71; see Figure 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 likely reflect 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 prioritised 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. It remains unclear, however, 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 prior experience with these costs.

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 (e.g., because they learned it through direct instruction) might fail to prioritise 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 preschoolers’ ability to estimate the relative difficulty of novel tasks⁴⁰, 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 ($n = 25/condition$; $M(SD) = 6.20(0.79)yr$ s; 60% female) made a teaching decision between two toys that differed only in costs (low-reward/high-cost toy vs. 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 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, then they should select the high-cost toy in both conditions.

Children chose the high-cost toy in the Exploration condition, replicating 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 Figure 4) nor between either condition and Experiment 1’s Different Costs condition (two-tailed Fisher’s exact tests: $p = 1.00$, $OR = 0.77$, and $p = 0.496$, $OR = 0.50$, for Exploration and Instruction, respectively). These results suggest that children can consider the costs of others’ learning and make decisions that maximise 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 best to teach, they did so because they anticipated that the learner would find them more difficult to discover on her own and not simply because they received help on these toys via the experimenter’s prompts (see Methods).

Discussion

Humans are incredibly sophisticated social learners. Part of what allows us to learn so effectively from others, however, is that we are also sophisticated social teachers: Without helpful teachers who prioritise teaching 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-cognitive capacities that underlie our commonsense intuitions about what is best to teach.

Complementing prior work that focused on the actual process of teaching a given concept, we presented a computational-level account⁴¹ that characterises 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 maximise their utilities^{8,13,14} and recent empirical work on early-developing abilities to reason about the utilities of others’ goal-directed actions^{31,32,34}. The Full model formalises our main hypothesis that children can consider the costs and rewards of learning from instruction and from exploration to decide what is best to teach. Children’s decisions in Experiment 1 were most consistent with the Full model compared to 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 teaching⁴⁴, 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 (i.e., 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^{45,32,31}, one might have expected children to have chosen the toy that was less costly to discover as best to teach. Yet children made the opposite choice, suggesting that they understood the goal of teaching (i.e., maximising 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 maximise the chances of identifying such competence in young children. First, we situated the task in a dyadic, pedagogical context with which U.S. children are highly familiar, where the learner’s goals (i.e., the learner wants to learn about both toys) and the teacher’s knowledge and motivation (i.e., 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 failing⁴⁶, children also anticipate the expected costs of discovery and prioritise what to teach in ways that help others 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 this age are unlikely to have acquired explicit rules about what to prioritise when teaching others, even the use of such heuristics that systematically consider both costs and rewards (and especially the tendency to preferentially teach high-cost options) would still be an

impressive feat. Yet the Conflict conditions in Experiment 1 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 demonstrates how smart, considerate teaching decisions can naturally emerge from early-developing social cognitive capacities (i.e., an understanding of what others want or know⁴⁷, as well as what is costly or rewarding for others⁸) 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 communication⁴⁸.

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. It cannot, however, 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 formalise 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, prior 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^{29,22}. 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. A helpful teacher, however, might also understand that the process of exploration and discovery itself can be beneficial for learning⁴⁹ and that pursuing a challenging goal may be intrinsically rewarding⁵⁰. 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 her 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 contexts⁵¹.

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

formalisation and are easy for even young children to understand. Adult teachers, however, also understand the long-term expected utility of generalisable information and skills (e.g., natural pedagogy⁹) and discovery costs that cannot be gleaned from visual complexity, such as causal complexity⁵² or the physical and mental demands of novel tasks⁴⁰. 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 (e.g., different preferences or competencies^{36,31,53}) may be more challenging and improve in tandem with children's Theory of Mind and executive control^{28,54,25}.

The ability to prioritise 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 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 may reflect the relative utilities of direct instruction given the nature of knowledge and skills to be

acquired; teaching may not be necessary if learning can easily occur through other means that do not require a “teacher” who is willing to increase others’ utilities at the expense of their own^{57,6}. Compared to highly abstract knowledge (e.g., algebra), practical skills (e.g., weaving) or social conventions (e.g., religious rituals) may be easier to acquire through observation and imitation. In such cases, knowledgeable individuals might slow their actions or simply tolerate onlookers as a way to facilitate others’ learning. Even the tendency to slow down or allow onlookers may be higher for goals that are more costly to achieve for the observer. It thus is 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 prioritisation of high-utility information is observed in non-pedagogical social learning contexts or in cultures where direct instruction is rare^{55,16}.

Theories of human learning have often contrasted the value of active exploration and discovery^{1,58,2,3,4,49} and the importance of social learning^{59,60,61,9,62,63}. These views have fueled theoretical debates about the relative efficacy of instruction vs. active learning^{64,11}, and 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 optimise 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 readily understand how to selectively share what they know to help others learn.

Methods

Implementation of 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 formalised as choosing the plan that maximises the learner’s expected utilities (see Figure 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 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 to 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 a 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 and Results 2, as well as Supplementary Note 1 for more details).

It is possible that children’s reasoning about rewards and costs are 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 upon 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 (i.e., 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 (e.g., $R_A < R_D$), or consider the utility of the teacher. See the Discussion for how the current model could be extended to incorporate these possibilities.

Additionally, we consider a probability (p_E) that the learner will 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 she will successfully activate the Unchosen Toy. For example, she may never explore this toy,

or even if she does, she 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 Figure 1), and chooses the plan with the higher overall expected utility for the learner. To generate model predictions reported in the main text (Figure 2c-d), we use $p_E = 0.5$, as well as a small degree of noise ($\alpha = 0.1$) in children’s choices. However, model fits are robust to different (1) distributions of the reward space, (2) ranges of reward values, and (3) 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 ($M(SD) = 6.29(0.90)yr$ s; 55% female; $n = 25$ per condition of all experiments) for Experiments 1 - 3 from a local children’s museum in the San Francisco Bay Area. In Experiment 1, participants were 150 5-, 6-, and 7-year-olds ($M(SD) = 6.35(0.97)yr$ s; 53% female). An additional 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, participants were 50 5-, 6-, and 7-year-olds ($M(SD) = 6.19(0.77)yr$ s; 54% female). An additional 3 children were excluded from analysis due to missing date of birth (1), parental interference (1), or experimenter error (1). In Experiment 3, participants were 50 5-, 6-, and 7-year-olds ($M(SD) = 6.20(0.79)yr$ s; 60% female). An additional 5 children were excluded from analysis due to experimenter error (2) or not completing the procedure (3).

No statistical methods were used to pre-determine 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 & 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 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 off-line from the video by the first author and a second researcher blind to the experimental hypotheses. See Supplementary Methods 1 for details.

Materials. We constructed seven rectangular toys ($25 \times 15 \times 15$ 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, and high-reward toys had a plastic orb that lit up different colours and spun around (see Figure 2b for a schematic of the toys used in each condition of Experiment 1 and Figure 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: $M_{DiscoveryTime}(SD) = 0.32(55.26)s$); figuring out the medium-cost toy took longer (Experiment 1: $M(SD) = 9.33(6.37)s$); figuring out the high-cost toys and extra-high cost toy took much longer (Experiment 1: high-cost toys: $M_{DiscoveryTime}(SD) = 81.22(49.42)s$; extra-high-cost toy: $M(SD) = 94.52(61.79)s$) and involved pressing many buttons. If children stopped exploring the high-cost or extra-high-cost toys before successfully activating them, the experimenter prompted children to continue and provided suggestions of what they could try next (e.g., "I wonder what would happen if you pressed two buttons at the same time."). These prompts were delivered as if 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 (order counterbalanced) to ensure they knew how to activate the toys and experienced the ease of activating them after discovery (i.e., 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 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.)

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/>.

References

- [1] Piaget, J. *The child's conception of the world* (Harcort, Brace: New York, 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 University Press, 2006).
- [3] Schulz, L. The origins of inquiry: Inductive inference and exploration in early childhood. *Trends in Cognitive Sciences* **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. *Behavioral and Brain Sciences* **38**, e31 (2015).
- [6] Caro, T. & Hauser, M. Teaching in nonhuman animals. *Quarterly Review of Biology* **67**, 74 (1992).
- [7] Wellman, H. & Cross, D. Theory of mind and conceptual change. *Child Development* **72**, 702–707 (2001).
- [8] Jara-Ettinger, J., Gweon, H., Schulz, L. E. & Tenenbaum, J. B. The Naïve Utility Calculus: Computational principles underlying commonsense psychology. *Trends in Cognitive Sciences* **20**, 589–604 (2016).
- [9] Csibra, G. & Gergely, G. Natural pedagogy. *Trends in Cognitive Sciences* **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. *Cognitive Psychology* **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. *Nature Human Behaviour* **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. *Developmental Psychology* **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 Development* **87**, 723–735 (2016).
- [17] Gweon, H. & Schulz, L. 16-month-olds rationally infer causes of failed actions. *Science* **332**, 1524–1524 (2011).
- [18] Was, A. M. & Warneken, F. Proactive help-seeking: Preschoolers know when they need help, but do not always ask for it. *Cognitive Development* **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. *Psychological Review* **120**, 779–797 (2013).
- [20] Koenig, M. A. & Harris, P. L. Preschoolers mistrust ignorant and inaccurate speakers. *Child Development* **76**, 1261–1277 (2005).

- [21] Gweon, H. & Asaba, M. Order matters: Children's evaluation of underinformative teachers depends on context. *Child Development* **89**, e278–e292 (2018).
- [22] Gweon, H., Shafto, P. & Schulz, L. Development of children's sensitivity to over-informativeness in learning and teaching. *Developmental Psychology* **54**, 2113–2125 (2018).
- [23] Sobel, D. M. & Letourneau, S. M. Children's developing knowledge of and reflection about teaching. *Journal of Experimental Child Psychology* **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. *Monographs of the Society for Research in Child Development* **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. *Journal of Cognition and Development* **17**, 264–284 (2016).
- [26] Baer, C. & Friedman, O. Fitting the message to the listener: Children selectively mention general and specific facts. *Child Development* **89**, 461–475 (2018).
- [27] Rhodes, M., Gelman, S. & Brickman, D. Children's attention to sample composition in learning, teaching and discovery. *Developmental Science* **13**, 421–429 (2010).
- [28] Bass, I. *et al.* Children's developing theory of mind and pedagogical evidence selection. *Developmental psychology* **55**, 286 (2019).
- [29] Gweon, H. & Schulz, L. From exploration to instruction: Children learn from exploration and tailor their demonstrations to observers' goals and competence. *Child Development* **90**, e148–e164 (2019).
- [30] Gergely, G., Nádasdy, Z., Csibra, G. & Bíró, 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 (New York, N.Y.)* **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. *Developmental Science* **6**, e12702–8 (2018).
- [35] Ronfard, S., Was, A. M. & Harris, P. L. Children teach methods they could not discover for themselves. *Journal of Experimental Child Psychology* **142**, 107–117 (2016).
- [36] Repacholi, B. & Gopnik, A. Early reasoning about desires: Evidence from 14-and 18-month-olds. *Developmental Psychology* **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 Development Perspectives* **8**, 90–95 (2014).
- [39] Atance, C. M. & O'Neill, D. K. Episodic future thinking. *Trends in Cognitive Sciences* **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 *Proceedings of the 39th Annual Conference of the Cognitive Science Society* (Cognitive Science Society Austin, TX, 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. *Journal of Political Economy* **97**, 1232–1254 (1989).
- [43] Birch, S. & Bloom, P. The curse of knowledge in reasoning about false beliefs. *Psychological Science* **18**, 382 (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. *Proceedings of the 40th annual meeting of the Cognitive Science Society* (2018).
- [45] Gergely, G. & Csibra, G. Teleological reasoning in infancy: The naïve theory of rational action. *Trends in Cognitive Sciences* **7**, 287–292 (2003).
- [46] Martin, A. & Olson, K. R. When kids know better: Paternalistic helping in 3-Year-Old children. *Developmental Psychology* (2013).
- [47] Wellman, H., Cross, D. & Watson, J. Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development* **72**, 655–684 (2001).
- [48] Goodman, N. D. & Frank, M. C. Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences* **20**, 818–829 (2016).

- [49] Gureckis, T. M. & Markant, D. B. Self-directed learning a cognitive and computational perspective. *Perspectives on Psychological Science* **7**, 464–481 (2012).
- [50] Inzlicht, M., Shenhav, A. & Olivola, C. Y. The effort paradox: Effort is both costly and valued. *Trends in Cognitive Sciences* 1–13 (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 Development* **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. *Behavioral and Brain Sciences* **49**, 1615 (2013).
- [55] Hewlett, B. S., Fouts, H. N., Boyette, A. H. & Hewlett, B. L. Social learning among Congo Basin hunter-gatherers. *Philosophical Transactions of the Royal Society B: Biological Sciences* **366**, 1168–1178 (2011).
- [56] Legare, C. H. Cumulative cultural learning: Development and diversity. *Proceedings of the National Academy of Sciences* **114**, 7877–7883 (2017).
- [57] Tomasello, M. Born (and bred) to help. In *Why We Cooperate*, 1–25 (MIT Press, Cambridge, MA, 2009).
- [58] Bruner, J. S., Jolly, A. & Sylva, K. *Play: Its role in development and evolution* (New York: Basic Books, 1976).

- [59] Vygotsky, L. S. *Mind in society: The development of higher psychological processes* (Harvard University Press, 1980).
- [60] Boyd, R., Richerson, P. J. & Henrich, J. The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences* **108**, 10918–10925 (2011).
- [61] Tomasello, M., Carpenter, M., Call, J., Behne, T. & Moll, H. Understanding and sharing intentions: The origins of cultural cognition. *The Behavioral and Brain Sciences* **28**, 675–91– discussion 691–735 (2005).
- [62] Legare, C. H. & Nielsen, M. Imitation and innovation: The dual engines of cultural learning. *Trends in Cognitive Sciences* **19**, 688–699 (2015).
- [63] Heyes, C. Who knows? Metacognitive social learning Strategies. *Trends in Cognitive Sciences* **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. *Psychological Science* **15**, 661–667 (2004).

Acknowledgements

We thank Carol Dweck, Michael Frank, Ellen Markman, Michael Henry Tessler, Mika Asaba, and Natalia Vélez for their helpful conversations and insightful comments. We thank Grace Bennett-Pierre, Avani Singh, Fernanda Kramer, Angelina Garron, and Nirali Chandaria for their help with data collection and coding. We are grateful to the Palo Alto Junior Museum & Zoo, the Tech Museum of Innovation in San Jose, and the children and families who participated in this research. This work was funded by: John Templeton Foundation Varieties of Understanding grant (HG), James S. McDonnell Scholar Award (HG), and NSF Graduate Research Fellowship (SB). These funding sources had no role in

the conceptualisation, design, data collection, analysis, decision to publish, or preparation of the manuscript for publication.

Author contributions

S.B. and H.G. conceived of and designed the experiments. S.B. collected and analysed the data. J.J-E contributed to designing, implementing, and conducting formal model comparisons with 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.

Figure 1. Full model equation and table of model space. (a) Costs and rewards considered in each teaching plan in the Full model. Activation Reward and Discovery Reward were assumed to be equal. Discovery Cost was always greater than the Activation Cost except for the low-cost toy whose Activation and Discovery costs were 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) Each row shows different models. Each column shows the components included in the models to calculate the utility of each teaching plan. A check mark indicates that the component was included in the model; a cross indicates that it was not included. This visualises how each model captures different considerations that participants might use to decide what to teach.

Figure 2. Experiment 1: Stimuli, behavioural results, and model predictions. (a) Choices made by children in Experiment 1 ($n = 25/condition$): Red represents the number of children who selected the red toy as shown in the top row of (b), and 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 toys used in each condition of Experiment 1; (c) Full model's predictions of children's choices; (d) Predictions made by each alternative model.

Figure 3. Experiment 1: Model comparison by noise-parameter. The log-likelihood ratios for the Full Model vs. the alternative models, $\log(L(\text{Data}|\text{FullModel})/L(\text{Data}|\text{AlternativeModel}))$, as a function of different noise parameters. Values above 0 (meaning that the likelihood ratio is greater than 1 – grey horizontal line in the plot) reflect that the Full model explains the data better than the alternative model; values below 0 (meaning that the likelihood ratio is less than 1) indicate that the alternative better explains the data than the Full. The Full model better explains the data compared to our original 4 alternative (lesioned) models for all noise-parameter values. For the model comparison reported in the main text we use a small degree of noise ($\alpha = 0.1$), shown in green; all comparisons favoured the Full model by at least five orders of magnitude (all likelihood ratios $> 1e5$).

Figure 4. Experiments 2-3: Behavioural results Children’s responses in Experiment 2’s Teach and Play conditions (left; $n = 25/\text{condition}$) and Experiment 3’s Exploration and Instruction conditions (right; $n = 25/\text{condition}$). Yellow represents the number of children selecting the toys along the bottom of the graph as best to teach (or play with) and red represents the number of children selecting the toys along the top of the graph. Error bars are bootstrapped 95% CIs. In Experiment 2, children more often chose the low-reward/high-cost toy in the Teach than 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]), and 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, $p = 0.742$, $OR = 0.65$).

Supplementary Information:

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

Sophie Bridgers¹, Julian Jara-Ettinger², and Hyowon Gweon¹

¹ Department of Psychology, Stanford University

² Department of Psychology, Yale University

Supplementary Methods 1: Experiments 1-3

In this section, we provide a more detailed description of the procedures used in Experiments 1-3.

Coding

Most parents provided parental permission to video-record their child’s participation in the experiment, so these children’s responses were recorded off-line from the video by the first author and a second researcher blind to the experimental hypotheses. For a small number of children without video consent (8 in Experiment 1, 1 in Experiment 2, and 0 in Experiment 3), the experimenter recorded their responses during the experiment. Thus their final choice and follow-up responses were not included in the reliability analyses described below, while their explanations were included.

Procedures

Exploration Phase

All conditions (except the Instruction condition in Experiment 3) started with an Exploration Phase. Children were introduced to two toys and encouraged to explore them in order to figure out how they worked. Children were always given the low-cost toy first, because piloting revealed that children may be less motivated to explore the second toy if the first was difficult to figure out. In the Rewards-only condition of Experiment 1, both toys were low-cost, so children were always given the high-reward toy first.

If children stopped pressing buttons before successfully activating a medium-cost, high-cost or extra-high-cost toy, the experimenter encouraged children to continue. For the medium-cost toy, children could receive one prompt: “Why don’t you keep trying the buttons? I’m sure the toy does something, we just have to figure it out”. For the high-cost and extra-high-cost toys, children received one prompt each time they paused their exploration in the following order: (1) “Why don’t you keep trying the buttons? I’m sure the toy does something, we just have to figure it out” (same prompt as for the medium-cost toy), (2) “Hmm, I wonder what would happen if you pressed two buttons.”, and (3) “Hmm, I wonder what would happen if you tried the big button with one little button at the same time.” Thus, children could receive 0-3 of these prompts depending on whether they stopped exploring before activating the toy, and if so, how many times they stopped. For the extra-high-cost toy, some children received an additional prompt: For this toy, there was a slight delay between when the correct buttons were pressed and the music played, so when children were exploring the toy, if they were pressing the buttons quickly, the experimenter asked, “What if you hold them (the buttons) down for longer?”. The experimenter delivered these prompts as spontaneous suggestions and with curious uncertainty to prevent children from thinking she knew how the toys worked. Critically, the experimenter never explicitly told or showed children how to activate either toy.

For the low-cost toys, all children immediately pressed the single button and successfully activated the causal effect. In Experiment 1, the average time for discovery (i.e., time between the child’s first button press to the first activation of the causal effect) for low-cost toys was $M(SD) = 0.32(55.26)s$, and no children received prompts. For the medium-cost toy, the average discovery time was $M(SD) = 9.33(6.37)s$, and 6 out of the 25 children who explored this toy received a prompt. For high-cost toys, the average discovery time was $M(SD) = 81.22(49.42)s$, and the average number of prompts children received was $M(SD) = 2.25(1.00)$. Finally, for the extra-high-cost toy, the average discovery time was $M(SD) = 94.52(61.79)s$, and the average number of prompts children received was $M(SD) = 2.80(1.41)$. We collapsed across conditions in Experiment 1 that involved a high-cost toy and ran two linear regressions both with choice as the outcome variable; the mean discovery time for high-cost toys was the single predictor variable in one regres-

sion and the average number of prompts for high-cost toys was the single predictor variable in the other. These analyses revealed that neither the average discovery time nor the number of prompts predicted choice (discovery time: $R^2 = 0.017$, $F(1, 69) = 1.19$, $\beta = -0.001$, $p = 0.28$; prompts: $R^2 = 0.006$, $F(1, 69) = 0.38$, $\beta = -0.034$, $p = 0.54$). We focused on the high-cost toys for these exploratory analyses because many more children explored these toys compared to the medium-cost and extra-high-cost toys (i.e., $n = 75$ for the high-cost toys vs. $n = 25$ for the medium-cost and extra-high-cost toys).

Note that in the Instruction condition of Experiment 3, the Exploration Phase was replaced by a *Demonstration* Phase. In this phase, the experimenter explicitly demonstrated how the two toys worked; critically, children did not have an opportunity to explore the toys nor to activate them. For instance, for the high-cost toy, the experimenter said, “If you press this big button here and this little button here at the same time, the toy plays music”; for the low-cost toy, she said, “If you press this big button here, the toy plays music”. After activating the toy, she added, “Isn’t that cool? Nothing else works. This is the only way to make the toy go.” The demonstration was repeated twice per toy, and the order in which the experimenter demonstrated the toys was counterbalanced across subjects.

Choice Phase

In the Choice Phase (in all conditions except the Play condition in Experiment 2), the experimenter asked children which of the two toys she should teach her friend. The exact phrasing was as follows: “I have a friend who asked to borrow these toys and she is going to play with them later all by herself, but she does not know anything about them. So I thought that before she plays with them all by herself, I could help her and teach her how just *one* of these toys works. How should I help her? Which toy should I teach her?”

In the Choice Phase of Experiment 2’s Play condition, instead of choosing which toy should be taught, children chose a toy for themselves to *play* with: “I need to work on something over there for a minute, while I’m gone you can choose one of these toys to play with. Which toy do you want to play with?”.

The toy children pointed to or verbally named was coded as their choice; if children spontaneously changed their answers, then their second choice was coded. The same coding was used for the Follow-up Questions described below. For children’s toy choices in all experiments, agreement between the first author’s and blind researcher’s coding was 100%. As an exploratory measure, children were asked to explain their choice. We provide these results in a separate section below (Choice Explanation Results).

Follow-up Questions

After the critical Choice Phase, we asked two follow-up questions in the following order: (1) “What was harder to figure out? This toy, this toy, or were they both the same?” and (2) “What was cooler? The [music/lights] on this toy, the [music/lights] on this toy, or were they both the same?” Children were also asked to explain their choice. Note that in the Instruction Condition of Experiment 3, children did not explore the toys themselves. Therefore, instead of asking which toy was harder for them to figure out, we showed children a picture of a 6-year-old child (introduced as Mia or Oliver, gender matched to the participant) and said: “[Mia/Oliver] played with these toys earlier all by him/herself, but I was not there to help him/her; which toy do you think was harder for [Mia/Oliver] to figure out?” Children were then asked which toy’s causal effect *they* thought was cooler just as in the other conditions described above. See Supplementary Figure 2A for children’s responses to these questions in Experiments 1-3.

Children’s responses to the first (“harder”) follow-up question suggest that our low-cost toys were indeed perceived as easier to figure out than our medium-cost, high-cost, and extra-high-cost toys. When the costs were matched (i.e., Experiment 1’s Rewards-only condition), the majority of children said the toys were equally hard. Reliability between coders was 100% for all experiments.

Children’s answers to the second (“cooler”) follow-up question also reflect an overall preference for the lights over the music effect. Reliability between coders was 100% for Experiments 1 and 3, and 98% for Experiment 2. All disagreements were resolved by the first author. However, we treat this question as an exploratory measure; while children’s responses could reveal their perception of the relative rewards of the toys’ causal effects, they may additionally reflect other factors, such as

their choice in the Choice Phase or the perceived difficulty of the toys. Furthermore, our reward manipulation was arguably more subjective than our cost manipulation (i.e., hardly any children thought the low-cost toys were harder to figure out than the high-cost toys, but some children perceived the music, our intended low-reward effect, as cooler than the lights, our intended high-reward effect). Therefore, to measure children’s evaluations of the relative rewards of the toys’ causal effects more precisely, we ran a separate norming study (see Norming Experiment below and Supplementary Figure 2B).

Supplementary Results 1

Further Analyses of Individual Conditions in Experiment 1

The main test of our theory was how well our Full model (compared to the alternatives) predicted the pattern of children’s responses across all six conditions of Experiment 1. However, we can also look at children’s behaviour within each condition. Translating our theory to the null-hypothesis testing framework, we use one-tailed binomial tests in each condition where the model makes a directed prediction and two-tailed tests otherwise. In the Rewards & Costs condition, children were significantly more likely to select the high-reward/high-cost than the low-reward/low-cost toy as best to teach (21/25, 84%; one-tailed Binomial test, $p < 0.001$, bootstrapped 95% CI [0.68, 0.96]). In the Different Costs condition, children were significantly more likely to select the high-cost/low-reward than the low-cost/low-reward toy (21/25, 84%; one-tailed Binomial test, $p < 0.001$, bootstrapped 95% CI [0.68, 0.96]). In the Different Rewards condition, children were marginally more likely to select the high-reward/low-cost (lights) than the low-reward/low-cost (music) toy (17/25, 68%; one-tailed Binomial test, $p = 0.054$, bootstrapped 95% CI [0.50, 0.86]). In the conflict conditions, children showed no significant preference for either toy in the Medium-Cost Conflict condition (11/25, 44% selected the low-reward/medium-cost toy; one-tailed Binomial test, $p = 0.345$, bootstrapped 95% CI [0.24, 0.65]) nor in the High-Cost Conflict condition (15/25, 60% selected the low-reward/high-cost toy; two-tailed Binomial test, $p = 0.424$, bootstrapped 95% CI [0.40, 0.79]). Finally, in the Extra-High Cost Conflict condition, children were significantly more

likely to select the low-reward/extra-high-cost toy than the low-reward/low-cost toy (19/25, 76%; one-tailed Binomial test, $p = 0.007$, bootstrapped 95% CI [0.58, 0.92]).

Zooming in on the Conflict conditions, we also see that children in the Medium-Cost Conflict condition were more likely to select the high-reward toy than children in the Extra-High Cost Conflict condition (two-tailed Fisher’s exact test, $p < 0.042$, $OR = 0.26$). Additionally, a logistic regression with the three conflict conditions as a predictor of children’s teaching decisions revealed a significant linear effect across these conditions ($\beta = -0.226$, $t(74) = -2.348$, $p = 0.0216$), providing evidence that children’s tendency to select the toy higher in rewards as best to teach decreased as the difference between the two toys’ discovery costs increased. Note that the Medium Cost and Extra-High Cost Conflict conditions were run during the review process. The preregistration for these conditions can be found at: <https://osf.io/5dmga/>.

Choice Explanation Results

We were curious to see whether the content of children’s explanations for why they chose a toy to teach or play with would provide additional insight into how they made these decisions. We did not have a priori hypotheses about this data and coded it as an additional, exploratory measure.

Coding Criteria

We categorised children’s explanations for why they chose a particular toy as appealing to the *reward* of the effect and/or the *cost* of discovery. Explanations that mentioned a toy’s causal effect (e.g., “Because it spins”), how a toy made the child feel (e.g., “Because it makes me happy”), or a child’s explicit preference (e.g., “Because I like it”) were coded as appealing to reward and not otherwise. Explanations that mentioned a toy’s button(s) (e.g., “Because that one has a lot of buttons”), the ease or difficulty of the toy (e.g., “Because it’s harder”), or the learner’s knowledge state about how the toy works (e.g., “Because she doesn’t know which buttons”) were coded as appealing to cost and not otherwise. Explanations could meet the criteria for both dimensions (e.g., “Because it’s much cooler and easier”) or neither (e.g., “Because it’s the first one”). Children who did not provide an explanation (i.e., did not respond or said “I don’t know.”) were coded as

giving no response. As there were very few responses coded as *both* or *neither*, we describe the results mainly in terms of appeals to cost or reward (also see Supplementary Figure 1).

Coding agreement was 100% for Experiments 1 and 3, and 98% for Experiment 2. All differences were resolved by the first author. We qualitatively describe the data below.

Explanation Results

In Experiment 1, children spontaneously appealed to differences in cost and reward in ways that reflected the main dimension(s) of contrast across the toys used in each condition. They appealed mainly to reward in the Different Rewards condition (reward: 21/23, 91%; cost: 2/23, 9%) and mainly to cost in the Different Costs condition (reward: 3/25, 12%; cost: 20/25; 80%). Even children who chose a low-reward or a low-cost toy often appealed appropriately to the rewards or costs of the toys, respectively (e.g., “Because it makes nice noises.”; “Because it’s easier than that one.”)

Note that in the Rewards & Costs and the three Conflict conditions, the toys differed in both rewards and costs, so children could in principle appeal to either dimension of contrast to explain their choice. In the Rewards & Costs condition, when one of the toys was higher in both reward and cost, children, as a group, appealed to both: 15/25 (60%) appealed to reward and 9/25 (36%) appealed to cost. In the Medium Cost Conflict condition, when high rewards were pit against medium costs, children as a group, like the Rewards & Costs condition, appealed to both: 13/24 (54%) appealed to reward and 11/24 (46%) appealed to cost.

In the High Cost and Extra-High Cost Conflict conditions, where high rewards were pit against high and extra-high costs, respectively, children showed a tendency to appeal more to cost (High Cost: 19/24, 79%; Extra-High Cost: 18/22, 82%) and less to reward (High Cost: 4/24, 17%; Extra-High Cost: 3/22, 14%). Children’s tendency to appeal to cost in these conditions, suggests that this contrast may have been more salient to them than the contrast between rewards, as well as more salient than the cost contrast in the Medium Cost Conflict condition. In the High and Extra-High Cost Conflict conditions, some children even explained their choice of the high-reward/low-cost toy in terms of cost (i.e., “Because it’s easier”). It is interesting to speculate whether these children

might have had an aversive reaction to the toys that were difficult to figure out or if they were possibly thinking about the experimenter and her cost of teaching. We discuss in the main text how our model could be extended to incorporate the teacher’s utility. Though we tried to guard against this consideration in our current paradigm, it is possible some children were thinking about the teacher.

In Experiment 2, children appealed to both reward (10/25, 40%) and cost (14/25, 56%) in the Teach condition, while in the Play condition, children mainly appealed to reward (18/21, 86%) and less to cost (2/21, 10%), suggesting that their teaching decision was driven by the discovery costs, whereas their personal preference for a toy was driven mainly by rewards. In Experiment 3, children rarely appealed to reward (4/22, 18%, in the Exploration condition and 3/22, 14%, in the Instruction condition), but mainly to cost regardless of prior experience with the toys (17/22, 77%, in Exploration and 18/22, 82%, in Instruction).

More generally, across all experiments, children’s explanations were systematic and relevant to the dimensions we manipulated across conditions. As noted in the main text, two children in the High Cost Conflict condition of Experiment 1 chose to teach the high-reward/low-cost toy (instead of the low-reward/high-cost toy) because they wanted to give the learner the opportunity to explore it. When asked to explain their choice, one child said, “...because it’s easier and then she could figure out hard things”, and the other said, “How about we let her figure this one (the high-cost toy) out by herself because this is going to be a challenge for her and you could help her with this one (the low-cost toy)”. These two children were some of the oldest in our sample, pointing to the possibility that this nuanced understanding of what or how to teach (i.e., the benefit of scaffolding or the value of providing a challenging task) develops later in childhood. Such intuitions might be mediated by the teacher’s assessment of the difficulty of discovery for the learner. While we remain cautious about interpreting these infrequent responses, they raise interesting questions for future work.

Supplementary Figure 1: The number of children who appealed to rewards, costs, rewards & costs (both), or neither when explaining their choice of toy to teach (or to play with) in each condition of Experiments 1-3. $N = 25$ per condition but not all children provided explanations (Rewards & Costs: $n = 25$; Different Costs: $n = 25$; Different Rewards: $n = 23$; Medium Cost Conflict: $n = 24$; High Cost Conflict: $n = 24$; Extra-High Cost Conflict: $n = 22$; Teach: $n = 25$; Play: $n = 21$; Exploration: $n = 22$; Instruction: $n = 21$).

Supplementary Methods and Results 2: Norming Experiment

In order to determine the reward function of our models, we ran a separate study to measure which causal effect, lights or music, children found more rewarding. A separate group of children ($N = 25$; $M(SD) = 6.15(0.93)$ yrs; 60% female) who did not participate in Experiments 1-3 were introduced to the low-reward/low-cost and high-reward/low-cost toys (i.e., the toys only differed in reward and were both easy to activate). Children first explored the toys and then were asked: “What do you think is cooler? The lights in this toy or the music in this toy?” Children were not given the option to answer “same”, so if children said they were both cool, the experimenter followed-up, “If you had to choose just one, which one would you say is cooler?” (order counterbalanced across children). Eighty-percent of children chose the lights as the “cooler” effect (bootstrapped 95% CI [0.64, 0.96]). See below (Reward Function) and Supplementary Figure 2C for details.

Supplementary Figure 2: Children’s responses to the follow-up questions. (A) The number of children who selected each of the causal effects as cooler and each of the toys as harder to figure out by condition in Experiments 1-3; red corresponds to choosing the toy (or effect) displayed on the top of each graph; yellow corresponds to choosing the toy (or effect) displayed on the bottom of each graph. “Same” judgements, meaning children said the two toys/effects were equally hard/cool, are shown in grey. All 25 children per condition answered both questions, except for the Cooler question in the Medium Cost Conflict condition ($n = 24$) and the Play condition ($n = 23$). (B) The number of children ($N = 25$) in the Norming Experiment who selected the lights (in red) or music (in yellow) as the cooler effect; eighty-percent chose the lights (bootstrapped 95% CI [0.64, 0.96]).

Supplementary Note 1: Computational Model

Model Implementation

Our main (Full) model assumes children pursue the teaching plan with the highest expected utility for the learner. We define the utility of each teaching plan as the linear sum of the following components: the expected cost the learner incurs (Activation Cost, C_A) and the expected reward she obtains (Activation Reward, R_A) when activating the toy she was taught (Chosen toy); the expected cost the learner incurs to discover how the other untaught (Unchosen) toy works (Discovery Cost, C_D); and the reward she obtains upon discovery (Discovery Reward, R_D). Figure 1 in the

main text includes the exact equations. Below we provide more details about how each component of the utility function (the cost and the reward functions) and the uncertainty parameter (p_E) were implemented.

Cost Function

High-cost toys had 7 buttons, and the extra-high-cost toy had 12 buttons; for both types of toys, 2 of the buttons had to be pressed together to generate an effect. The medium-cost toy had 7 buttons, only one of which had to be pressed. The low-cost toys had one button. The cost functions were determined by assuming that each button press results in a constant cost set to 1, assuming that the cost of activating or discovering how to activate a toy was physical in nature (i.e., button-presses). Using a different value does not change our results because all other parameters are defined in terms of this basic cost unit.

Activation cost: C_A was set to 1 for the low-cost toys and 2 for the high-cost and extra-high-cost toys.

Discovery cost: C_D was set to the expected number of buttons the learner would need to press to discover how to activate the toy. We assume that the learner would first try simple activation sequences (pressing one button at a time) and then try more complex and costly sequences (i.e., trying combinations of two buttons) once the simpler hypothesis space is depleted. These assumptions are consistent with children’s actual exploratory behaviours: Most children began by pressing each button once and only attempted combinations of two buttons once they had pressed each button individually. C_D was set to 1 for low-cost toys. For the medium-cost, high-cost, and extra-high-cost toys, both an analytic solution, as well as simulation provide an expected discovery cost of 4, 29, and 79, respectively.

Analytically, the expected cost is the expected number of actions (button presses) the learner needs to take to activate the toy. For the high-cost and extra-high cost toys, under the assumption that children first press each button individually, the learner will always incur an initial cost of n for a toy with n buttons (pressing each button). Once the learner has tested these simplest hypotheses, the next space of hypotheses consists of pairs of buttons. For a toy with n buttons, this space has

$\binom{n}{2}$ hypotheses. In this case, different learners might need different numbers of attempts – some might be lucky and find the correct combination of two buttons quickly, others might be unlucky and have to try almost all combinations, and most might fall somewhere in between. Each press of a pair of buttons produces a cost of 2. Thus, the total expected discovery cost is given by:

$$C_D = n + 2p(t=1) + 4p(t=2) + 6p(t=3) + \dots + 2 * \binom{n}{2} p(t = \binom{n}{2})$$

where the first term (n) corresponds to the initial cost of trying each button once, and the rest of the terms correspond to the expected cost when searching in the hypothesis space with two buttons. In this second term, $p(t=x)$ is the probability that the learner finds the correct combination on the x -th attempt. This formula can be summarised as:

$$C_D = n + \sum_{i=1}^{\binom{n}{2}} 2 * ip(t=i)$$

If the learner attempts all combinations in a random order, then the probability that they succeed on any given trial is the same. Because there are $\binom{n}{2}$ possibilities, then $p(t=x) = \frac{1}{\binom{n}{2}}$ for all values of x . Substituting this probability, we obtain:

$$C_D = n + 2 \sum_{i=1}^{\binom{n}{2}} \frac{i}{\binom{n}{2}}$$

Because the sum of integers 1 to n is given by the formula $(n * (n+1))/2$ the expected discovery cost becomes:

$$C_D = n + 2 \frac{\binom{n}{2} * ((\binom{n}{2}) + 1)}{2} * \frac{1}{\binom{n}{2}}$$

which reduces to:

$$C_D = n + \binom{n}{2} + 1$$

Therefore, when $n = 7$ (because the high-cost toys had 7 buttons), $C_D = 7 + 22$, with a final

value of $C_D = 29$. When $n = 12$ (i.e., the extra-high-cost toy with 12 buttons), $C_D = 12 + 67$, with a final value of $C_D = 79$.

For the medium-cost toy, where the activation sequence is just one button, the formula is:

$$C_D = \frac{\binom{n}{1} + 1}{2}$$

When $n = 7$ (because the medium-cost toy had 7 buttons), $C_D = \frac{7+1}{2}$, with a final value of $C_D = 4$. These analytical solutions can also be verified through Monte Carlo sampling. The code for deriving the cost is available at <https://osf.io/wunbq/>.

Note that C_A and C_D for the low-cost toys are both minimal (1), but for the high-cost and extra-high-cost toys, C_D is significantly higher than C_A (29 or 79 vs. 2). For the medium-cost toy, C_D is slightly higher than C_A (4 vs. 1). Thus, demonstrating a high-cost or extra-high-cost toy to a learner increases the overall utility of the corresponding teaching plan by replacing the high C_D (that the learner would have incurred if taught the low-cost toy) with a comparatively low C_A .

Reward Function

The toys used in our experiments either had an orb that lit up in different colours and spun around (intended to be the high-reward effect) or played music (intended to be the low-reward effect). The Activation Reward, R_A , and Discovery Reward, R_D , of a particular toy were assumed to be equal and were treated as variables to allow for the possibility that different participants assign different subjective rewards to each causal effect. As we describe in the main text, we modelled the distribution of rewards using beta distributions multiplied by a constant to cover a wider range of values.

Beta distributions are determined by two parameters: α and β . We searched for the parameters of two beta distributions such that the probability that a sample from the lights reward space would be higher than a sample from the music reward space 80% of the time (as determined by the Norming Experiment). For simplicity, we only considered distributions that were mirror images of each other. That is, we assumed that one distribution had parameters (x,y) and the second distribution had parameters (y,x). This simplifies the search over parameters.

Supplementary Figure 3: (a): Space of beta parameters. The x-axis shows potential α parameters and the y-axis shows potential β parameters. For each combination we considered the probability that the pair of Beta distributions with parameters (α, β) and (β, α) would yield a higher reward from the first distribution at a rate between 75% and 85% of the time. The grey region shows the space of parameters that match our empirical data from the Norming Experiment. (b)-(d): Beta distributions generated using the parameters on the left, centre, and right end of the grey region from panel (a). We use the Beta distributions shown in panel (b) for generating our model predictions reported in the main text.

For any pair of parameters, we assessed the fit through Monte Carlo sampling (using $n = 10000$ samples for each combination of parameters; see <https://osf.io/wunbq/> for code). Because Monte Carlo produces approximate probability estimates and because of possible noise in our Norming Experiment results (80% of children preferring the lights effect and 20% preferring the music), we considered any set of parameters to fit our data when the probability estimates fell between 75% and 85%. Supplementary Figure 3 shows the results from this parameter search; the grey region in panel (a) shows the space of parameters that fit our data from the Norming Experiment.

To generate the predictions reported in our main text, we use the values from the left end of the region ($\alpha = 1.875$, $\beta = 1$). This is because as the α and β parameters increase, beta distri-

butions have less variance; consequently, these distributions can only be consistent with children’s preferences in the Norming Experiment when the reward difference is sufficiently small for the two low-variance distributions to overlap (see Supplementary Figure 3b-d). Thus the main text presents the values that allow for the highest potential difference in the toys’ rewards.

Our model predictions are robust to a range of parameters that define the reward space. In the middle range, we tested parameters $\alpha = 6$, and $\beta = 4.15$; on the right end, we tested parameters $\alpha = 10$, $\beta = 7.5274$. Summaries of model predictions using different values of α and β parameters from the middle, and the right end of the region are available at <https://osf.io/wunbq/>.

Exploration Certainty Parameter

While teaching the Chosen toy ensures that the learner will activate it with minimal costs, the child participant may be unsure whether the learner will actually explore the Unchosen toy or even if she does, that she will succeed in activating it. To reflect this uncertainty, in the model we included an additional parameter p_E that reflects the probability that a learner interacts with the Unchosen (Explored) toy. To generate our model predictions reported in the main text, we consider a 50% probability that the learner will explore the Unchosen toy ($p_E = 0.5$). Again, our model fit is robust to a wide range of exploration certainty parameter values. We test the robustness using the beta distribution we present in the main text (Supplementary Figure 3b) and exploration certainty parameters: $p_E = 0.3, 0.5, 0.7$, and 0.99 (Supplementary Figure 4).

Alternative (“lesioned”) Models

Our full model assumes that children consider both the expected costs and rewards of each teaching plan (which includes considerations of both the Chosen and the Unchosen toys). However, it is also possible that children instead use simpler strategies and only reason about a subset of these variables. Thus in addition to the full model, we implemented simpler (“lesioned”) models that correspond to four plausible alternative hypotheses.

First, since the learner can easily obtain the rewards of both toys if she knows how they work, children might only be concerned about the relative difficulty of figuring out the toys. If so, they

Supplementary Figure 4: The Full model’s and alternative models’ predictions for the six conditions of Experiment 1 at different values of p_E : 0.3, 0.5, 0.7, and 0.99

might only consider the costs the learner will incur when discovering the Unchosen toy, and when activating the Chosen toy (Costs-only model). Alternatively, once children know how each of the toys work, they might ignore the learner’s knowledge state and only focus on teaching what is cool. If so, they might only consider the rewards the learner obtains after successfully activating each toy (Rewards-only model). Yet another possibility is that children consider both the costs and rewards but are only concerned with what is being taught and do not think about the experience the learner will have exploring the untaught toy (Instruction-only model). Finally, children might take what is being taught for granted and focus instead on the costs and rewards of having to explore and figure out the other toy (Exploration-only model). Supplementary Figure 5 summarises the different considerations of the full and alternative models. We assess the relative performance of these different models by comparing how well each model’s predictions correspond to children’s choices.

Model Fits

To evaluate how well each model captures children’s responses, we calculated the likelihood of each model generating the pattern of results across conditions in Experiment 1. Since our models are idealised formalisations, we added a small probability $\alpha \in (0, 1)$ of noise in children’s choices, to consider the possibility that some children get distracted and respond randomly. Thus, we assumed that children sample a response from the model’s distribution with probability $1 - \alpha$ and that they respond randomly with probability α . In the main text and here, we present results with the noise parameter $\alpha = 0.1$ (see Model Implementation section). Our qualitative conclusions, however, hold for a wide range of values of this noise parameter, as well as the exploration uncertainty parameter (see <https://osf.io/wunbq/> for model code and additional analyses showing the robustness of our findings with respect to this parameter, as well as Supplementary Figure 6).

After computing each model’s likelihood, we computed the ratio between the Full model’s likelihood and each alternative model’s likelihood. For each alternative model, a value below 1 indicates that it explains the data better than the Full model, and a value over 1 suggests that the Full model better explains the data. The Full model had the highest likelihood of generating the observed data; all comparisons favoured the Full model by at least 5 orders of magnitude (all likelihood ratios $> 1e5$).

To better understand why the Full model best captures children’s choices, we can consider why the alternative models failed to predict participant performance in various conditions. Here we focus on their most salient limitations. The Costs-only, Rewards-only, and Instruction-only models fail to capture children’s choices in the Conflict conditions because they choose to teach the high-reward/low-cost toy at the same rate across all three conditions, and either over-predict (Rewards-only and Instruction-only) or under-predict (Costs-only) children’s tendency to select this toy. The Exploration-only model’s predictions vary with the discovery costs of the toys; thus it makes graded predictions across the Conflict conditions but over-predicts children’s tendency to select the higher-cost toy. This is because without considering the utility of instruction, the decision to teach is based on the utility of exploration and this is much lower for the low-reward/higher-cost toys in these conditions than for the high-reward/low-cost toy. The Full model also makes graded

predictions across these conditions, and because it considers the utility of both instruction and exploration, its predictions better capture children’s responses.

For the non-conflict conditions, the Costs-only model fails to capture children’s choices in the Different Rewards condition: It predicts children will be evenly split between the two toys because they are both easy to activate, whereas the Full model predicts a preference for the toy higher in reward. The Rewards-only model’s critical failure was in the Different Costs condition, as it predicts chance performance, while the Full model predicts a preference for the high-cost toy. The Instruction-only model also fails here but by choosing to teach the easier toy; this is because without considering the other toy’s discovery cost, the factor that determines the choice is the Activation Cost, which is slightly lower for the low-cost toy ($C_A = 1$) than for the high-cost toy ($C_A = 2$). Finally, the Exploration-only model most notably fails to predict responses in the Different Rewards condition; the toys’ costs were matched, so this model predicts children will teach the low-reward toy (so as to let the learner discover the high-reward toy alone); however, the Full model predicts a preference for the high-reward toy, which is also consistent with children’s choices. Altogether, the Full model was the only model that accurately predicted responses across all six conditions in Experiment 1. Figures 1d and 2b in the main text provide a summary of each model’s predictions.

Model without Integration: Rewards-or-Costs model

Motivation

The space of alternative models we considered allowed us to ask whether children were considering both the expected costs and rewards of the toys. The failure of the Costs-Only and Rewards-Only models to capture the pattern of children’s choices across the conditions of Experiment 1 suggests that both factors contributed to children’s decisions. However, an assumption of our Full Model, not tested by these alternatives, is that these decisions involve a trade-off between costs and rewards (i.e., the ability to integrate costs and rewards into a single representation of utility). One, however, could imagine a model that takes both rewards and costs into account but rather than integrating them, compares them each independently. Here, we implement such a model.

Supplementary Figure 5: (A) Number of children in each condition of Experiment 1 selecting either the red or yellow toy to teach; all CI’s are bootstrapped using 10,000 samples. Rewards & Costs: 84% selected high-reward/high-cost toy (yellow), one-tailed Binomial test, $p < 0.001$, 95% CI [0.68, 0.96]; Different Costs: 84% selected high-cost/low-reward toy (yellow), one-tailed Binomial test, $p < 0.001$, 95% CI [0.68, 0.96]. Different Rewards: 68% selected high-reward/low-cost toy (red), one-tailed Binomial test, $p = 0.054$, 95% CI [0.50, 0.86]; Medium-Cost Conflict: 44% selected low-reward/medium-cost toy (yellow), one-tailed Binomial test, $p = 0.345$, 95% CI [0.24, 0.65]; High-Cost Conflict: 60% selected low-reward/high-cost toy (yellow), two-tailed Binomial test, $p = 0.424$, 95% CI [0.40, 0.79]; Extra-High Cost Conflict: 76% selected low-reward/extra-high-cost toy (yellow), one-tailed Binomial test, $p = 0.007$, 95% CI [0.58, 0.92]). Children were more likely to select the high-reward toy in the in the Medium-Cost Conflict compared to the Extra-High Cost Conflict condition (two-tailed Fisher’s exact test, $p < 0.042$, $OR = 0.26$); children’s tendency to select the toy higher in rewards decreased across the three conflict conditions as the difference in the discovery costs increased ($\beta = -0.226$, $t(74) = -2.348$, $p = 0.0216$). (B) Full model’s predictions (left) and Rewards-or-Costs model’s predictions (right) of children’s choices.

Implementation

This model compares the expected rewards of the toys, compares the expected discovery costs of the toys, and teaches the toy that is higher in rewards or higher in costs. More specifically, it considers both $R_D(ToyX) - R_D(ToyY)$ and $C_D(ToyX) - C_D(ToyY)$ but without integrating the two; if one is 0, it uses the other (i.e., Different Costs and Different Rewards conditions); if both are non-zero and have the same sign, the model selects the toy with higher rewards and costs (i.e., Rewards & Costs condition); if both are non-zero and in conflict (i.e., have opposite signs), it chooses randomly between the two toys, equivalent to randomly choosing a dimension on which to focus (i.e., the Conflict conditions). In these situations of conflict, one could also imagine a model that prioritises costs over rewards or vice versa, but for simplicity, this model has no preference. The model also chooses randomly if both differences are zero, but we do not empirically test this case. Supplementary Figure 5 shows the behavioural results alongside the predictions of this model, what we call the “Rewards-or-Costs” model, and the predictions of our Full model. Noise is $\alpha = 0.1$ for both models and the exploration parameter is $p_E = 0.5$ for the Full model, the same values

assigned to these parameters to generate predictions in the main text.

This model without integration (Rewards-or-Costs model) does quite well in predicting children’s behaviours in the Rewards & Costs, Different Costs, and Different Rewards conditions because in these conditions, one can just focus on one dimension and make the same choices as predicted by our Full model. In the Conflict conditions, however, one toy is higher in rewards and the other is higher in costs; thus the model randomly chooses a toy and this preference does not change depending on the magnitude of difference in discovery costs. Therefore, it predicts that children’s preference for each toy should not differ across these three conditions. The predicted proportion of children selecting the toy higher in discovery costs in these conditions is a little above 50% because the model assumes that some children actually prefer the music reward (i.e., for some children, these conditions are similar to the Rewards & Costs condition where one toy is both cooler and harder to figure out).

In contrast to these model predictions in the Conflict conditions, the empirical data suggest a significant increase in children’s tendency to choose the toy that is higher in discovery cost as a function of the magnitude of the difference between the two toys’ discovery costs ($\beta = -0.226$, $t(74) = -2.348$, $p = 0.0216$). Thus, the Full model appears to better capture the qualitative structure and relative ordering of children’s preferences across conditions.

We note however that when we compare the likelihoods of the data under our Full model vs. Rewards-or-Costs model, we find that the likelihood ratio slightly favours the Rewards-or-Costs model. In other words, the Full model provides a better qualitative fit to the structure of children’s choices across conditions, but the Rewards-or-Costs model provides a better quantitative fit in its point estimates of children’s choices within each condition. The Rewards-or-Costs model achieves a slightly better quantitative fit because in the Conflict conditions, the Full model predicts relatively large differences between these conditions whereas the observed differences in children’s choices are quite small; the Rewards-or-Costs model makes the same predictions for all three conflict conditions, but critically this value falls at a sweet spot that makes the data relatively more likely than does the Full model’s predictions under this particular model comparison method. More specifically, the Full model underestimates the proportion of children who select the toy higher in discovery cost in the

Supplementary Figure 6: Log-likelihood ratios ($\log(L(\text{Data} - \text{Full Model})/L(\text{Data} - \text{Alternative Model}))$) as a function of different noise parameters. Values above 0 (meaning that the likelihood ratio is greater than 1) reflect that the Full model explains the data better than the alternative model; values below 0 (meaning that the likelihood ratio is less than 1) indicate that the alternative better explains the data than the Full.

Medium Cost and High Cost conflict conditions, whereas the Rewards-or-Costs model's prediction for these two conditions is closer to the empirical proportions. If you compare the likelihoods of these models at different levels of noise, however, the Full model does better than the Rewards-or-Costs model at higher noise levels (see Supplementary Figure 6). At higher levels of noise, the Full model's predictions are made less extreme: It predicts more children will select the toy higher in discovery cost and the Reward-or-Costs model predicts more children will select the toy higher in rewards, so the likelihood ratio shifts to favouring the Full model.

To summarise, although the Rewards-or-Costs model provides a quantitative fit to the data

close to that of our Full model, it fails to capture the empirical differences between the Conflict conditions. Prior empirical work has shown that young children and even infants integrate rewards and costs when reasoning about others’ behaviours^{1,2}, so it is plausible that the 5- to 7-year-olds in our experiments were capable of such integration. Moreover, it is difficult to estimate the precise shape of children’s subjective cost and reward functions. Our model predictions regarding the exact proportion of children’s toy choices in each condition depend on how these functions are defined. In the current study, we relied on one (and our current best) way to approximate these functions, but future work could further improve the fit of our model by better estimating these functions.

Model with Observation Reward

Motivation

Our full model well captures children’s decisions across a wide range of exploration certainty and noise parameters. However, at very high levels of exploration certainty, the predictions for the Conflict conditions (high-reward/low-cost toy vs. a toy that is either medium, high, or extra-high in discovery cost) overwhelmingly favour the higher-cost toy, such that there is little to no difference across conditions (compared to children’s increasing preference for the higher cost toy across these conditions). The model makes these predictions because the importance of teaching the high-reward/low-cost toy (which would guarantee the high reward for the learner) becomes less significant when it is highly likely that the learner will explore this toy; since the cost of discovery is so low, there is little need to teach the learner how to activate it. Instead, it is more beneficial to teach the toy higher in cost to save the learner the exploration cost and ensure that the learner obtains the reward, even though it is relatively lower than the reward from the other toy.

In light of these predictions, how should we interpret children’s increasing preference for the higher-cost toy across these conditions? One possibility is that children are indeed agnostic to whether or not the learner will actually play with the Unchosen toy (represented by $p_E = 0.5$, as presented in the main text). However, one might question whether 50% exploration chance is too low to capture children’s reasoning, since the experimenter told children that the learner would play with both toys. In fact, another possibility is that a higher exploration certainty parameter

		Model Parameters				
Models		Activation Reward	Activation Cost	Discovery Reward	Discovery Cost	Observation Reward
	Full Model	✓	✓	✓	✓	✓
	Costs-only	X	✓	X	✓	X
	Rewards-only	✓	X	✓	X	✓
	Instruction-only	✓	✓	X	X	✓
	Exploration-only	X	X	✓	✓	X

Supplementary Figure 7: Model space. Each row shows different models. Each column shows the components included in the models to calculate the utility of each teaching plan. Check mark indicates that the component was included in the model; cross indicates that it was not included. This visualises how each model captures different considerations that participants can use to decide what to teach.

better captures children’s reasoning, but the current model is missing a variable that contributes to children’s decisions about what is best to teach.

As discussed in the main text, our model can easily be extended to include other considerations. One plausible component that was not included in the full model is the utility of the learner observing the teacher’s demonstration of the Chosen toy (i.e., the utility of *observation*).

Implementation

In watching the teacher demonstrate this toy, the learner does not incur a cost (i.e., the cost of observation is 0), but does get to enjoy the reward of the toy’s effect. The reward of observation for the Chosen toy is equal to its reward of activation. The “extended full” model that includes this term computes the learner’s expected utility under a given teaching plan as the linear sum of the following components: By being taught the Chosen toy, the learner incurs no cost but obtains the

Observation reward (R_O); upon activating the Chosen toy, the learner incurs a small Activation Cost (C_A) and obtains the Activation Reward (R_A); in exploring the other, Unchosen Toy, the learner incurs a Discovery Cost (C_D), and obtains the Discovery Reward (R_D) upon successful discovery. Again, we can consider different levels of exploration certainty (p_E), and compare this model (as well as the four alternative models: Costs-only, Rewards-only, Instruction-only, and Exploration-only models) to children’s choices in Experiment 1. This extends the space of model parameters to include the Observation Reward in the Full Model, Rewards-Only, and Instruction-Only models.

Supplementary Figure 7 summarises the different considerations of the full and alternative models. The Costs-only and Exploration-only models are identical to those of the same name considered in the main text, while the Rewards-only and Instruction-only models include the Observation reward of the Chosen toy. Again, we find that the Full model with Observation reward best captures children’s performance compared to the four alternatives, and furthermore, model fits are robust to different values of exploration chance and noise. Even at high-levels of exploration chance, this model generates predictions that qualitatively match children’s decisions in the Conflict conditions of Experiment 1.

Both Full models, either with or without the Observation reward, are theoretically plausible and make reasonable predictions for our results. The current experiment does not distinguish between these two models. While teasing the two models apart is out of the scope of our current work, future experiments that directly manipulate exploration certainty could determine whether a consideration of the utility of observation influences children’s decisions. More generally, we hope that our work will motivate studies that incorporate a richer set of considerations, including the teacher’s own utility, which can also be easily incorporated into the model.

Supplementary References

- [1] Jara-Ettinger, J., Gweon, H., Schulz, L. E. & Tenenbaum, J. B. The Naïve Utility Calculus: Computational principles underlying commonsense psychology. *Trends in Cognitive Sciences* **20**, 589–604 (2016).

- [2] Liu, S., Gonzalez, G. & Warneken, F. Worth the wait: Children trade off delay and reward in self- and other-benefiting decisions. *Developmental Science* **6**, e12702–8 (2018).