

# Children consider others' expected costs and rewards when deciding what to teach

Sophie Bridgers<sup>1</sup> (sbridge@stanford.edu), Julian Jara-Ettinger<sup>2</sup> (jjara@mit.edu) & Hyowon Gweon<sup>1</sup> (hyo@stanford.edu)

<sup>1</sup>Department of Psychology, Stanford University, Stanford, CA 94305

<sup>2</sup>Department of Brain and Cognitive Sciences, MIT, Cambridge, MA 02139

## Abstract

Humans have an intuitive sense of how to help and inform others even in the absence of a specific request. How do we achieve this? Here we propose that even young children can reason about others' expected costs and rewards to flexibly decide what is best for others. We asked children to choose one of two toys to teach to another agent while systematically varying the relative costs and rewards of discovering each toy's functions. Children's choices were consistent with the predictions of a computational model that maximizes others' utilities by minimizing their expected costs and maximizing their expected rewards. These results suggest that even early in life, children draw rational inferences about others' costs and benefits, and choose to communicate information that maximizes their utilities.

**Keywords:** social cognition, prosocial decision-making, communication, pedagogy, naive utility calculus

## Introduction

Life is a series of small and big decisions. While some of these decisions involve choosing what is best for ourselves, frequently they are made with other people in mind. Humans, even in early childhood, often deviate from egoistic behavior and act to benefit others (e.g., Fehr & Fischbacher, 2003; Warneken & Tomasello, 2006).

Concerns for others not only motivate commendable sacrifices but also underlie some of the most mundane communicative acts. Imagine, for instance, that a guest at your house is going to the restroom to wash her hands. To complete this goal, you know she will need to turn on the lights and operate the faucet, and you can predict whether she will have any trouble with these tasks. Based on this, you might prioritize communicating the things that are trickier to figure out (e.g., an unusual faucet, "first turn the lever and pull it towards you"), while leaving simpler things untold (e.g., a normal light switch, "flip up the switch to turn on the lights"). Despite the apparent simplicity of this decision, choosing what to communicate based on what you think others need to know and what they can figure out by themselves is far from trivial.

There are many ways to help and inform others, but our time and resources are limited. Therefore, we must prioritize what to share and how to help. How do we constrain such decisions? Offering help or information can be straightforward when it is clear what our social partners want or need. Under these circumstances, even preverbal infants help others by handing objects out of reach (Warneken & Tomasello, 2006), providing unknown information (Lizzkowski, Carpenter, & Tomasello, 2009), and fulfilling others' desires (Repacholi & Gopnik, 1997). However, naive agents might not know what kind of help they need, or that they need help at all. For instance, your guest has no way of knowing that your faucet

is unusual until she attempts to turn it on. As a knowledgeable host, you can anticipate and pre-empt the trouble she might have by telling her how the faucet works in advance. As such, communicative prosociality often relies on our ability to proactively anticipate, select, and provide information others need before they even realize they needed it.

Our example above illustrates an important consideration in proactively informing others: what others are capable (or incapable) of achieving on their own. Though humans learn a great deal through self-guided exploration (e.g., Bonawitz et al., 2011; Schulz, 2012; Sim & Xu, 2014), truly naive agents (e.g. your guest in the bathroom) may tediously struggle through trial-and-error only to arrive at a poor solution. In these cases, learning from knowledgeable others is particularly valuable. They not only can provide new information, but also can prioritize information that would be too difficult, time-consuming, or perhaps even impossible for the learner to acquire by herself.

Prospectively reasoning about the information others need requires a sophisticated use of Theory of Mind; one must reason about their social partner's knowledge, goals, and competencies, to infer what would be useful for them to know (Strauss, Ziv, & Stein, 2002; Tomasello & Carpenter, 2007). Research suggests children readily consider others' mental states to selectively provide information that is useful for others. For instance, 5- to 7-year-olds adjust the amount of information they teach based on the learner's prior knowledge (Gweon, Shafto, & Schulz, 2014) and goals (Gweon, Chu, & Schulz, 2014), omitting unnecessary or irrelevant information to communicate just enough to support accurate inference. Children's tendency to communicate what is useful and relevant might reflect sensitivity to the value of information; their resistance to provide unnecessary information might reflect their desire to reduce the overall costs of information transfer.

In these studies, however, providing less information meant lower costs both for the children themselves as teachers (in terms of demonstrating the evidence) and for the learner (in terms of processing the evidence). It is thus unclear whether children were simply driven by the desire to reduce their own costs of teaching, or by the desire to reduce the learner's costs of learning. If the latter is true, children should be able to prioritize teaching information that decreases the learner's costs of learning even when children's own costs of teaching are equated.

Here, we show that children decide how to help others by considering the expected costs and rewards from the helpee's perspective. More specifically, we propose that children's choices of potential prosocial actions reflect the use of a *naive*

*utility calculus* (Jara-Ettinger, Gweon, Tenenbaum, & Schulz, 2015); children reason about others' costs and rewards and incorporate them into a single concept of utility to prioritize actions that have the highest net utility for others.

From offering physical aid to providing useful information, helping others can take various forms (Tomasello, 2009). Here, we focus on children's information sharing to test our hypothesis. Prior research on children's inferences and evaluations in pedagogical contexts suggests that children hold strong expectations about the information knowledgeable informants ought to select for naïve learners (e.g., Bonawitz et al., 2011; Gweon, Pelton, Konopka, & Schulz, 2014; Gweon, Shafto, & Schulz, 2014), and even act as helpful informants themselves (Gweon, Chu, & Schulz, 2014; Gweon, Shafto, & Schulz, 2014). Thus, children's information selection as teachers is an ideal case study for examining children's prosocial communicative decisions. In the following sections, we first describe our behavioral experiment, then present our theoretical predictions alongside their formal instantiation in computational models, and finally compare these predictions to our empirical data.

## Experiment

In the current study, children first learned about two novel, causal toys. They were then told that a naïve learner wanted to figure out both toys, and were asked which toy would be best to teach her if they had to choose just one. We systematically manipulated the relative difficulty for discovering how to activate the toys (costs) and the attractiveness of the toys' causal effects (rewards) across four experimental conditions. Critically, the costs for demonstrating (teaching) each toy were approximately the same, because once children knew how they worked, both toys were easy to activate. Additionally, unlike previous studies, children never actually demonstrated the toy they selected. We made these decisions in order to minimize the effect of children's own costs and rewards on their decisions.

We selected 5 – 7 years as our target age range for two reasons. First, children this age can consider others' mental states to decide what to communicate as teachers (Gweon, Shafto, & Schulz, 2014). Second, they can infer and integrate the expected costs and rewards of others' actions in their reasoning (see Jara-Ettinger, Gweon, et al., 2015). Thus, we expected that children this age might be capable of simulating another person's utility to make decisions on their behalf.

## Methods

**Participants** We recruited 126 5-, 6-, and 7-year-olds ( $M(SD) = 6.36(0.96)$  yrs, 54% female) from a local museum. An additional 7 children were excluded from analysis due to difficulty understanding English (5) and unknown date of birth (2). We randomly assigned children to the Play condition or one of four experimental (teach) conditions: (1) Easy-Dull v. Hard-Cool; (2) Easy-Cool v. Easy-Dull; (3) Easy-Dull v. Hard-Dull; (4) Easy-Cool v. Hard-Dull.

**Stimuli** We constructed four toys that differed in their relative

cost of discovery (the difficulty of spontaneously figuring out how to activate the toy; Easy vs. Hard), and the relative reward of their causal effect (its salience; Dull vs. Cool). Easy (low-cost) toys had a single button that generated an effect; figuring out these toys was nearly immediate (2s on average) and usually involved one action. Hard (high-cost) toys had one large button and 6 identical small buttons, 5 of which were inert and one that had to be pressed simultaneously with the large button to generate an effect; figuring out these toys took a long time (82s on average) and usually involved many failed actions. However, after discovery, both toys were quite easy to activate. Dull (low-reward) toys played music, and Cool (high-reward) toys had a plastic orb that lit up different colors and spun around. The names of the experimental conditions describe the two toys used. Figure 2d provides a schematic of the toys.

**Procedure** All conditions began with the *Discovery Phase*. The experimenter introduced the participant to two novel toys. She said that she did not know how the toys worked, and needed the child's help to figure them out. Children explored both toys until they activated them. For Hard toys, if children stopped exploring before successfully activating the toy, the experimenter encouraged them to continue with a series of prompts (e.g., "Hmm, I wonder what would happen if you pressed two buttons."). The experimenter delivered these prompts as spontaneous suggestions and with uncertainty to prevent children from thinking she knew how the 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).

Next came the *Choice Phase*. In the experimental conditions, the experimenter told children that she had a friend who knew nothing about the toys but was going to play with them later all by herself. The experimenter said that before her friend played alone with the toys, she would teach her friend how *just one* of the toys worked. Children were then asked: "How should I help her? Which toy should I teach her?" Children selected a toy and were asked to explain their choice. In the Play condition, the experimenter asked children to choose a toy with which they wanted to play (there was no mention of a friend), and allowed the child to play with it briefly.

Last was the *Manipulation-check Phase*. Here, children were asked (1) which toy was *harder* to figure out, and (2) which toy had a *cooler* effect, with the option of answering "same" for both questions.

## Computational Modeling and Predictions

To help us interpret our results, we implemented a set of formal computational accounts of our main and alternative hypotheses that capture how children should behave under different considerations. Following recent theoretical and empirical work on social reasoning via a *naïve utility calculus*, these models are based on the assumption that people reason about others in terms of their utilities (Jara-Ettinger, Gweon, et al., 2015; Jara-Ettinger, Tenenbaum, & Schulz, 2015). In

the context of our experiment, we assume that participants estimate each teaching plan’s utility, and then choose the plan with the highest net utility for the learner. Here, a *teaching plan* refers to the participant’s decision of which toy the experimenter should teach to a naïve learner; since the learner will play with both toys, deciding what is best to teach involves a consideration of both the toy that will be taught and the toy that the learner will figure out on her own.

A learner’s expected utility can be described as the difference between her expected rewards and costs. In our experiment, a helpful agent would maximize the learner’s overall utility by making a choice that maximizes her expected rewards (showing her a novel toy with high value) and minimizes her expected costs (showing her a novel toy that would require a lot of time and effort to figure out alone).

If children prioritize providing information that would maximize the learner’s expected utility, their teaching decisions should reflect the relative costs and rewards of the toys under consideration. More specifically, we test our predictions across four conditions. First, when one toy is both more rewarding than the other and expected to be more difficult (i.e., incur higher costs) for the learner, children should show a strong preference for teaching this toy and leave the other toy (low reward, low cost) for the learner to discover on her own (Easy-Dull v. Hard-Cool condition). When the two toys are equally costly to figure out but differ in reward, children should prefer the toy with a higher reward (Easy-Cool v. Easy-Dull). If the two toys are equally rewarding but differ in cost, children should choose the toy that is higher in cost (Easy-Dull v. Hard-Dull). Finally, when one toy is higher in cost and the other higher in reward, it is unclear which teaching decision would have the higher net utility for the learner, and so children might be split between the two toys (Easy-Cool v. Hard-Dull).

**Model Implementation** In computing the utilities, our main (full) model considered the learner’s expected utilities under each teaching plan as the linear sum of the following components: the cost the learner would incur and the reward she would obtain when activating the toy she was taught (activation cost,  $C_A$ , and activation reward,  $R_A$ , respectively), the expected cost the learner would incur to discover how the other toy works and the reward she would obtain upon discovery (discovery cost,  $C_D$  and discovery reward,  $R_D$ , respectively), and the reward the teacher would obtain from activating the toy she teaches (here we assume that the cost of teaching each toy is negligibly different since both toys are easy to activate, but the results are the same if this assumption is removed).

That is, given two toys, *Toy X* and *Toy Y*, our model compares two teaching plans:  $Teach_X$  (teach how *Toy X* works and leave *Toy Y* for the learner to figure out) and  $Teach_Y$  (teach how *Toy Y* works and leave *Toy X* for the learner), and computes the utility of each plan through the equations shown in Table 1.

The cost functions were determined by assuming that each button press results in a constant cost set to 1. (Using a different value does not change our results because all other parameters are defined in terms of this basic cost unit). Thus, a toy’s activation cost equals the number of button presses required to activate it. A toy’s expected discovery cost (as estimated by participants) was set to the expected number of buttons the learner would need to press to discover how to activate the toy, assuming that the learner would first try simple activation sequences (pressing one button at a time) and only try more complex and costly sequences (trying combinations of two buttons) once the simpler hypothesis space was depleted. Indeed, when participants interacted with the high-cost toys, most children began by pressing each button once and only attempted combinations of two buttons once they had pressed each button individually.

The toys’ rewards were treated as variables, allowing for the possibility that different participants assign different rewards to each toy. The music reward was set within a range that captured rewards large enough that could surmount the expected costs of discovery of either toy, as well as rewards that were too small for a teacher to deem the effect worthy of discovery (from 1 to 45; as the expected discovery cost of the more costly toy is  $\sim 22$ ). Our model results are robust to different reward ranges). For each possible music reward, we set a uniform distribution over the range of rewards for the light-up toy. This range was selected so that the probability that the lights’ reward surpasses the music’s reward matched the empirical data obtained from children’s choice of which effect was “cooler” in the *Manipulation-check Phase*.

**Alternative Models** While we assume that children consider both the expected costs and rewards of each teaching plan (which includes both the toy that will be taught and the toy that will be left untaught), 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, children might only be concerned about the relative difficulty of the toys, and so only consider the costs the learner will incur when discovering the untaught

	Activation reward	Activation cost	Discovery reward	Discovery cost	Teacher reward
$U(Teach_X) =$	$+R_A(Toy_X)$	$-C_A(Toy_X)$	$+R_D(Toy_Y)$	$-C_D(Toy_Y)$	$+R_A(Toy_X)$
$U(Teach_Y) =$	$+R_A(Toy_Y)$	$-C_A(Toy_Y)$	$+R_D(Toy_X)$	$-C_D(Toy_X)$	$+R_A(Toy_Y)$

Table 1:  $R_A$  and  $R_D$  determine the toy’s positive reward from activation and discovery, respectively (which are the same),  $C_A$  determines the negative activation cost for a knowledgeable agent, and  $C_D$  determines the negative discovery cost for a naïve agent.

toy, and when activating the taught toy (costs-only model). Alternatively, once children know how each of the toys work, they might only focus on teaching what is cool and just 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 trying to activate the untaught toy (taught-only model). Finally, children might take what is being taught for granted and focus instead on the costs and rewards of having to figure out the other toy (untaught-only model). Table 2 summarizes the different considerations of the full and alternative models. We can assess the relative performance of these different models by comparing how well each model’s predictions correspond to children’s choices.

		Model Parameters				
		Activation reward	Activation cost	Discovery reward	Discovery cost	Teacher reward
Models	Full Model	✓	✓	✓	✓	✓
	Costs-only	✗	✓	✗	✓	✗
	Rewards-only	✓	✗	✓	✗	✓
	Taught-only	✓	✓	✗	✗	✓
	Untaught-only	✗	✗	✓	✓	✗

Table 2: Model space. Each model captures different considerations that participants can use to decide what to teach.

## Results and Discussion

In the *Manipulation-check Phase*, most children thought Hard toys were harder than Easy toys, Cool toys were cooler than Dull toys, and that they were equally hard or cool when these dimensions were held constant; this confirmed that our manipulation of the toys’ relative costs and rewards was successful. However, compared to the difficulty manipulation, there was less agreement about the toys’ relative “coolness” (i.e., some children thought the music was cooler than the lights). Figure 1 provides a summary of children’s responses.

Children’s choices of which toy was best to teach varied systematically with respect to the relative costs and rewards of the toys. When one toy was higher in both discovery cost and reward (Easy-Dull v. Hard-Cool), children strongly preferred the experimenter teach that toy (21/25; two-tailed binomial test,  $p = 0.001$ ). When the toys differed only in reward (Easy-Cool v. Easy-Dull), children showed a mild preference for the experimenter to teach the cooler toy (18/26; two-tailed binomial test,  $p = 0.076$ ).<sup>1</sup> When the toys differed only in

<sup>1</sup>Children’s explanations for why they selected the Easy-dull toy

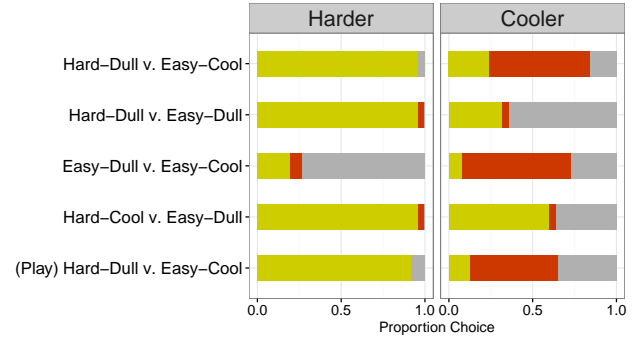


Figure 1: The proportion of children who selected one toy as harder (left panel) and one toy as cooler (right panel) in each condition. Order of toy names per condition reflects order of bars. ‘Same’ judgments are in grey.

discovery cost (Easy-Dull v. Hard-Dull), children selected the harder toy (21/25; two-tailed binomial test,  $p = 0.001$ ), and finally, when cost and reward were pitted against each other (Easy-Cool v. Hard-Dull), children’s choices were split (10/25; two-tailed binomial test,  $p = 0.424$ ).

One might think children were at chance in the Easy-Cool v. Hard-Dull condition not because one toy was higher in rewards and the other higher in cost, but rather because the two toys were equally rewarding. Perhaps, adding buttons to the “Dull” toy made it more interesting, leading more children to choose the Hard-Dull toy. To address this possibility, we compared the results from this condition to the Play condition. In this condition, children overwhelmingly preferred to play with the Easy-Cool toy (20/25; two-tailed binomial test,  $p = 0.004$ ), showing a significant difference from their tendency to teach this toy in the Easy-Cool v. Hard-Dull experimental condition (two-tailed Fisher’s Exact test,  $p = 0.009$ ). These results confirm that children’s responses in the Easy-Cool v. Hard-Dull condition were driven by a consideration of both the expected rewards and costs for the learner.

The pattern of children’s choices across conditions suggests they were indeed reasoning about the relative expected costs and rewards of figuring out how different toys worked for a naïve learner and considered these factors when deciding what should be taught. Figure 2a provides a summary of children’s responses in the Choice-phase.

**Model Fit** To evaluate model performance we calculated the likelihood of each model generating the empirical data. Since our models are idealized formalizations, we added a small probability  $\alpha \in (0, 1)$  that any given participant may get distracted and respond at random. Thus, for each model we assumed that children sample a response from the model’s distribution with probability  $1 - \alpha$  and that they give a random response with probability  $\alpha$ . Here we present results with the noise parameter  $\alpha = 0.1$  (see Model Implementa-

revealed that many considered the music (“Dull” effect) cooler than the lights (“Cool” effect), indicating that they were also choosing the toy with the higher expected rewards.

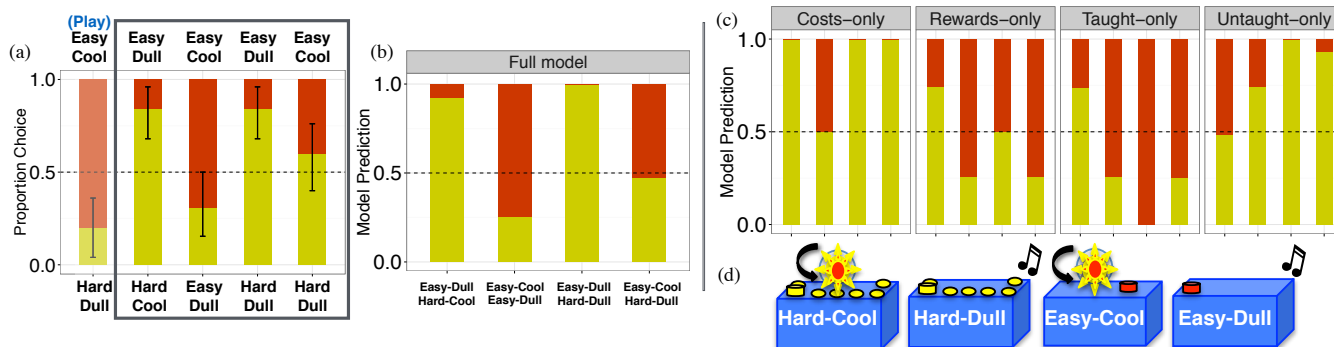


Figure 2: (a) Proportion of children selecting each toy in the *choice phase* (teaching conditions outlined in black and play condition indicated in blue). Error bars are bootstrapped 95% CI. (b) The full model predictions. (c) The predictions of (from left to right): Costs-only, Rewards-only, Taught-only, and Untaught-only models (d) Schematic of the four toys used varying in cost (Easy/Hard) and reward (Cool/Dull). Condition names reflect the two toys used.

tion section). Our qualitative conclusions, however, hold for a wide range of noise and persistence parameters (see <https://osf.io/p29zr/> for model code and predictions).

The full model had the highest likelihood of generating the observed data across conditions. After computing each model’s likelihood, we computed the ratio between the full model’s likelihood and each alternative model’s likelihood. Thus, 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. All model comparisons favored the full model by at least 3 orders of magnitude (all likelihood ratios  $> 1e4$ ). In short, our empirical data provided compelling support for the full model over the four alternative models.

To better understand why the data are strong evidence for the full model, we consider the alternatives. The alternative models failed to predict participant performance in various conditions. Here we focus on their most salient limitations, which shed light on why the full model captures participants’ judgments.

First, the *costs-only* model failed to capture data in the Easy-Cool v. Hard-Dull condition. According to this model, participants should always select whichever toy is most difficult for the learner to discover alone. If this were true, then participants should have preferred to teach the high-cost toy, but instead, their choices were split, as the full model predicted. The *rewards-only* model’s critical failure was in the Easy-Dull v. Hard-Dull condition, where it predicted that participants should have no preference, when participants in fact preferred to teach the high-cost toy. The *taught-only* model also failed to predict the Easy-Dull v. Hard-Dull condition, as it predicted that the easier toy should be taught (even for a knowledgeable agent, the low-cost toy is slightly easier to activate than the high-cost toy). Last, the *untaught-only* model most notably failed to predict responses in the Easy-Cool v. Easy-Dull condition. Because the costs were matched, this model predicted that children should teach the less rewarding

toy (so as to let the learner discover the high-reward alone). Instead, participants preferred to teach the more rewarding toy, as the full model predicted. Altogether, the full model was the only model that accurately predicted responses in all four teaching conditions. Figure 2b-c provides a summary of each model’s predictions.

## General Discussion

We propose that children consider others’ expected rewards and costs of learning when deciding what to helpfully communicate to a naïve learner, in a way consistent with a *naïve utility calculus*. Children’s behavior in our teaching conditions was unambiguously most consistent with a model that took into account the costs and rewards of both the taught and untaught toys, suggesting children think about what the learner ought to know and what would be more difficult for her to discover on her own. When the toys only differed in value, children selected to teach the cooler toy, maximizing the utility for the learner by increasing her expected rewards. Similarly, when the toys only differed in discovery cost, children selected the harder toy, maximizing the utility for the learner by decreasing her expected costs. When cost and reward were pitted against each other, children were split, suggesting a desire to simultaneously increase the learner’s rewards (teaching the cooler toy) and reduce her costs (teaching the harder toy).

It is possible that children who taught the Hard toys did so not because the learner would find them more difficult but simply because children received help on these toys (i.e., experimenter prompts). However, we think this is unlikely. If children were using this simpler strategy, they should prefer the Hard toy in all conditions where it was used, but in the Easy-Cool v. Hard-Dull condition, children were at chance. Children’s explanations provide additional support that the learner’s expected costs and rewards were indeed informing their choices: Children spontaneously appealed to differences in cost and reward, and their tendency to do so varied appro-

priately with the dimensions of contrast across the particular toys used. For example, they appealed more to reward in the Easy-Cool v. Easy-Dull condition but more to cost in the Easy-Dull v. Hard-Dull condition. In sum children in our experiment did appear to reason about the relative utilities of teaching one toy or the other and selected the teaching plan with the higher overall utility for the learner.

Simulating an unknown person's expected costs and rewards across two goals (toys) and integrating these considerations to determine the best teaching plan requires complicated perspective taking with many unobservable variables. We did not find any age differences in our experiment, but it is possible that younger children would perform differently on this task. Perhaps, a consideration of reward comes online before a consideration of cost, or perhaps the ability to integrate these variables improves across development. Comparing younger children's teaching decisions to older children's and the predictions of our different models would shed light on how a *naïve utility calculus* might develop.

Additionally, it is an open question how much of children's ability to reason about the rewards and costs of another person's exploration depended on their own experience interacting with the toys. Even adults struggle to take the perspective of others who know less or believe something different than they do (Sommerville, Bernstein, & Meltzoff, 2013), so one might expect that after successfully activating the toys, children would fail to recognize that a toy would be hard for a naïve person to figure out. We are currently examining whether children's teaching selections differ when they are explicitly taught how the toys work rather than discover how to activate them through play.

In our experiment and model, the teacher is assumed to take on the learner's utility function and to select what is best to teach based on the expected costs and rewards for the learner. However, teaching is a dyadic situation wherein the teacher's own costs and rewards are presumably also a consideration. Critically in our experiment, we aimed to see if children understood that teaching can reduce the costs and increase the rewards of learning for another person. In order to examine this, we controlled the potential costs and rewards of teaching by making the toys, once known, simple to activate and having children make decisions about what someone else should teach, rather than act as teachers themselves. Exploring children's behavior as actual teachers, as well as situations where certain toys are more rewarding or costly to teach will give us a deeper understanding of how we reason about the utility of communicating certain things over others in real-world social interactions.

Deciding how to act prosocially on behalf of others is no trivial feat. By considering the expected costs and benefits for others of the aid we could provide, we can constrain our decisions of how best to help and inform others. While learning from knowledgeable teachers is useful for acquiring accurate beliefs about the world, not all knowledge is equally useful; some information leads to higher rewards (e.g., knowing

where to forage for food), and some leads us to bypass significant costs or even risks (e.g., learning how to hide from predators). The results of the present study provide compelling evidence that children are able to reason about the expected rewards and costs of learning from the perspective of the learner, and that these considerations inform their decisions of what is most useful to communicate.

## Acknowledgments

We thank Avani Singh for help with data coding, and Ellen Markman, Noah Goodman, and Josh Tenenbaum for their helpful feedback. We also thank the Palo Alto Junior Museum & Zoo and the families who participated. This work was supported by the Varieties of Understanding grant from the John Templeton Foundation to H.G. and a NSF Graduate Research Fellowship to S.B.

## References

- Bonawitz, E., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. E. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition, 120*.
- Fehr, E., & Fischbacher, U. (2003). The nature of human altruism. *Nature, 425*.
- Gweon, H., Chu, V., & Schulz, L. E. (2014). To give a fish or to teach how to fish? children weigh costs and benefits in considering what information to transmit. In *Proceedings of the 36th annual conference of the cognitive science society*.
- Gweon, H., Pelton, H., Konopka, J., & Schulz, L. E. (2014). Sins of omission: Children selectively explore when agents fail to tell the whole truth. *Cognition, 132*.
- Gweon, H., Shafto, P., & Schulz, L. E. (2014). Children consider prior knowledge and the cost of information both in learning from and teaching others. *Proceedings of the 36th Annual Conference of the Cognitive Science Society*.
- Jara-Ettinger, J., Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2015). Children's understanding of the costs and rewards underlying rational action. *Cognition, 140*.
- Jara-Ettinger, J., Tenenbaum, J. B., & Schulz, L. E. (2015). Not so innocent: Toddlers inferences about costs and culpability. *Psychological science, 26*(5).
- Liszkowski, U., Carpenter, M., & Tomasello, M. (2009). Twelve-month-olds communicate helpfully and appropriately for knowledgeable and ignorant partners. *Cognition, 108*.
- Repacholi, B. M., & Gopnik, A. (1997). Early reasoning about desires: Evidence from 14- and 18-month-olds. *Developmental Psychology, 33*.
- Schulz, L. (2012). The origins of inquiry: Inductive inference and exploration in early childhood. *Trends in Cognitive Sciences, 16*.
- Sim, Z., & Xu, F. (2014). Acquiring inductive constraints from self-generated evidence. *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*.
- Sommerville, J. A., Bernstein, D. M., & Meltzoff, A. N. (2013). Measuring beliefs in centimeters: Private knowledge biases preschoolers' and adults' representation of others' beliefs. *Child Development, 84*.
- Strauss, S., Ziv, M., & Stein, A. (2002). Teaching as a natural cognition and its relations to preschoolers developing theory of mind. *Cognitive Development, 17*.
- Tomasello, M. (2009). *Why we cooperate*. MIT Press.
- Tomasello, M., & Carpenter, M. (2007). Shared intentionality. *Developmental Science, 10*.
- Warneken, F., & Tomasello, M. (2006). Altruistic helping in human infants and young chimpanzees. *Science, 311*.