REPORT

COGNITIVE DEVELOPMENT

Ten-month-old infants infer the value of goals from the costs of actions

Shari Liu,[1,3]* Tomer D. Ullman,[1,2,3] Joshua B. Tenenbaum,[2,3] Elizabeth S. Spelke[1,3]

Infants understand that people pursue goals, but how do they learn which goals people prefer? We tested whether infants solve this problem by inverting a mental model of action planning, trading off the costs of acting against the rewards actions bring. After seeing an agent attain two goals equally often at varying costs, infants expected the agent to prefer the goal it attained through costlier actions. These expectations held across three experiments that conveyed cost through different physical path features (height, width, and incline angle), suggesting that an abstract variable—such as "force," "work," or "effort"—supported infants' inferences. We modeled infants' expectations as Bayesian inferences over utility-theoretic calculations, providing a bridge to recent quantitative accounts of action understanding in older children and adults.

When we observe people's actions, we see more than bodies moving in space. A hand reaching for an apple is not just one object decreasing its distance from another; it can indicate hunger (in the person who is reaching), helpfulness (if the person is reaching on behalf of someone else), or compromise (if the person reaching would prefer a banana, but not enough to go buy one). This fast and automatic ability to interpret the behavior of others as intentional, goal-directed, and constrained by the physical environment is often termed "intuitive psychology" (1–4). We used behavioral experiments and computational models to probe the developmental origins and nature of this ability.

Over the past two decades, research has revealed that the building blocks of our intuitive psychology are present as early as the first year of life. Despite infants' limited experience, their interpretations of other people's actions are guided by assumptions about agents' physical properties (5), intentions and goals (6), mental states (7–10), causal powers (11), and dispositions to act efficiently (7, 12, 13). This wealth of findings does not reveal, however, whether infants' capacities depend on a host of distinct local abilities (14–16) or on a single coherent system supporting inference, prediction, and learning (3, 17–19).

We tackled this question in a case study, based on a computationally precise proposal for a coherent, abstract, and productive system for action understanding (Fig. 1). Previous studies suggest that infants are sensitive to the costs of agents' actions (3, 7, 12, 13) and can infer agents' preferences (6, 9). Decision theorists for hundreds of years have recognized these as the two central factors guiding the decisions of rational agents (20–22). We asked whether infants can integrate these dimensions to infer agents' goals: Do infants use the cost that an agent expends to attain a goal state in order to infer the value of that goal state for the agent?

Such an inference has been proposed to rest on three nested assumptions that together constitute a "naïve utility calculus" (23), analogous to classical economic thinking. First, agents act

[1]Department of Psychology, Harvard University, Cambridge, MA, USA. [2]Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology (MIT), Cambridge, MA, USA. [3]Center for Brains, Minds, and Machines (CBMM), McGovern Institute for Brain Research, MIT, Cambridge, MA, USA.
*Corresponding author. Email: shariliu01@g.harvard.edu

Fig. 1. A schematic of our computational model. (**A**) The forward direction defines the agent as a rational planner that calculates the utilities of different actions from their respective costs and rewards and then selects an action stochastically in proportion to its utility. In this case, the overall utility for approaching the triangle is higher than for approaching the square, so the central agent (circle) will likely choose triangle over square. (**B**) An observer (i) assuming this model and priors over the costs of different actions can (ii) observe a series of actions and then (iii) infer a posterior distribution over the hidden values of an agent's costs and rewards given its actions. (iv) These posteriors can then be used to predict the actions of the agent in a new situation, in which the same goal states can be reached with different actions.

to maximize their utility U, under constraints (*2*, *4*, *24*, *25*). Second, this utility separates into rewards and costs, two distinct components that can be individual targets of inference (*26*). That is, if $R(S)$ is the reward of a goal state S, and $C(A)$ is the cost of an action, then an agent acts to maximize the following

$$U(A, S) = R(S) - C(A) \qquad (1)$$

Third, the cost of an action is not arbitrary but depends on properties of both the agent and the situation: properties that jointly determine how much effort the agent might need to exert to carry out that action.

These assumptions can be formalized as generative models that successfully predict the quantitative and qualitative behavior of adults and older children (*4*, *23*, *26*). In these models, observers who reason that other agents are maximizing their expected utility according to Eq. 1 can use what they know about rewards and costs to predict the agents' future actions. Inverting this process, observers can use the agents' overt actions to infer their hidden rewards and costs, according to

$$P(R, C \mid A) \propto P(A \mid C, R) \cdot P(R, C) \qquad (2)$$

where $P(R, C \mid A)$ is the posterior distribution over the rewards and costs of an agent. By Bayes' theorem, this distribution is proportional to the product of $P(A \mid C, R)$—the likelihood of the agent choosing action A given rewards R and costs C, given by a rational planning procedure (*4*, *23*)—and $P(R, C)$, a prior distribution over costs and rewards.

Do infants apply the logic of cost-reward reasoning? Past research suggests that infants are sensitive to the relative value of different goal objects for an agent who chooses to approach one object over another (*6*, *27*) as well as the relative efficiency of the actions taken by an agent who approaches a goal object (*12*, *13*, *28*). Past studies do not reveal, however, whether infants have a unified intuitive psychology in the form of a generative model, or separate representations for variables such as cost and reward that become unified only later in development, as children gain experience exerting themselves to achieve goals or communicating with others about their desires and actions. It is also an open question whether infants consider cost and reward in terms of abstract variables—such as work, effort, desire, or value—or whether their understanding is restricted to perceptual features of actions, such as the distance or duration an agent travels or the number of times it selects a particular goal. In physical action contexts, effort often covaries with perceptible properties such as the length or duration of a path traveled, but effort depends ultimately on the amount of force that the agent must exert over time and distance (the amount of work the agent must do). Likewise, value often covaries with the number of times an agent selects a goal but ultimately depends on how strongly the agent desires a goal relative to the cost of achieving it or its value relative to other options.

We designed and conducted three experiments to test whether infants learn about the reward agents place on goals from cost, working backward from the assumption that agents maximize utility and inferring relative rewards from observed actions under varying costs. We then used the data from these experiments, together with the findings from past experiments (*6*, *7*, *13*), to test a variety of computational models of infants' performance, including models with integrated versus isolated, and abstract versus cue-based, representations of costs and rewards (model description is provided in the supplementary materials). Our empirical and computational findings support the view that a productive system grounded in cost-reward trade-offs guides action understanding toward the end of the first year of life.

We tested $n = 80$ 10-month-old infants in three experiments with prespecified designs, procedures, sample sizes, and analysis plans (*29*). In all experiments, infants first saw an agent move to and refuse to move to each of two target goals under conditions of varying cost. Then, infants watched test events in which the agent chose either the higher- or the lower-value target when both were present at equal cost. If infants infer the reward of the targets to the agent from the effort undertaken to reach them, then they should be more surprised when the agent chooses the lower-value target, looking longer at the test trials displaying that action (*30*).

In experiment 1 ($n = 24$ infants), we leveraged events widely used in studies of early action understanding, in which animated characters jump efficiently over barriers of variable heights to arrive at goal objects (*3*, *7*, *13*, *31*) and indicate their preferences by selecting one goal over another (*6*, *9*). During familiarization, infants watched six trials that consisted of four different events involving a central agent and one of two target individuals on a level surface (Fig. 2A and movie S1). In each event, the target jumped and made a noise, and the agent responded by turning to face and beginning to approach the target, whereupon a barrier fell onto the stage
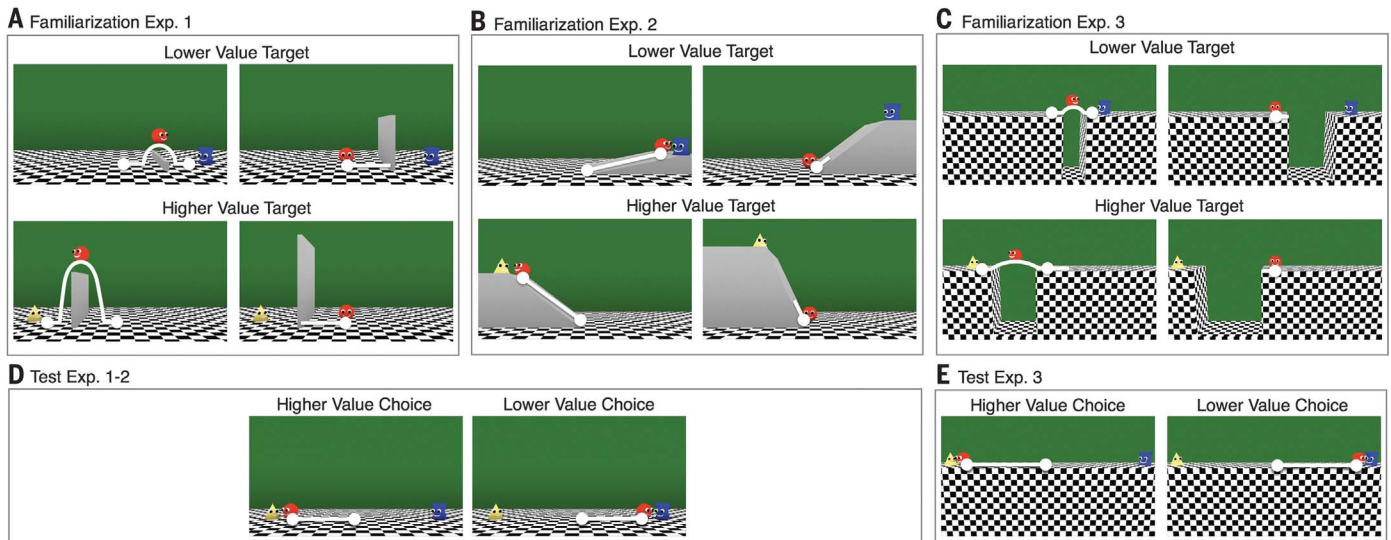
Fig. 2. Structure of experiments 1 through 3. (**A** to **C**) During familiarization, the central agent (circle) accepted a low cost and refused a medium cost for the lower-value target (square) and accepted a medium cost and refused a high cost for the higher-value target (triangle). Other than the sizes of the barriers, ramps, and trenches, and the consequent trajectories of motion, the pairs of events displaying approach or refusal of approach to the two targets were identical. (**D** and **E**) At test, the agent looked at each of the two targets and chose either the lower- or higher-value target. White circles indicate start- and end points of action, and white lines indicate trajectories.

Fig. 3. Boxplots of average looking time toward the higher- and lower-value choice during test in experiments 1 through 3. White diamonds indicate means, with error bars indicating within-subjects standard errors. Horizontal lines indicate medians, boxes indicate middle quartiles, and whiskers indicate points within 1.5 times the interquartile range from the upper and lower edges of the middle quartiles. Light gray points connected across boxes indicate looking times from individual participants. Beta coefficients indicate effect sizes in standard deviations, and asterisks indicate significance relative to prespecified (experiments 1 and 2) and preregistered (experiment 3) alphas (*P < 0.05). Statistical analyses are provided in the text and supplementary materials.

directly in the agent's path. On two of these events (one for each target), the agent looked to the top of the barrier, made a positive "Mmmm!" sound, backed up, and then jumped over the barrier, landing next to the target. On the other two events, the agent looked to the top of the barrier, made a neutral "Hmmm…" sound, backed away, and returned to its initial position. The critical distinction between these events concerned the height of the barrier and therefore the length, height, and speed of the jump that the agent undertook so as to clear it (all jumps were equated for duration). For one target, the agent jumped over a low barrier and declined to jump a medium barrier; for the other target, the agent jumped the medium barrier and declined a tall barrier. After this familiarization, the agent appeared between the two equidistant targets on a level surface. Infants viewed two pairs of looped test events (Fig. 2D and movie S4), order counterbalanced, in which the agent looked at each of the targets and then repeatedly approached either the higher- or the lower-value target. Our prespecified dependent measure was average log-transformed looking time (32) across test trials. In experiment 1, we predicted differential looking at the test events but did not prespecify the direction of this difference.

Infants looked longer at test trials in which the agent chose the target for whom it had jumped a lower barrier (mean = 28.41 s, SD = 14.85), relative to the target for whom it had jumped a higher barrier (mean = 21.79 s, SD = 12.29) [95% confidence interval (CI) (0.062, 0.591), b coefficient (B) = 0.327, SE = 0.130, standardized β coefficient (β) = 0.502, $t(24)$ = 2.523, P = 0.019, two-tailed, mixed effects model with random intercept for participant] (Fig. 3) (30). These findings suggest that infants inferred the rewards that the central agent placed over the targets from the cost the agent was willing to expend to reach these targets, and they therefore expected the agent to choose that target at test. Nevertheless, experiment 1 does not show whether infants used the physical effort undertaken by the agent, or variables that merely correlate with effort (such as distance or speed), in their predictions.

To control for distance and speed of travel, experiment 2 (n = 24 infants) used ramps of three different incline angles to convey cost (Fig. 2B and movie S2). On each familiarization trial, a target appeared on the top of one ramp, and the agent looked up the ramp and either climbed to the target or returned to its starting position. The agent climbed the shallow ramp and declined to climb the medium ramp for one tar-

get and climbed the medium ramp and declined the steep ramp for the other target. The methods were otherwise the same as in experiment 1. Consistent with our prespecified directional prediction, infants again looked longer at the test events in which the agent approached the lower-value target (mean = 30.94 s, SD = 13.31) than test events in which the agent approached the higher-value target (mean = 27.05 s, SD = 17.55) [95% CI (0.028, 0.472), B = 0.250, SE = 0.109, β = 0.408, $t(24)$ = 2.294, P = 0.015, one-tailed, mixed effects model with random intercept for participant] (Fig. 3) (30). This finding further suggests that infants understand agents' actions in accord with abstract, general, and interconnected concepts of cost and reward, but narrower explanations remain. In experiments 1 and 2, the agent was confronted with an obstacle to its forward motion (a barrier or ramp), and the size of the obstacle covaried with the cost of the agent's action, requiring the agent to move further upward to attain the higher-value target. Because infants become sensitive to the effects of gravity on objects on inclined planes well before 10 months of age (33), they may learn that agents will move to greater heights or overcome higher obstacles for more rewarding targets, without invoking a more abstract representation of physical effort. Experiment 3 was undertaken to explore these interpretations.

In experiment 3 (n = 32 infants), the agent was separated from each of the two targets during familiarization not by an obstacle but by a horizontal gap in the supporting surface (Fig. 2C and movie S3). Infants first saw a ball roll off the edge of a narrow, medium, and wide gap and shatter (movie S6). During familiarization, these three trenches, requiring jumps of variable lengths and speeds but of equal durations and heights, were interposed between the agent and target; the agent moved to the edge of a trench, looked at the far side, and then jumped over a narrow trench for one target (and refused the medium trench) and a medium trench for the other target (and refused the widest trench). The methods were otherwise unchanged (Fig. 2E and movie S5). The methods and analyses for experiment 3 were preregistered at https://osf.io/k7yjt (29) and tested the same directional prediction as that in experiment 2. Infants again looked longer at the lower-value choice (mean = 23.05 s, SD = 13.58) relative to the higher-value choice (mean = 17.47 s, SD = 10.69) [95% CI (0.020, 0.501), B = 0.260, SE = 0.119, β = 0.403, $t(32)$ = 2.185, P = 0.018, one-tailed, mixed effects model with random intercept for participant] (Fig. 3) (30).

Regardless of whether an agent cleared higher barriers (experiment 1), climbed steeper ramps (experiment 2), or jumped wider gaps (experiment 3) for one target over the other, infants expected the agent to choose that target at test. Across all experiments, infants looked longer at the lower-value action (mean = 26.99 s, SD = 14.13) than the higher-value action (mean = 21.64 s, SD = 13.94) [95% CI (0.139, 0.415), B = 0.277, SE = 0.070, β = 0.424, $t(80)$ = 3.975, P < 0.001, one-tailed, mixed effects model with random

intercepts for participant and experiment], supporting our general hypothesis that infants infer the values of agents' goals from the costs of their actions. Although past research had shown that infants represent the goal of an agent's action from observations of an agent's choices between two objects (6) and expect agents to give different emotional responses when agents complete versus fail to complete their goals (31), the present experiments provide evidence that infants develop ordinal representations of reward even when the number of choices and expressed emotions are equated across the actions and only the costs of the actions vary. Moreover, they show that infants do not simply attribute higher reward to goals that agents pursue for a longer duration or attain with greater frequency because these variables were equated as well. The findings also provide evidence for longstanding suggestions that infants represent physical cost as a continuous variable that agents seek to minimize (3, 13): Infants make appropriate cost assessments even when the specific physical features that distinguished lower- from higher-cost actions—including the relative length, curvature, duration, or speed of motion trajectories—systematically varied. Together, experiments 1 through 3 suggest that infants represent cost and reward as interconnected, abstract variables that they apply to a wide range of events.

The discovery that infants infer the rewards of goals from the costs of achieving them provides empirical support for the thesis that an abstract and productive system guides infants' analysis of agents and their actions (3, 17, 19). Specifically, we suggest that the cognitive machinery supporting infants' intuitive psychology includes a mental model both of how agents plan actions in the forward direction, in accord with maximizing their utilities (Eq. 1 and Fig. 1A) (23), and a procedure for inverting this model, in accord with the computational framework of inverse planning (Eq. 2 and Fig. 1B) (4). Applying this general framework to our specific experiments, we posit that infants have developed a model of action planning before the experiment: They assume that agents value some goal objects more than others and that agents engage in costlier actions to achieve goals with higher reward. When infants see the agent take costlier actions to arrive at one target than to arrive at another, infants invert this model to infer the relative reward of the two targets to that agent. Then, when infants see the agent flanked by the two targets in a situation in which costs are equal, they apply their knowledge of the targets' relative value to the agent to run their planning model for that agent forward, predicting the target that it will approach. We have implemented this hypothesis in a computational model that accounts not only for the findings of the present experiments but also for a range of past studies of early action understanding (6, 7, 13). Furthermore, we compared this model with an array of simpler models that focus only on relative costs

or rewards in isolation, or on particular cues to effort or value. We found that only the full model with abstract variables for costs and rewards can account for all of the findings (fig. S3 and supplementary materials).

The present studies raise key questions for future research. First, the cognitive architecture underlying infants' assessment of cost remains to be explored. Our experiments suggest that infants are responding to an abstract notion of cost, rather than specific physical path features such as vertical motion (controlled for in experiment 3), horizontal motion (controlled for in experiment 1), or raw path length (controlled for in experiment 2). We do not know, however, whether infants represent the abstract costs of actions by drawing on a concept of experienced effort or exertion within the domain of naïve psychology, or by leveraging an intuitive concept of force or work done (the integral of force applied over a path) from the domain of naïve physics (34, 35), or perhaps both. Also, our experiments investigated only one class of goal states and target-directed actions, leaving open the breadth and generality of infants' intuitive psychology. In particular, cost can be defined in terms of work or effort to produce physical forces, but there are other kinds of costs: Agents could consider variables such as the mental effort of planning (36, 37) and the risks of choosing certain actions, neither of which involves applications of force. It is an open question whether these other variables trade off against reward in infants' intuitive psychology the way that physical work or effort does. Last, our studies do not speak to the origins of these abilities. Although 10-month-old infants cannot perform the actions from our experiments or communicate with others about them, their productive system for reasoning about costs and rewards may arise through their experiences observing the actions of other agents or performing actions within their repertoire, such as lifting their arms or balancing their bodies against the force of gravity. Alternatively, this system of intuitive psychology may guide infants' action understanding from the beginning. Testing these possibilities would address fundamental questions concerning the nature, origins, and interrelations between our intuitive psychology and intuitive physics.

However these questions are answered, the present study suggests that our propensity to understand the minds and actions of others in terms of abstract, general, and interrelated concepts begins early. Before human infants learn to walk, leap, and climb, they leverage mental models of agents and actions: forward models of how agents plan, and inverse models for working backward from agents' actions to the causes inside their minds.

REFERENCES AND NOTES

1. F. Heider, M. Simmel, *Am. J. Psychol.* **57**, 243–259 (1944).
2. D. C. Dennett, *The Intentional Stance* (The MIT Press, 1987).
3. G. Gergely, G. Csibra, *Trends Cogn. Sci.* **7**, 287–292 (2003).
4. C. L. Baker, R. Saxe, J. B. Tenenbaum, *Cognition* **113**, 329–349 (2009).
5. R. Saxe, T. Tzelnic, S. Carey, *Cognition* **101**, B1–B8 (2006).
6. A. L. Woodward, *Cognition* **69**, 1–34 (1998).
7. G. Gergely, Z. Nádasdy, G. Csibra, S. Bíró, *Cognition* **56**, 165–193 (1995).
8. K. H. Onishi, R. Baillargeon, *Science* **308**, 255–258 (2005).
9. Y. Luo, R. Baillargeon, *Cognition* **105**, 489–512 (2007).
10. Y. Luo, S. C. Johnson, *Dev. Sci.* **12**, 142–149 (2009).
11. P. Muentener, S. Carey, *Cognit. Psychol.* **61**, 63–86 (2010).
12. A. E. Skerry, S. E. Carey, E. S. Spelke, *Proc. Natl. Acad. Sci. U.S.A.* **110**, 18728–18733 (2013).
13. S. Liu, E. S. Spelke, *Cognition* **160**, 35–42 (2017).
14. A. L. Woodward, *Curr. Dir. Psychol. Sci.* **18**, 53–57 (2009).
15. M. Paulus, *Psychol. Res.* **76**, 760–767 (2012).
16. C. M. Heyes, C. D. Frith, *Science* **344**, 1243091 (2014).
17. E. S. Spelke, K. D. Kinzler, *Dev. Sci.* **10**, 89–96 (2007).
18. J. B. Tenenbaum, C. Kemp, T. L. Griffiths, N. D. Goodman, *Science* **331**, 1279–1285 (2011).
19. R. Baillargeon, R. M. Scott, L. Bian, *Annu. Rev. Psychol.* **67**, 159–186 (2016).
20. D. Bernoulli, *Econometrica* **22**, 23–36 (1954).
21. J. S. Mill, *Utilitarianism* (Oxford University, 1863).
22. J. Bentham, *An Introduction to the Principles of Morals and Legislation* (Clarendon Press, 1879).
23. J. Jara-Ettinger, H. Gweon, L. E. Schulz, J. B. Tenenbaum, *Trends Cogn. Sci.* **20**, 589–604 (2016).
24. C. G. Lucas et al., *PLOS ONE* **9**, e92160 (2014).
25. A. Jern, C. Kemp, *Cognition* **142**, 12–38 (2015).
26. J. Jara-Ettinger, H. Gweon, J. B. Tenenbaum, L. E. Schulz, *Cognition* **140**, 14–23 (2015).
27. Y. Luo, R. Baillargeon, *Psychol. Sci.* **16**, 601–608 (2005).
28. G. Csibra, G. Gergely, S. Bíró, O. Koós, M. Brockbank, *Cognition* **72**, 237–267 (1999).
29. Materials and methods are available as supplementary materials.
30. Before conducting experiment 1, we were unsure whether infants would express their expectations by looking at the unexpected event for longer or shorter times, and so we prespecified a two-tailed alpha. After experiment 1, however, we had strong reason to believe that infants would look longer at the test actions that they found less probable, and so we prespecified (for experiment 2) and preregistered (for experiment 3) one-tailed alphas and report these numbers in the text. However, results of all three experiments were also statistically significant using two-tailed significance thresholds, with $P = 0.031$ and $P = 0.036$ for experiments 2 and 3, respectively.
31. A. E. Skerry, E. S. Spelke, *Cognition* **130**, 204–216 (2014).
32. G. Csibra, M. Hernik, O. Mascaro, D. Tatone, M. Lengyel, *Dev. Psychol.* **52**, 521–536 (2016).
33. I. K. Kim, E. S. Spelke, *J. Exp. Psychol. Hum. Percept. Perform.* **18**, 385–393 (1992).
34. E. Téglás et al., *Science* **332**, 1054–1059 (2011).
35. T. D. Ullman, E. Spelke, P. Battaglia, J. B. Tenenbaum, *Trends Cogn. Sci.* **21**, 649–665 (2017).
36. A. Shenhav et al., *Annu. Rev. Neurosci.* **40**, 99–124 (2017).
37. W. Kool, S. J. Gershman, F. A. Cushman, *Psychol. Sci.* **28**, 1321–1333 (2017).

Science

Ten-month-old infants infer the value of goals from the costs of actions

Shari Liu, Tomer D. Ullman, Joshua B. Tenenbaum and Elizabeth S. Spelke

Ranking valuations on the basis of observed choices

Obliged to make a choice between two goals, we evaluate the benefits of achieving the goals compared with the costs of the actions required before deciding what to do. This seems perfectly straightforward, and it is unsurprising to learn that we can also apply this reasoning to others; that is, someone that we see choosing a goal that requires a more costly action must value that goal more highly. What is remarkable, as Liu *et al.* report, is that preverbal children can reason in this same fashion.

Science, this issue p. 1038

ARTICLE TOOLS	http://science.sciencemag.org/content/358/6366/1038
SUPPLEMENTARY MATERIALS	http://science.sciencemag.org/content/suppl/2017/11/21/358.6366.1038.DC1
REFERENCES	This article cites 41 articles, 5 of which you can access for free http://science.sciencemag.org/content/358/6366/1038#BIBL
PERMISSIONS	http://www.sciencemag.org/help/reprints-and-permissions

Use of this article is subject to the Terms of Service