Supplementary Materials for

Ten-month-old Infants Infer the Value of Goals from the Costs of Actions

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

correspondence to: shariliu01@g.harvard.edu

This PDF file includes:

Other Supplementary Materials for this manuscript includes the following:

Materials and Methods:

Experiment 1

Methods

Participants. Our final sample included 24 healthy, full-term infants (15 female, M_{age}=9.95m, range=9.43-10.53). Eight more infants were tested, excluded, and replaced (1 for fussiness that prevented study completion, 1 for technical failure, 5 for inattentiveness during test events, and 1 for experimental error). Sample size and exclusion criteria were fixed prior to the start of data collection, and decisions concerning exclusions were made by researchers unaware of the order of the test events viewed by the infant (and therefore blind to the data that the infant provided). All participants were recruited from the greater Boston area and tested at the Laboratory for Developmental Studies at Harvard University with parental informed consent. Families received a small thank-you gift (e.g. a t-shirt or toy) for participating. All study protocols were approved by the Committee on the Use of Human Subjects at Harvard University.

Materials and Design. All animated events were created in Blender (*38*), synchronized with a custom audio track in iMovie, and presented using Keynote on a 101.6cm by 132.1cm LCD projector screen. Two speakers flanking the screen played all stimuli-related sounds. Infants' looking time data were coded online using Xhab64 (*39*) software and offline using jHab (*40*).

The experiment consisted of 3 pairs of familiarization trials, 1 pre-test trial, and 2 pairs of test trials. All familiarization (Movie S1) and test trials (Movie S2) began with an attention-getting animation and sound (3.0s), followed by looped sequences of events. *Familiarization* sequences consisted of 4 videos (8s or 8.9s; see below) and *test* sequences consisted of a single video (5.2s), each played on a loop with black screens (0.5s) interspersed between the events. Events featured three agents with eyes and a smiling mouth: a central red spherical agent, and two target agents (a cone and cylinder). All familiarization events featured the agent paired with one of the two targets at a time; all test events featured both targets together. The timing of all actions was held constant within familiarization and test blocks. The location of the targets was constant across participants, and the identity of the higher value target was counterbalanced across participants.

During *familiarization* (Movie S1), the agent responded to the call of one of the targets by accepting or refusing to jump over a small (1 units tall), medium (6 units), or large (10 units) barrier that fell with a thud between the agent and target. In events where the agent accepted the cost (8.9s), it looked up at the barrier, made a positive "Mmmm!" sound, and leapt over it to reach the target. The agent always acted efficiently (Gergely et al., 1995) adapting the height of its jump to the height of the barrier; all jumps were accompanied with a popping sound. In events in which the agent refused the cost (8.0s), it looked up at the barrier, made a mildly negative "Hmmmm…" sound, and backed away, returning to the center of the screen. Each block of familiarization events consisted of 4 videos, wherein the agent accepted a small cost and refused a medium cost for one target, and accepted a medium cost and refused a large cost for the other target. (For convenience, we describe events where the agent does not take a costly action as "refusing" or "declining" to take an action, which is how these events appear to adults. But for the purposes of our hypothesis about infants' action understanding, and for our computational model, what matters in these cases is only that the agent is presented with a costly action, considers it and does not take it.) Each familiarization trial consisted of looped blocks either in the above order (small, medium, medium, large) or the opposite order. Across the 6 familiarization trials, both orders were presented 3 times in an ABABAB pattern. The identity of

the higher value target and the first familiarization block were counterbalanced across participants.

A single non-looped *pre-test* event following familiarization featured a still image of the two targets without the agent. During *test* (Movie S4), the agent reappeared, rotated left then right while saying "Hmmm…", and then approached one of the targets: either the higher- or lower value agent. The same sound accompanied each approach. Across 2 pairs of test trials presented in alternation, the agent approached the Triangle target twice and the Square target twice. The first test trial (higher- or lower value approach) was counterbalanced across participants.

Procedure. Infants were seated on their caregivers' laps approximately 1.5m away from the screen. Caregivers were instructed to keep their eyes closed and to refrain from interacting with their infants throughout the experiment, and were monitored for compliance.

After calibrating infants to the screen using a toy, the researcher began the experiment. The researcher had access to a video feed of the infants' faces, a computer screen indicating the current trial, and a third screen indicating when to conclude a trial. The researcher ran the experiment and coded looking time online while unaware of the order of events (and therefore unaware of the infant's differential reactions to the displays), but could determine the start of each trial as well as the timing of actions (e.g. when the central agent approached one of the two targets) based on auditory cues.

Across both the familiarization and test phases of the experiment, the researcher began coding a trial immediately following the attention getter, and concluded the trial once the infant had attended to the screen for 60s cumulatively or looked away for 2s consecutively. During pre-test, the researcher waited until the infant looked towards each target agent at least once, and then began the test trials. These criteria were fixed prior to the start of data collection.

Coding and analysis. Videos of all test sessions were coded offline by observers who were unaware of the order of events that infants viewed, using the same thresholds as online coding, and reviewed for predetermined exclusion criteria (fussiness that prevented study completion, online coding error, experimenter error, technical failure, and parental interference). Further, if infants were determined to have missed a critical part of the test trial (i.e., looked away for the entirety of the approach at test), then that test pair was marked and excluded from subsequent analyses. If infants missed a critical portion of both test pairs, then they were dropped from the sample and replaced. To assess the reliability of the offline-coded data, 100% of the test trials were recoded independently by an additional researcher who was unaware of test pair order. The two coders agreed on trial cutoffs for 95% of the test trials, and the intraclass correlation (ICC) between the two raters was 0.994, 95% CI [0.991, 0.996]. Thus, the primary offline coding data were used in our analyses.

The primary dependent measure was log-transformed looking time (*32*) but plots and descriptive statistics feature raw values for ease of interpretation.

All models were fit in R (*41*). Linear mixed models were fit using the lme4 package (*42*). Detection of influential observations was conducted using the influence.ME package (*43*). Plots were produced using the ggplot2 package (*44*). To explicitly take into account repeated measures, all mixed models included participant identity as a random intercept. Three classes of models were fit: (1) null models, featuring participant identity as the only predictor, (2) hypothesis-driven models, which included additional manipulated factor(s), and (3) exploratory models, which included additional non-hypothesis driven factors. We leveraged likelihood ratio tests (LRTs) and the Aikake Information Crtierion (AIC) to evaluate model fit and parsimony..

All degrees of freedom from mixed effects models were calculated using the Satterthwaite approximation method. Bracketed values indicate 95% confidence intervals.

We predicted that if infants can infer value from effort and use this information to predict the agents' subsequent actions, then they will differentiate between the more probable outcome (when the agent chooses the higher value goal) and the less probable one (when the agent chooses the lower value goal) by showing a looking preference in either direction.

Results

Hypothesis-driven Results. A model with the single predictor of test trial (higher- or lower value) revealed that infants looked longer at the lower value action (M=28.41s, SD=14.85) than the higher value action (M=21.79s, SD=12.29), B=0.327, SE=0.130, β=0.502, $t(24)$=2.523, p=.019, [0.062, 0.591]. This model outperformed a null model by a LRT, $X^2(1)$=5.648, p=.017. A leverage analysis using Cook's Distance revealed 1 influential observation in this model. Removal of this case produced an inferentially equivalent result, B=0.263, SE=0.119, β=0.454, $t(23)$=2.221, p=.037, [0.021, 0.506].

Exploratory Results. A first exploratory analysis tested for an effect of test pair order by including an interactive effect between test trial presentation order and trial type. Infants discriminated between the test events to a similar degree regardless of whether they were assigned to watch the lower- or higher value event first, B=0.212, SE=0.255, β=0.325, $t(24)$=0.829, p=0.415, [-0.310, 0.733]. Removing one influential case produced an inferentially equivalent result, B=0.354, SE=0.226, β=0.610, $t(23)$=1.569, p=0.130, [-0.107, 0.816]. A second model was fit with the additive effects of test pair order and test trial type, summed looking time during familiarization, sex, the identity of the higher value character, and the order of the first block of familiarization. Infants' looking preferences were not predicted by any of the exploratory factors, with all CIs containing 0, ps>0.1. Removal of one influential case produced inferentially equivalent results. The best model out of the above 4 was the hypothesis-driven model (AIC=92.500).

Experiment 2

Methods

Participants. Our final sample included 24 healthy, full-term infants (15 female, M_{age}=9.88m, range=9.47-10.43). Ten more infants were tested, excluded, and replaced (1 for fussiness that prevented study completion, 1 for technical failure, 2 for online coding errors, 2 for parental interference, and 4 for inattentiveness during test events).

Materials, design, and procedure. All materials, design, and procedure were identical to those from Exp. 1 except that during familiarization (Movie S2), each target appeared at the top of a ramp and the agent either refused or accepted to climb the ramp to reach it. To manipulate action cost while controlling for path length, the angle of the ramps varied (11.51°, 39.26°, and 64.09°) such that the agent began 10 Blender units away from the top of the ramp. When the agent accepted the cost (6.8s), the agent moved up the ramp once, slid back down, and then moved all the way up to the target. When the agent refused the cost (5.5s), the agent moved up the ramp once, slid back down, and then turned away from the target, back towards the center for the screen.

Coding and analysis. All coding and analysis procedures were identical to those from Exp. 1. To assess the reliability of the offline-coded data, 100% of the test events were recoded independently by an additional researcher who was unaware of test pair order. The two coders agreed on trial cutoffs for 98% of the test trials, and the intraclass correlation (ICC) between the two raters was 0.978, 95% CI [0.967, 0.985]. Thus, the primary offline coding data were used in

our analyses. Because of our strong directional prediction, all reported p-values in hypothesis-driven results in Exp. 2 are one-tailed. All other reported *p*-values are two-tailed.

Results

Hypothesis-driven Results. As in Exp. 1, infants looked longer at the lower value action (M=30.84, SD=13.79) than the higher value action (M=27.05s, SD=17.55), B=0.250, SE=0.109, β=0.408, $t(24)$=2.294, p=.015, [0.028, 0.472]. This model outperformed a null model by LRT, $X^2(1)$=4.760, p=.029. No influential cases were detected.

Exploratory Results. An exploratory model testing explicitly for presentation order revealed that infants differentiated between the test events differently depending on whether they were assigned to watch the lower value versus the higher value approach first, B=0.521, SE=0.190, β=0.851, $t(24)$=2.741, p=.011, [0.133, 0.909]. No influential cases were detected. which revealed that whereas infants who saw the lower value test event first looked longer at the lower value (M=28.70s, SD=11.12) versus higher value (M=20.01s, SD=14.83) test trials, B=0.511, SE=0.134, $t(24)$=3.797, p=.001, [0.233, 0.788], infants who saw the higher value choice did not differentiate between the lower- (M=32.97s, SD=15.39) and higher value test events (M=34.08, SD=17.78), B=-0.011, SE=0.134, $t(24)$=0.079, p=.938, [-0.288, 0.267]. An additional model testing for effects of summed attention during test, sex, the identity of the higher value target, and the first familiarization loop revealed that no further effects other than one of first familiarization loop, where infants assigned to watch a sequence of low to high cost first looked longer overall at test, B=0.440, SE=0.185, β=0.851, $t(24)$=2.374, p=.023, [0.062, 0.819]. Removal of 3 influential observations from this model yielded inferentially equivalent results. The best model of the four reported was the simpler exploratory model with the single interactive effect (AIC=78.219).

Experiment 3

Methods

Participants. Our final sample included 32 healthy, full-term infants (15 female, M_{age}=10.03m, range=9.57-10.50). Six more infants were tested, excluded, and replaced (1 for fussiness that prevented study completion, 3 for online coding errors, 1 for parental interference, and 1 for inattentiveness during test events). Sample size was determined from a simulation power analysis over data from Exp 1-2. The design, procedure, and analyses of this experiment were pre-registered via the Open Science Framework (https://osf.io/k7yjt/).

Materials, design, and procedure. All materials, design, and procedure were identical to those from Exp. 1-2 except as follows. During familiarization (Movie S3), each target appeared at the far end of a trench (15 units deep), and the agent either accepted this cost by jumping across it or refused to jump. To manipulate action cost while controlling for movement against gravity, the width of the trench varied (5, 10, and 15 units) such that the agent began 13, 18, or 23 units away from the target. When the agent accepted the cost (8.4s for the small cost, 9.3s for the medium), the agent looked down at the bottom of the trench, looked at its target, and backed up and leapt across the trench. When the agent refused the cost (7.2s), it looked down at the bottom of the trench, looked at its target, and backed up the turned away, back towards the center for the screen. The agent backed away a longer distance and accelerated more when accepting the medium versus the small cost. During test (Movie S5), infants saw the same choice events as in Exp. 1-2, but the events took place on a platform at the same height used for familiarization. Prior to familiariazation, infants watched an additional video (Movie S6) , in which a ball rolled across the platform and off the small, medium, and wide trench (5.8s, 5.7s, and 5.0s), identical in width and depth to those during familiarization and situated at the center of the screen. During

each segment of the video, the ball traveled in a parabolic trajectory after it rolled off the edge and shattered against the far wall of the trench. The location of impact depended on the width of the trench, such that the ball shattered at a lower point on the far wall for wider trenches. The side of the screen that the ball emerged from (left vs. right) was consistent across all videos within participants and was counterbalanced across participants.

Coding and analysis. All coding and analysis procedures were identical to those from Exp. 1-2. To assess the reliability of the offline-coded data, 100% of the test events were recoded independently by an additional researcher who was unaware of test pair order. The two coders agreed on trial cutoffs for 96% of the test trials, and the intraclass correlation (ICC) between the two raters was 0.995, 95% CI [0.993, 0.996]. Thus, the primary offline coding data were used in our analyses. All reported p-values in hypothesis-driven results in Exp. 3 are one-tailed. All other reported *p*-values are two-tailed.

Results

Hypothesis-driven Results. As in Exp. 1-2, infants looked longer at the lower value action (M=23.05s, SD=13.58) than the higher value action (M=17.47, SD=10.69), B=0.260, SE=0.119, β=0.403, $t(32)$=2.185, p=.018, [0.020, 0.501]. This model outperformed a null model by LRT, $X^2(1)$=4.452, p=.035. Removal of 1 influential case produced an inferentially equivalent result, B=0.286, SE=0.120, β=0.466, $t(31)$=2.379, p=.011, [0.043, 0.529].

Exploratory Results. An exploratory model testing explicitly for presentation order revealed that like in Exp. 2, infants differentiated between the test events differently depending on whether they were assigned to watch the lower value versus the higher value approach first, B=0.770, SE=0.195, β=1.194, $t(32)$=3.941, p<.001, [0.357, 1.165]. We detected 1 influential observation in this model and removed it from subsequent pairwise comparisons, which revealed that whereas infants who saw the lower value test event first looked longer at the lower value (M=27.95s, SD=14.11) versus higher value (M=15.05s, SD=9.06) test trials, B=0.645, SE=0.117, $t(31)$=5.519, p<.001, [0.407, 0.883], infants who saw the higher value choice showed a weak preference for the higher-value (M=20.62, SD=11.93) relative to the lower-value (M=16.56s, SD=9.84) test events, B=-0.236, SE=0.121, $t(31)$=-1.959, p=0.059, [-0.483, 0.010]. An additional model testing for effects of summed attention during test, sex, the identity of the higher value target, and the first familiarization loop revealed that no further effects other than one of total attention in seconds during familiarization, where infants who were more attentive during familiarization looked marginally longer overall at test, B=0.002, SE=0.001, β=0.261, $t(32)$=1.948, p=.060, [0.000, 0.005]. Removal of two influential cases from this model yielded an inferentially equivalent result, and revealed an additional finding that infants randomly assigned to conditions where the higher value target was on the right looked longer overall at test, B=0.368, SE=0.149, β=0.610 $t(30)$=2.464, p=.020, [0.066, 0.670] The best model of the four reported was the simpler exploratory model with the single interactive effect (AIC=114.24).

Results Across Exp 1-3. As reported in the main text, across all experiments, infants looked longer at the lower value action (M=26.99s, SD=14.13) than the higher value action (M=21.64s, SD=13.94), B=0.277, SE=0.070, β=0.424, $t(80)$=3.975, p<.001, one-tailed, [0.139, 0.415], mixed effects model with random intercepts for participant and experiment. Removal of one influential case produced an inferentially equivalent result, B=0.258, SE=0.068, β=0.406, $t(79)$=3.799, p<.001, one-tailed, [0.123, 0.393]. To test explicitly for differences in responses to test events across experiments, an additional model with an interactive effect between experiment and test event was fit and revealed no differences in looking preference across

experiments, and removal of one influential observation produced an inferentially equivalent result.

Two further models tested explicitly for the effect of presentation order on attention to the lower- vs. higher value test events, and additional predictors of summed attention during test, sex, the identity of the higher value target, and the first familiarization loop. These models revealed a robust effect of presentation order, B=0.528, SE=0.126, β=0.808, t(80)=4.179, p<.001, [0.277, 0.778], and for an effect of attention during familiarization, B=0.002, SE=0.001, β=0.209, t(79.65)=2.221, p=.030, [0.000, 0.006]. mixed effects model with random intercepts for participant and experiment. Removal of influential cases produced inferentially equivalent results. A comparison of model fit and parsimony revealed that the model including the single interactive effect provided the best description of the data (AIC=280.90).

Fifty out of the 79 infants who demonstrated a preference in either direction (63.29%; 1 infant in Experiment 2 looked for 60 seconds on all test trials) looked longer on average at the lower value test events, p=.033, 95% CI [0.269, 0.490], exact binomial test. The proportion of infants who looked in the predicted direction did not differ across Experiments 1 (17/24, 70.83%), Experiment 2 (14/23, 60.87%), and Experiment 16 (19/32, 59.38%), $X^2(2)$=1.022, p=.600. See Figure S3. Further non-parametric analyses on raw looking times in seconds supported the finding that infants looked longer at the lower value test events across all experiments, [2.165, 8.185], V=2241, p=.001, Wilcoxon signed rank test, and [0.720, 8.235], bootstrapped median difference in looking times across test events with 10,000 samples.

To compare attention during familiarization across the experiments, we fit a linear model including summed looking time in seconds as the dependent variable and experiment (1, 2, or 3) as a predictor. We found that infants looked longer during familiarization in Exp. 1 (M=261.96s, SD=61.29) than in Exp. 2 (M=216.50, SD=63.49), B=45.47, SE=19.22, β=0.653, t(77)=2.366, p=.021, [7.196, 83.74], and in Exp. 3 (M=209.70, SD=72.33), B=52.26, SE=17.98, β=0.751, t(77)=2.907, p=.005, [16.460, 88.056]. See Figure S2.

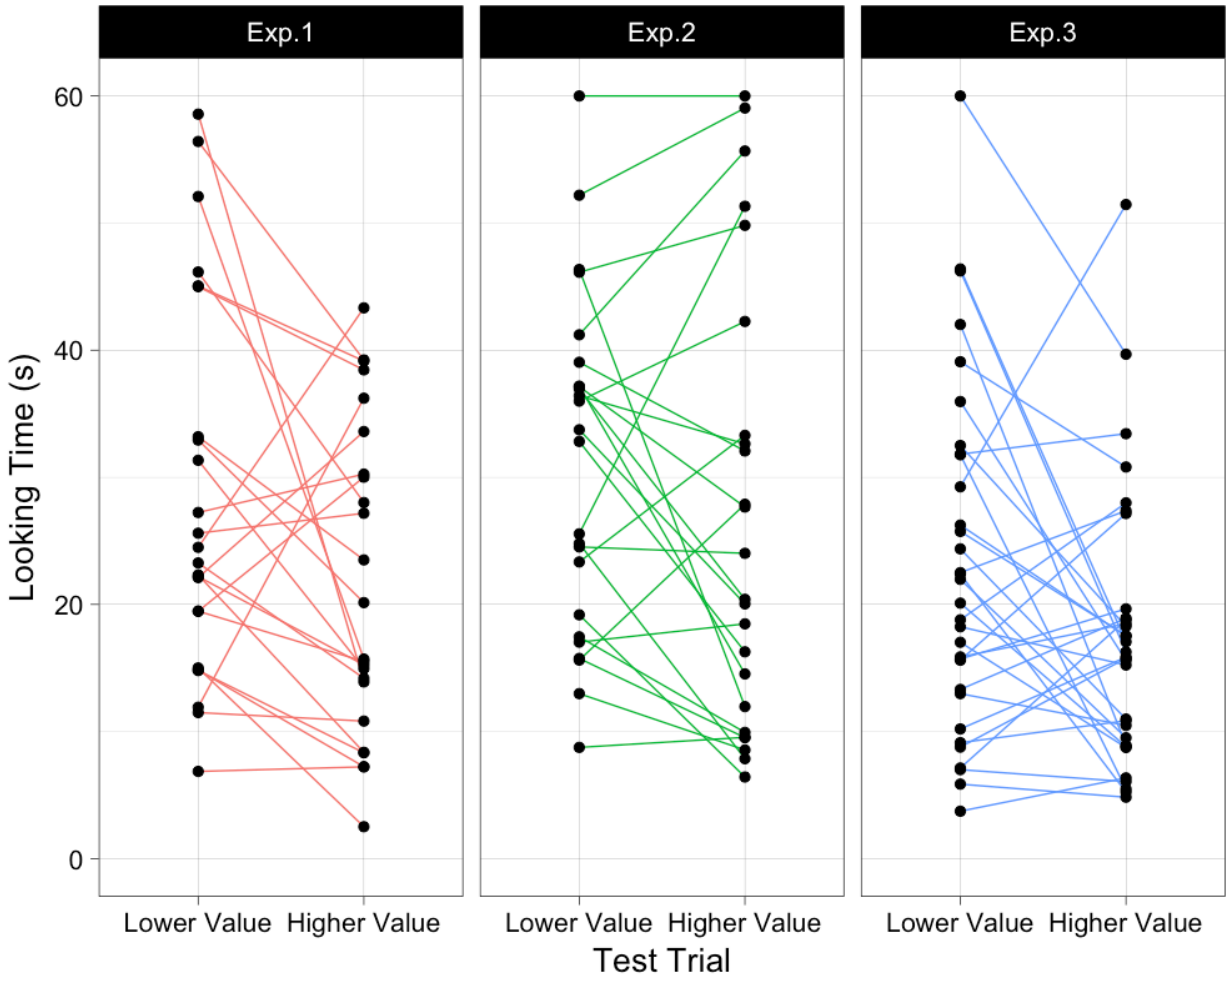

Fig S1. Looking times averaged across two test pairs towards the lower value and higher value test event from all participants in Experiments 1-3.

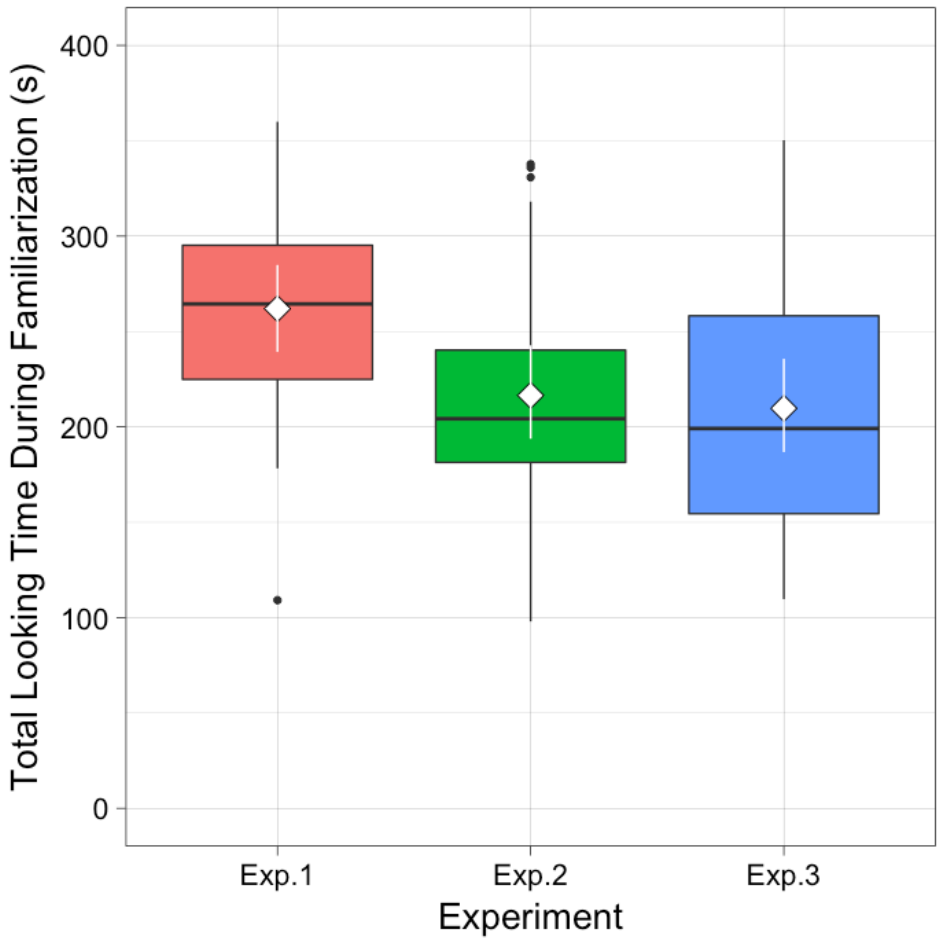

Fig S2. Total looking time in seconds during familiarization across Exp 1-3. Boxes indicate middle quartiles, vertical lines indicate points within 1.5 times the interquartile range from the 25th and 75th percentiles, and horizontal lines indicate medians. Means and 95% confidence intervals are plotted in white.

Movie S1. Sample familiarization event for Experiment 1.
Movie S2. Sample familiarization event for Experiment 2.
Movie S3. Sample familiarization event for Experiment 3.
Movie S4. Sample test event for Experiments 1-2.
Movie S5. Sample test event for Experiment 3.
Movie S6. Pre-familiarization event for Experiment 3.

Computational Modeling Details

We constructed several alternative models for predicting the actions of a central red agent, given previous stimuli showing its behavior. We evaluated the models both as accounts of our present three experiments testing infants' ability to integrate costs and rewards in action understanding, and as accounts of five previous experiments that tested infants' abilities to make inferences about rewards (*6*) or costs (*7*, *13*) individually. As our main proposed model, we constructed a probabilistic program for inferring rewards over possible goals through a Bayesian computation, based on the Naïve Utility Calculus (NUC) and inverse planning. The cost-function of this model was either based on distance (NUC-distance), or physical effort through the application of force (NUC-force). We next constructed three alternative models that were lesions of various aspects of this model. Only one of our two NUC models (NUC Force; Fig. S3B) accounts for all of these findings, including Experiment 2 in which infants saw the agent cover the same total distance at the same speed but on trajectories requiring different amounts of physical effort due to the force of gravity. In contrast, two models that do not integrate costs and rewards (Costs Only, Fig. S3E and Rewards Only, Fig. S3F), as well as two integrative but simpler models using only perceptual cues for value (Counts, Fig. S3D) or distance-based representations of cost (NUC Distance, Fig. S3C) instead of abstract effort-reward tradeoffs, account for just a subset of this body of empirical findings.

The eight experiments shown in Fig. S3 only begin to test the explanatory scope of our computational framework. This approach naturally extends to explain other findings (*23*), including infants' rational imitation of other agents' goal-directed actions (*45*, *46*), understanding of the transitivity of agents' preferences (*47*), inferences of agents' preferences from their non-random selections of objects (*48*), and inferences about the existence and location of an occluded object from the cost of an agent's action (*49*).

In the rest of this section we describe the full NUC model, then the alternative proposals. We then compare the predictions of the different models to the empirical results of eight different experiments with young children that all relate to reasoning about cost, reward and efficient behavior.

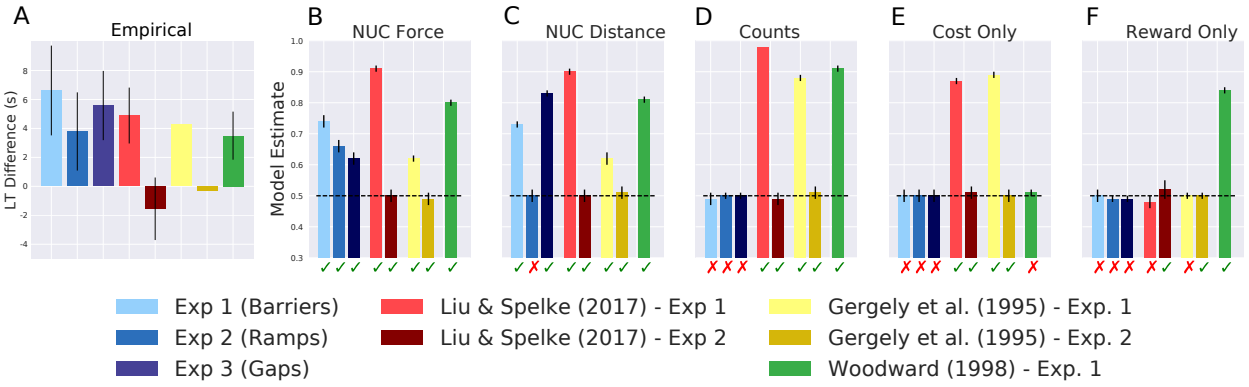

Fig S3. Comparison of model predictions (B-F) to the empirical results (A) of Experiments 1-3 in this paper, as well as additional experiments on infants' sensitivity to costs and rewards (*6*, *7*, *13*). Model estimates were generated from 5,000 MH samples with a 10-step interval and burn-in of 1,000. Error bars indicate 95% CIs (model estimates) and standard errors (where available from empirical findings). Check marks and crosses respectively indicate consistency or inconsistency with empirical findings.

Decision-making framework

Following the logic of "Bayesian Theory of Mind" (*4*, *23*, *50*, *51*), the primary model assumes agents have a planning procedure for generating actions that maximize their utility or expected utility, and that agents then use observed actions to invert this procedure to reason about the utilities and constraints of the planning agents.

Following a standard decision-making framework for rational planning (e.g. *48*) we assume that a rational planning agent:

(i) Divides the world into possible states *S*, such that for each state $s \in S$ there exists a set of possible actions A_s.

(ii) Has access to a transition function *T* such that:

$$T(S, A_s) \rightarrow P(S') , \quad (1)$$

where *P(S')* is the probability distribution over the states of the worlds S' that can result from taking action A_s in state S.

(iii) Has a utility function *U* such that:

$$U(A, S) = Reward(S) - Cost(A), \quad (2)$$

where *Reward* and *Cost* are functions that map from states and actions to real numbers. This separation of the utility into independent components follows recent work by Jara-Ettinger et al. (2016), who showed that both components can be separate targets of inference, and used as explanatory variables by young children across different scenarios.

(iv) Is guided by a decision-function $D(S, A_s) \rightarrow P(A)$ that selects a given action in a state, in order to maximize the expected utility U. For our simple scenarios, we assume a soft-max decision function (*53*):

$$P(\mathrm{A}_s = \mathrm{a} | U, S = s) = \frac{e^{\beta U(a,T(s,a))}}{\sum_j e^{\beta U(a,T(s,a))}} , (3)$$

where β is a noise-parameter (inverse-temperature) adjusting the agent's determinism. As $\beta \rightarrow 0$, the agent will behave in a more random fashion. As $\beta \rightarrow \infty$, the agent will tend to use greedy action selection. Intermediate values of β (around 1) approximate probability matching behavior. Thus, we place an equal prior probability on β taking one of two values, β = 3 (representing a near-optimal rational agent) and β = 0 (representing a random, non-rational agent). The exact value of the parameter for a rational agent is not important for the qualitative pattern of results discussed in the following sections.

The Specific Decision-making Environment

States: In our studies infants saw the agent making a single choice between staying put and moving towards one of two possible goal agents. The full set of possible states in Experiments 1-3 are $s \in \{Start, Target_A, Target_B\}$, where $Target_i$ can be the Square or Triangle goal agent. Not all states are available in every situation, depending on the stimuli (for example, if the stimuli shows the decision making agent near a ramp with the Square goal agent at the top, the sub-set of states for this environment does not include reaching the Triangle goal agent).

Actions: In Experiment 1 the actions in any situation were a subset of $\{Nothing, Left, Right, Jump\ tall, Jump\ medium, Jump\ short\}$.

In Experiment 2 the actions were a subset of $\{Nothing, Left, Right, Climb\ steep, Climb\ moderate, Climb\ shallow\}$.

In Experiment 3 the actions were a subset of $\{Nothing, Left, Right, Jump\ wide, Jump\ medium, Jump\ narrow\}$.

The set of possible actions and states was similarly altered for the other 5 experiments we considered, that were not part of the current study.

Cost functions: We do not know the specific cost that infants believe the agent incurs for jumping over the barriers or climbing the inclines, but we assume that the cost for jumping/climbing is greater than staying put, and that cost generally scales with distance. There are at least two potential models for how to take this distance into account when calculating cost. The first model (NUC-distance) ignores physical forces and effort, and uses only total distance traveled between states as the input to the cost function:

$$Distance_{cost}(S_1 \to S_2) \propto \int_C a\, ds, \quad (4)$$

where *C* is the trajectory from S_1 to S_2, and *a* is some constant factor. We considered each Blender unit of distance to be one unit of cost for the model.

The second model (Force) considers the physical work required to get from one state to the next. Specifically, we consider the work done in the presence of a conservative force field (gravity):

$$Work_{cost}(S_1 \to S_2) = \int_C F \cdot ds = \int_{t1}^{t2} F \cdot dt. \quad (5)$$

For Experiments 1&2, we considered the force of over-coming friction as well as gravity, meaning $F = f \cdot g \cdot \cos(\alpha) + g \cdot \sin(\alpha)$, where α is the angle opposite the vertical part of the incline, and f is the friction coefficient. The mass (m) was set arbitrarily to 1 and so does not appear. Without the friction component the non-accelerated movement of an agent on a non-inclined plane is effortless. We set the friction coefficient arbitrarily to *f = 0.5,* similar to wood sliding on wood or metal sliding on wood, though we note that the exact choice of *f* in the range 0 to 1 does not affect the qualitative results. For experiment 3 we considered in addition the impulse forces that accelerate and decelerate the agent. This instantaneous change in velocity can be related to the work done as a change in kinetic energy, meaning $W = \Delta E_k = \frac{1}{2}(V_1^2 - V_2^2)$, where V_1 and V_2 are the velocities before and after the application of the impulse. Both models would predict similar cost differences for Experiments 1 and 3, but they would diverge with regards to Experiment 2. $Distance_{cost}$ of the steep and shallow inclines is the same, while $Work_{cost}$ is sensitive to the incline of the ramp.

Rewards: We do not make a-priori assumptions about the reward associated with reaching the goal agents, except that they can potentially exceed the range of the costs:

$$\text{Reward(Target A)} \sim \text{Uniform}(0, \text{Reward}_{max}),$$
$$\text{Reward(Target B)} \sim \text{Uniform}(0, \text{Reward}_{max}).$$

Inference

Using the previous assumptions, we can calculate the posterior distribution over the rewards and costs conditioned on an observed action A. Applying Bayesian reasoning, this is a combination of the likelihood of the observed action (given by the planning procedure), and the prior distributions over costs and rewards:

$$P(Reward, Cost \mid A, S) \propto P(A \mid S, Cost, Reward)\, P(Reward, Cost) \quad (6)$$

We constructed a probabilistic program that samples rewards and costs from this posterior, conditioned on observed actions matching the actions of the agent . The program was written in Church (*54*), a probabilistic programming language based on Scheme.

This program can be used to infer the value of the reward for the different experiments (cases 1-8 as detailed above). Below we consider the specific case of Experiment 2 from the current paper (Ramps), using the NUC-force model. The sampling procedure in Church uses the Metropolis-Hastings algorithm, and Fig. S4 shows the resulting approximation to the posterior in Eq. 6 for Experiment 2 and the NUC-force model, with 5,000 samples at 10-step intervals with a burn-in of 1,000 samples. As can be seen in Fig. S4, the model shifts probability away from a uniform distribution for both goals, such that the reward for the higher-effort target is higher in expectation than the lower-effort one. For both the targets, most of the probability mass is in between the accepted and rejected costs, as expected.

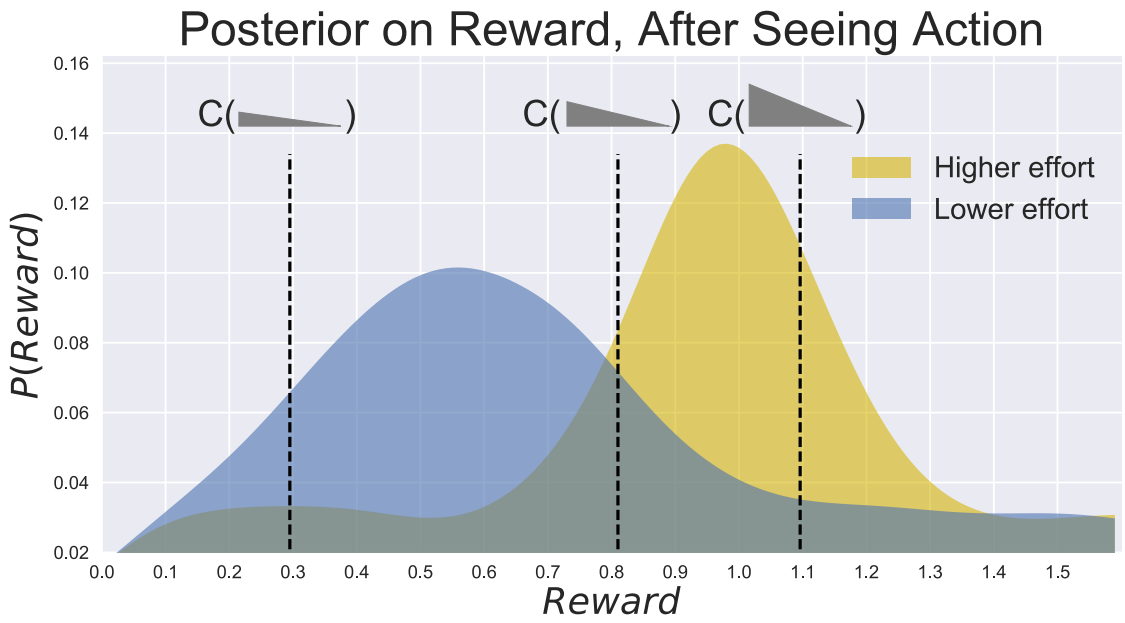

Fig S4. The model's inferred posterior probability distribution over possible reward values, after seeing the actions taken in Experiment 2 when using a Force-based cost function. Note that for the force-based model the cost of the different ramps is a combination of the effort of working against gravity and overcoming friction.

Action Prediction

The 'forward planning' part of the program implements Eq. 1-6 and the decision-making framework of the first section. This forward planning can use the posterior distribution inferred from observing the agent's action to predict its next actions, P(Action | Previous stimuli). This probability can be calculated for cases in which there is only one reward and multiple possible actions (as in Gergely et al. 1995, Experiments 1&2) or multiple rewards (the other experiments considered). In all model comparisons we used 5,000 samples at 10-step intervals with a burn-in of 1,000 samples for the inference of the reward distributions, and for the sampling of action predictions. We discounted the 'Do Nothing' action predictions, as the infants never saw these in the test stimuli, and were faced rather with a direct comparison between two actions.

To give an example of a particular way in which action prediction proceeds, we consider again Experiment 2 from the current paper under the NUC-force model. For the test stimuli, the situation includes no barriers, and both targets are obtainable. The possible actions are: going to target A, going to target B, or doing nothing. We assume that:

$$\text{Reward(Target A)} \sim P(\text{Reward(Target A)} \mid \text{Previous Stimuli}), \quad (7)$$

$$\text{Reward(Target B)} \sim P(\text{Reward(Target B)} \mid \text{Previous Stimuli}), \quad (8)$$

where $P(\text{Reward(Target)} \mid \text{Previous Stimuli})$ is calculated using the program approximating Eq. 6.

Fig. S5 shows the resulting prediction for the probability distribution over the agent's action. As can be seen in Fig. S5, the model predicts the agent will move towards the Target that it expended higher effort to reach earlier. Consequently, when the agent moves towards the lower effort Target, this goes against the prediction.

The amount of surprise the model predicts can be directly related to the inverse of the probability of a given event or action (*34*).

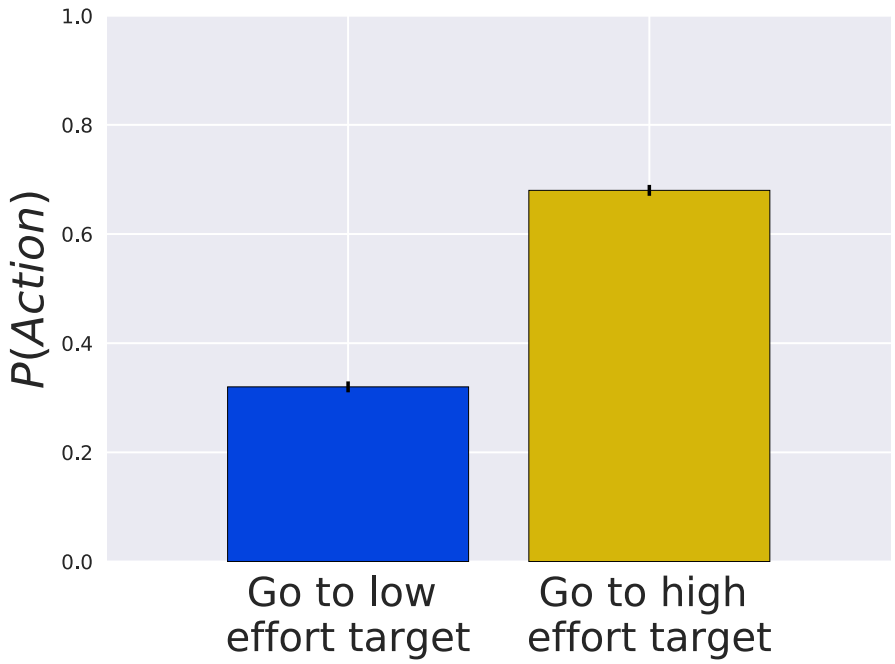

Fig S5. Probability of predicted next action by the NUC-force model for Experiment 2 (Ramps), for the test case in which both targets are equally accessible. Black lines show bootstrapped 95% confidence intervals.

Alternative Models

The NUC-force and NUC-distance models outlined so far assume that infants are able to use Bayesian reasoning to integrate inferences about cost, reward, and agent rationality to predict the next action of an agent. The NUC-distance model assumes that cost is proportional to perceptually observable distance, while the NUC-force model assumes that cost is proportional to physical effort.

But an alternative option is that infants are performing a much more low-level analysis of the scene, or reasoning about costs and rewards in a non-integrated way. Consider for example the findings from Experiment 1 in Woodward (1998), showing that 9-month old infants expected an agent to reach for a goal object A over B, after A and B had switched positions compared to habituation. A full NUC mental model could be deployed to infer that the reward for A is higher than the reward for B, and then correctly predict the action. But it is also possible that infants are reasoning along the lines of "A goal previously reached for will be reached for again, regardless of its position".

We thus consider 3 alternatives to the NUC model, all of which involve different forms of lower level or non-integrated reasoning, as follows:

1) **Count-based:** This is a perceptual-based account that relies on directly observable cues: reaching an item, and the distance to the item. This model is similar to the NUC in that agents are assumed to be rational planners that act to achieve goals under constraints. However, rather than using full Bayesian inverse planning to reason about rewards, it uses an easy-to-calculate proxy for the reward of a target: The model tallies the number of times a goal object has been approached or reached, and considers that tally in proportion to the total number of times any goals have been reached. Furthermore, the model considers the cost as proportional to the distance. Thus, for this model:

$$P(Reward_i | Stimuli) = \frac{\#(Reward_i\ approached)}{\sum_j \#(Reward_j\ approached)} \quad (9)$$

The cost is calculated as in Eq. 4, and the planning proceeds as in Eq. 1-3.

2) **Cost Only:** This model assumes all goal states are equally rewarding, but is sensitive to the cost in the form of distance as in Eq. 4. Such a model is able to reason about agents acting efficiently to reach goals, in the sense of minimizing distance as in Gergely et al. (1995), for example.

3) **Reward Only:** This model assumes rewards can vary and performs the correct Bayesian updating on the probability distribution over reward as in Eq. 6, but without the cost factor. Such a model is sufficient for correctly predicting behavior when reasoning about the choices of agent that shows a simple preference for one target over another, as in Woodward (1998), for example.

Comparison to Empirical Data from Eight Experiments

In this section we compare the outputted prediction of the NUC models as well as the alternative models, to eight different experiments examining infants' expectations about cost, reward and efficient behavior. The experiments considered were:

1) Experiment 1 in the current paper (Barriers).
2) Experiment 2 in the current paper (Ramps).
3) Experiment 3 in the current paper (Gaps).
4) Experiment 1 in Liu and Spelke (2017), in which six-month-old infants first view an agent jumping over barriers of varying sizes to get to a goal. In the test phase, the agent is blocked by a previously unseen small barrier and either makes a small jump (expected) or a large jump (unexpected) over the barrier to reach the goal.
5) Experiment 2 in Liu and Spelke (2017), which is identical to (3), except that the barriers are placed behind the goal such that they are not blocking the path of the agent. Infants' looking time was at chance when shown a small or large jump, presumably inferring that the agent is irrational based on previous behavior.
6) Experiment 1 in Gergely et al. (1995), in which 12-month-olds first view an agent jumping over a barrier to reach a goal. In the test phase, the barrier is removed and the agent either moves straight to the goal (expected) or performs the previously seen jump on the way to the goal (unexpected).
7) Experiment 2 in Gergely et al. (1995), which is identical to (5) except the barrier is placed behind the goal such that it is not blocking the path of the agent. Infants' looking time was at chance when shown the agent moving straight or making a jump on the way to the goal.
8) Experiment 1 in Woodward (1998), in which 9-month-old infants first saw an agent reaching for and grasping one of two possible goals. In the test phase, the positions of the goals were switched. The agent either reached to the same goal using a new trajectory (expected) or follow the same trajectory to reach a new goal (unexpected).

We use the same model set-up and parameters for all experiments, varying only the actions and stimuli observed, and the different cost and reward functions used depending on the model.

Following (*34*), we relate the amount of surprise in infants' looking time to 1-P(outcome) as predicted by the model. Specifically, we relate this measure to the difference in looking time (measured in seconds) between the unexpected event and the expected event. The results of the models and the comparison to the eight experiments are shown in Fig. S3.

Code

All the code implementing the computational models and the analysis of their output is available on Open Science Framework (https://osf.io/crx4d/).