Infants' interpretation of information-seeking actions

Bálint Varga (varga_balint02@phd.ceu.edu)
Gergely Csibra (csibrag@ceu.edu)
Ágnes Melinda Kovács (kovacsag@ceu.edu)
Department of Cognitive Science, Central European University,
Budapest, Hungary

Abstract

Although infants can frequently observe others gathering information, it is an open question whether and how they make sense of such activities since the mental causes and intended effects of these are hidden and underdetermined by the available evidence. We tested the hypothesis that infants possess a naive theory that leads them to grasp the purpose of information-gathering actions when they serve as sub-goals of higher-order instrumental goals. We presented 14-month-old infants with actions that were inefficient with respect to the agent's instrumental goal but could or could not be justified as information-seeking behavior via this theory. We expected longer looks in the condition where the detour could not be justified and the results were in line with our predictions. While this evidence is compatible with our hypothesis, further studies are in progress to rule out alternative interpretations of our findings.

Keywords: action understanding; information-seeking; theory of mind; cognitive development

Introduction

Research on infant action interpretation has provided a fair amount of evidence on how babies start to understand behavior as intended towards changing the *physical* (e.g. Gergely & Csibra, 2003; Liu et al., 2017) and *social* (e.g. Tauzin & Gergely, 2018; Vouloumanos, Onishi, & Pogue, 2012) environment. However, we do not know much about how infants attribute goals to actions that are aimed to change the acting agents' own *epistemic* states: that is, how they interpret information-seeking actions. Ranging from eye-movements and active locomotion to sophisticated experimentation, information-gathering is ubiquitous in the human realm. These exploratory behaviors rarely seem random, as observers we can often recognize specific questions that inquisitive agents seem to untangle in a systematic manner (Bennett, 1976).

Since information-gathering acts are both common sub-routines of many everyday instrumental tasks (that change the observable state of affairs) and are also used as means to satisfy curiosity, infants, at least intuitively, have vast opportunity to observe such activities – still, it is an open question whether and how they make sense of them. From an evolutionary point of view, one could argue that interpreting and predicting information-gathering is no less beneficial than making sense of any other kind of goal-directed behavior, at least if such interpretative capacities are indeed adaptations for competition, exploitation (Whiten & Byrne, 1997) or cooperation (Moll & Tomasello, 2007). On the other hand, understanding information-seeking poses unique challenges for observers, as their mental causes and intended effects are both hidden and underdetermined by the available evidence: there is an enormous set of potential variables that an observed agent might be interested in, while the same actions could provide very different sorts of information for an agent.

This inverse problem of action interpretation is not exclusive to information-seeking: as it is often formulated, in principle, any finite set of goals are compatible with an infinite number of action means, and any stream of behavior is compatible with an infinite number of goals (e.g. Davidson, 1963). To solve this inductive problem, several accounts have introduced the idea that what ultimately constrains how we conceptualize behavior is a set of assumptions of rationality (e.g. Baker et al., 2017; Baldwin, Loucks, & Sabbagh, 2008; Gergely & Csibra, 2003).

For instance, according to the *teleological stance* theory put forward by Gergely and Csibra (2003) infants possess a naïve theory of rational action which guides them to expect that (1) actions have a fundamental function of achieving certain goal states in the world and that (2) acting agents will consistently apply to the most efficient (least costly) means to reach these states, given the constraints of the particular situation. If the observed action means are justified by environmental constraints, then the action, the outcome, and the environmental constraints produce a well-formed teleological schema, and the observed outcome is attributed as the agent's goal. In addition, this teleological theory is hypothesized to be productive to the degree that it allows the prospective inference of goal states and retrospective inference to situational constraints as well.

A related proposal states that adults and infants are not only able to represent action costs but can also integrate these variables with agents' preferences as utilities and take them as the underlying causes of observed behavior (Jara-Ettinger et al., 2016). The suggestion is that the naïve theory that lies beneath human action understanding represents agents as utility maximisers, that is, rational planners who calculate the utilities of candidate actions (via subtracting costs from rewards) and select the one with the highest positive utility. With this *naïve utility calculus* at hand, observers could predict agents 'future actions based on the rewards and costs they assign to them and also infer backwards to the costs and rewards from observed behavior.

What is special to information-gathering acts, however, is that their intended, epistemic effects are also *unobservable*. To clarify, we take epistemic effects as *becoming informed about something*. Information, at least according to one prominent account, is the reduction of uncertainty over a set of possibilities (Shannon, 1948). Implied by this definition, the information that the very same state or event provides is not an intrinsic property of it, but always relative to the prior hypothesis space of the information-gathering agent. If these considerations are valid, it is reasonable to assume that the interpretation of information-seeking acts would need to involve some kind of quantification over mutually exclusive possibilities that the observed agent entertains. Moreover, this formulation suggests that a core problem an observer has to face when interpreting information-seeking is the specification of the hypothesis space the acting agent has in mind and aims to reduce its uncertainty about. Arguably, if there were certain rules that govern the way agents tend to gather information, representing these could provide inductive constraints to solve this problem.

Accepting the dichotomy between *interested* and *disinterested* sampling (Chater, Crocker, & Pickering, 1998), we propose that the former is the kind of information-gathering that is governed by specific principles, possibly exploitable by even young observers. [1] Interested sampling is at work when an agent has a higher-order instrumental decision problem that she tries to maximize. In this case, the gathered information has value if it changes the agent's expectations regarding which is the best option within this decision problem. For example, if one is faced with the decision whether to pack an umbrella for a trip, watching the forecast is valuable to the extent it helps making this decision. Disinterested sampling, in contrast, is not attached to any instrumental goal or task: the agent seeks information in spite of the fact that it lacks any specific knowledge about its potential material or strategic benefits. Rather, its goal is merely to identify the correct hypothesis in a given hypothesis space. Here, information can be valued according to the extent it brings closer to this desired epistemic state.

We argue that observers can exploit the fact that interested sampling is attached to some higher-order instrumental goal with a corresponding decision problem. Specifically, we hypothesize that the following set of assumptions, constituting a *naïve theory of information-seeking*, could be used to reliably predict and attribute interested information-seeking in most cases: (1) *having no sufficient access to goal-relevant variables gives rise to uncertainty*, and that (2) *agents will engage in actions towards the sub-goal of reducing this task-related uncertainty,* but (3) *only if the expected rewards of these actions are higher than their cost*.

The application of such a theory would admittedly rest on a set of capacities, however, some of these were indicated to be present already in infancy, including the inference of instrumental goals and preferences (e.g. Csibra, 2008; Liu & Spelke, 2017; Woodward, 1998), and the inference of third-party perceptual access to goal-relevant aspects of a scenario, together with the causal knowledge that perception leads to knowledge (e.g. Luo & Baillargeon, 2007; Luo & Johnson, 2008; Surian, Caldi, & Sperber, 2007).

On the other hand, the theory would also entail the attribution of uncertainty over a set of possibilities and it is far from clear whether infants can ascribe such mental states to agents (although evidence on early grasp of at least a rudimentary form of uncertainty was indicated by Liszkowski, Carpenter, & Tomasello, 2008). Finally, assumption (3) would require an extension of the cost-benefit calculations supposedly present in a naïve utility calculus (Jara-Ettinger et al., 2016), namely the evaluation of information according to the gain it provides in terms of expected utility, for which no evidence is available as yet.

Here, we aimed to test assumptions (1) and (2) only, asking whether infants will attribute uncertainty to agents, and use this inferred mental state for the attribution of information-seeking goals to these agents 'subsequent actions.

Methods

We tested 14-month-old infants' responses to an action that was inefficient with respect to the agent's ultimate instrumental goal of reaching an object, but in a particular case could be justified as efficient information-seeking behavior, geared with the proposed naïve theory. The experiment consisted in two familiarization phases followed by two test trials (see Figure 1 for an overview).

First, we demonstrated six times that a red, ball shaped agent prefers to selectively approach a target object (a duck), independent of its location, in contrast to a different, blue object.[2] Based on previous findings (Woodward, 1998), we expected infants to infer, given these events, that the agent had the goal or preferred to approach the duck. Importantly, the two objects had identical parts that played a role later, in the test phase.

In the second phase of familiarization, we showed two different situations, two times each. The aim of these phases was to familiarize infants with the different perspectives that can be taken in a novel environment. In one scenario, the agent came out of a house, looked at the duck jumping to one of two boxes through a tube, then followed it by jumping in the same tube. In the other situation, the agent was standing at a platform below, facing away from the tubes and boxes, while the non-target object appeared behind her and jumped to one of the boxes with a bumping sound. The agent turned around, looked at the non-target object in the box, then turned back, without following it to the tube.

[1] These two forms of inquiry have also been labelled as *situation-specific* and *informational* (Coenen, Nelson, & Gureckis, 2019), or *instrumental* and *non-instrumental* (Gottlieb & Oudeyer, 2018).

[2] Although, for simplicity, we refer to these as "objects", in fact, they possessed agentive features (eyes and self-propelled motion), only in order to make the scene more engaging. Our predictions, however, were independent of these features.

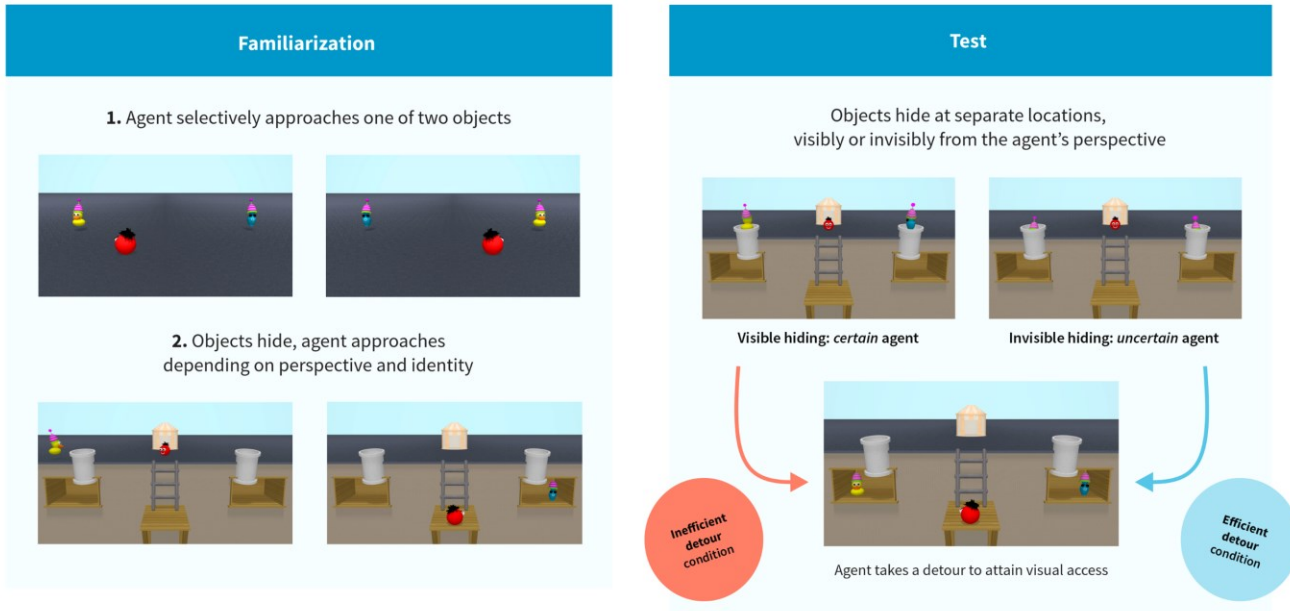

Figure 1: Overview of the experimental design with frames from the corresponding stimuli.

Finally, infants were shown two test trials, that differed in one critical aspect. In one case, when the agent came out of the house, she could fully see each object before they simultaneously jumped into the boxes. We assumed that this perspective could enable infants to attribute certainty to the agent regarding the location of the target object. Then, the agent took a detour: she ascended a ladder to the platform, taking a position where she had clear view on which object was in which hole. This detour, however, was hard to justify in this condition as the agent already had access to this information. Hence, we label this as the *Inefficient* condition.

In the other, *Efficient* condition, when the agent came out of the house, only the identical, indiscernible parts of the objects were visible, after which the objects jumped to the boxes. That is, the location of the goal-object fell within a range of possibilities from the perspective of the observed agent (location A or B), but not from the perspective of the infant. This state of limited information made two mutually exclusive actions available to the agent (go to location A or B) to have exactly the opposite utilities in the two equally probable states, rendering the average expected utilities ultimately equal for both possible actions. That is, given the perspective of the agent and assumption (1), participants could attribute uncertainty to the agent regarding object location. Then, again, the agent made the very same detour as she did in the *Inefficient* condition, which was now justifiable by the assumptions (1) and (2) of the proposed naïve theory.

We paused the animation at the end of the detour in both conditions and measured infants' looking time and pupil dilation. We expected longer looks in the *Inefficient* detour condition, where the detour could not be justified by the agent's uncertainty about the location of the target object. For pupil size, we predicted a difference in the same direction during the detour.

Participants

A total of 24 full-term, typically developing 14-month-old infants were included in the analysis (N = 24; 11 girls; M_{age} = 14 months, 18.2 days; SD_{age} = 7.5 days). All infants were recruited in Budapest for research purpose, on the basis of birth records.

24 additional infants were tested but not included due to crying or fussiness (17), experimenter error (5) and equipment failure (2). Infants received a small toy gift for their participation at the end of the experiment.

All data collection methods and procedures were approved by the United Ethical Review Committee for Research in Psychology, in Hungary.

Materials and Design

Animations were created and rendered in Maya®, a 3D computer animation software (Autodesk Inc.). The resulting stimuli were displayed on the inbuilt monitor of a Tobii T60 XL eye tracker device (51.1 x 32.3 cm, 1920 x 1200 pixels resolution), controlled by PyHab (Kominsky, 2019), a stimulus presentation system for PsychoPy (Peirce et al., 2019). Additional PyGaze (Dalmaijer et al., 2014) scripts were used to control the eye tracker and record its output. Eye tracker data was recorded with 60 Hz sampling rate.

First, infants saw six familiarization trials of the agent approaching the target object over the distractor. During these trials the location of the two objects (left or right) alternated in a fixed AABBBA pattern for all subjects. The target object was always the duck for each participant.

Then, infants were presented with both versions of the second type of familiarization trials (with either the target or the distractor object hiding in the boxes), two times each, in either an ABAB or ABBA order.

Finally, all infants were presented with both test conditions (*Efficient* and *Inefficient*) once, but their order and the location of the target and distractor objects were counterbalanced across subjects. Attention-getters, depicting colorful geometrical figures on a gray background, were inserted between each trial, except in the first familiarization phase, where they were inserted after every second trial.

All stimuli, data and analysis scripts can be found at https://osf.io/vm78t/

Procedure

Following consent procedures, the experiment took place in a sound-proof room with dimmed lights. Participants were seated on their caregiver's lap, at about 60 cm distance from the display. Caregivers wore opaque glasses that prevented them from seeing the stimuli. They were instructed to keep the child seated on their laps and not to interact with them. The experimenter was seated behind a curtain and monitored infants' behavior from a separate screen via a video camera.

After a few seconds of eye-tracker calibration, infants watched six and four trials of the two types of familiarization phases, respectively (trials were fixed-duration, 8 and 12 seconds long). Then, infants were exposed to two test trials (26 seconds each). In both test trial, at the end of each movie, the last frame remained on, and looking times were recorded until the participants looked away for 2 consecutive seconds or have looked at the screen for a total of 30 seconds.

The entire session was video recorded. To extract exact looking time durations for data analysis, looking behavior was coded offline with PsyCode (Gervain, Bonatti, & Filippin, 2009).

Pupillometric analysis

Preprocessing was based on the methodology of Jackson and Sirois (2009). Pupil estimates from each eye were regressed onto estimates from the other eye and linear interpolation was also conducted in order to fill in missing samples. Before, to remove artifacts, a Hampel-filter was applied to the mean of the right and left pupil data to detect and remove outliers. Specifically, if a sample differed from the median of its time window (composed of the sample and its six surrounding samples, three per side) by more than three standard deviations, it was replaced with the median.

Pupil dilation change was baseline-corrected according to the following rationale. First, we wanted to be sure that any potential luminosity difference between conditions could not affect the dilation of the pupil at the relevant time window of the analysis. Second, we also wanted to identify the earliest point in time when subjects in both conditions tended to notice the agent's action and moved their gazes towards it, which would also meet the previous criterion, as at this point their eyes would be directed at the same area of the display and their pupils would be exposed to identical light sources. Therefore, that first 300-ms-long epoch within the duration of the agent's detour was identified in which the difference between conditions in terms of both X and Y mean gaze positions has converged below a threshold of 60 pixels, which was approximately 1.5° visual angle in our setup, corresponding to the spatial resolution of the eye-tracker device used in the study (Holmqvist, 2017). We assumed that 300 ms would be sufficient for the physiological response of the pupil to adapt to the lighting conditions of the stimuli. This analysis yielded a baseline period from the 1.728^{th} to the 2.028^{th} seconds of the detour ($M_{x\text{-difference}} = 59.2$ pixels, $M_{y\text{-difference}} = 42.2$ pixels). This period corresponded to the agent's jump towards the ladder, and our subsequent period of interest covered the remaining 4 seconds of the detour, during which the agent ascended the ladder and looked to the objects.

For each trial and for each participant, a baseline pupil diameter was computed by averaging the data samples in this 300 ms, then the baseline-corrected pupil diameter was calculated as the difference between raw pupil diameter and this average.

Results

All statistical analyses were conducted in R (R Core Team, 2013). Our main dependent measure for looking behavior was log-transformed looking time (based on Csibra et al., 2016) but descriptive statistics and plots feature raw values for ease of interpretation (see Figure 2). On average, infants looked longer in the *Inefficient* (M = 18.22 s, SD = 9.57 s) condition, compared to the *Efficient* (M = 13.10 s, SD = 9.95 s) condition, and this difference turned out to be statistically significant: $t(23) = -3.18$, two-tailed, $p = .004$, $d = 0.65$. Furthermore, the estimated Bayes Factor ($BF_{10} = 10.34$) suggests that the data were 10.34:1 in favor of the alternative hypothesis (a noninformative Jeffreys prior was placed on the variance of the normal population, while a Cauchy prior was placed on the standardized effect size, as suggested by Rouder et al., 2009).

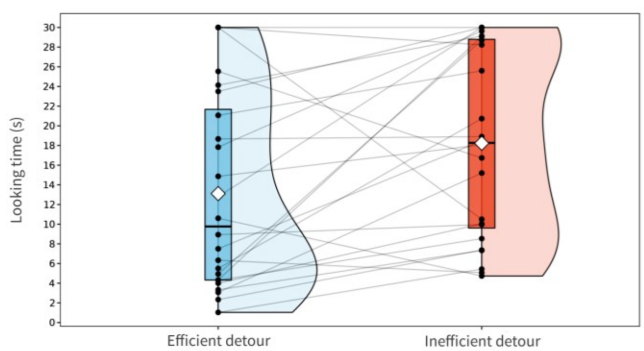

Figure 2: Boxplots of raw looking times in seconds for test trials (*N = 24*). White diamonds indicate means, bold horizontal lines indicate medians, boxes indicate interquartile ranges, half-violin density plots represent the distributions, linked black dots indicate individual data points across conditions.

To analyze pupil dilation, mean pupil size was computed by participant and condition. The mean baseline-corrected pupil size during the detour differed in the same direction as looking time, with larger pupil diameters in the *Inefficient* condition: $M_{Efficient} = 0.19$, $SD_{Efficient} = 0.24$, $M_{Inefficient} = 0.27$, $SD_{Inefficient} = 0.19$ (see Figure 3 for the graphical depiction of this segment). To test for a significant difference between the two conditions, while controlling for multiple comparisons, a permutation analysis was conducted, based on a technique that was originally developed for EEG-data analysis (Maris & Oostenveld, 2007), but was later adapted to eye-data analysis as well (e.g. Oakes et al., 2013).

This method, first, selects those timepoints in which the pupil-dilation difference between the two conditions is below a critical p-value (.05, two-tailed). Second, it clusters these together on the basis of temporal adjacency. Third, it repeatedly shuffles the data between conditions, and, for each iteration (here, we used 1000 iterations), it carries out the first two steps on this shuffled data as well. Finally, for each original cluster, it calculates the proportion of random clusters that resulted in a larger cumulative t-statistic than the observed one. If this proportion is smaller than the critical alpha-level (.05), it can be concluded that the data in the two conditions are significantly different.

In our case, however, this analysis revealed no such significant difference.

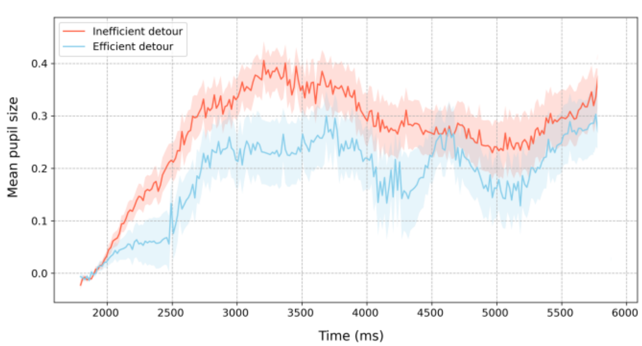

Figure 3: The temporal dynamics of average baseline-corrected pupil size in the two conditions after the offset of the baseline period. Lighter areas around the mean indicate standard error.

Discussion

We presented 14-month-old infants with actions that were inefficient with respect to the agent's demonstrated instrumental goal but could be justified as information-seeking under a naïve theory with the following content: (1) having no sufficient access to goal-relevant variables gives rise to uncertainty; (2) agents will engage in actions towards the sub-goal of reducing this task-related uncertainty. We found that infants looked longer when the same action could *not* be justified based on these assumptions. Since this variable is usually interpreted as scaling with surprise, this evidence is compatible with the hypothesis that infants can attribute information-seeking goals to actions that comply with the principles (1) and (2) above.

Interestingly, we found no significant difference in terms of pupil-dilation between the two conditions. Our time window of interest corresponded to 4 seconds of the agent's detour. Were we to extend this window to the period after the detour, we might have found such a difference, assuming that making sense of the inefficient detour would correspond to higher cognitive load, one correlate of which was proposed to be pupil dilation (e.g. van der Wel, & van Steenbergen, 2018). However, analyzing later segments would be misleading in our case, as we could not control for the idiosyncratic eye-movements of subjects in response to the still display that followed the detour.

Notably, there might be an alternative interpretation of our findings. Since in the *Inefficient* condition the agent had initial knowledge of the target object's location, it is possible that infants had a specific expectation of direct approach, which was then violated. In contrast, they might have had no particular expectation to be violated in the *Efficient* condition, yielding relatively shorter looking times in the latter. However, it is worth mentioning that, under this reading, we can still claim that infants were able to assign a goal to the agent, take her perspective, attribute her knowledge based on this perspective, and predict her behavior given this attribution, all within one task.

Note that we ensured that the agent *did* see some portion of the objects in both test conditions, and we only manipulated the visibility of those object parts that were relevant for identity disambiguation. Therefore, it is not obvious how simple computations based on line of sight or registration (e.g. Butterfill & Apperly, 2013) could account for the observed differences. Rather, the looking time results suggest that infants had different interpretations of the events depending already on this slight visual manipulation, indicating relatively sophisticated perspective-taking abilities.

That said, we are planning a control study to rule the above detailed alternative interpretation out. It is worth to point out that including the remaining two combinations of the independent variables, that is, visible or invisible hiding always followed by direct approach (instead of detour), would not help to falsify this alternative. Rather, we aim to control for the possible confound that the visibility of the target object's hiding introduces an unilateral expectation of direct approach in the two conditions.

Therefore, in the forthcoming study, we intend to keep the design the same, except we introduce one additional non-target object (again, partly identical to the other objects). Then, only the two non-target objects will appear in both test conditions, hiding visibly or invisibly, as in the original study. Thus, if the objects are visible, infants should no longer expect direct approach, as the target is not in the scene. However, we would still predict them to look longer in the condition with visible hiding, in response to the agent's detour, as the detour remains inefficient – only not because the location of the target is already known, but because it is known to be *not* in the scene, therefore the detour provides no information about its whereabouts. In contrast, the detour remains efficient in the condition with invisible hiding, as, from the agent's point of view, the target could still be in one of the hiding locations.

Even if follow-up work generates stronger evidence for the assumptions (1) and (2), the third principle, stating that information will only be gathered if its expected rewards are higher than the cost of the actions leading to it, remains to be tested. Such an assumption would make the naïve theory conceptually continuous with a more general intuitive model of goal-directed behavior and psychology, previously proposed in the literature, in which agents are expected to maximize utility (Baker et al., 2017; Gergely & Csibra, 2003; Jara-Ettinger et al., 2016).

More formally, adapting the analysis of Chater et al. (1998), the expected value of an interested information-seeking action *IS* with the cost *C* and different possible epistemic *results*, given the agent's initial knowledge *K*, can be given by the equation:

$$EV(IS) = \left[\sum_{result} P(result|K) \times EV(result)\right] - C(IS)$$

Where *EV(result)*, that is, the expected value of learning a result can be given by the difference between the maximum expected utility of choosing with and without the information it provides. This can be captured by an equation where $\mathcal{A}$ is the set of *a* actions and *s* are the possible world states that are considered in a given decision problem, and *U* is the utility of the outcome of action *a* in a possible world state *s*:

$$EV(result) = [max_{a\in\mathcal{A}} \sum_{s} P(s|result) \times U(a, s)] - [max_{a\in\mathcal{A}} \sum_{s} P(s) \times U(a, s)]$$

If such a model is included in young observers' naïve utility calculus, several predictions should follow. Ceteris paribus, the value of information-seeking should be seen higher, (i) the lower the cost of the action, (ii) the larger is the number of the goal-relevant possible world states it disambiguates, and (ii) the higher is the difference between the utilities of possible outcomes it helps to decide among. We will aim to test these predictions in further studies. Notably, and congruent with this theoretical framework, a recent study of Aboody, Zhou & Jara-Ettinger (2021) has shown that preschoolers are ready to infer others epistemic states and corresponding valuation of information based on the cost they are willing to pay for acquiring it.

While the question of how infants interpret information-seeking might be considered interesting in its own right, addressing it could also provide an opportunity to study how decision problems, alternative possibilities and information value are understood early in development. These capacities might correspond to the preconditions of other remarkable cognitive achievements as well.

For example, relevance theory (Sperber & Wilson, 2002) aims to provide a solution to the interpretational problem of communication by positing that hearers infer the intended meaning of an utterance by some sort of relevance-based comprehension procedure: computing cognitive benefits via following the path of least effort and testing interpretative hypotheses in order of accessibility until the expectations of relevance are satisfied. The authors claim that this algorithmic procedure can be successful because speakers reliably communicate relevant information. Similarly in this aspect, the Rational Speech Act model of linguistic communication captures the speaker's need to be informative, by assuming that the her utility scales with the social benefit of providing epistemic help to the listener (Goodman & Frank, 2016). However, inferring what information is relevant for others is not a trivial task. One could speculate that recognizing an agent's instrumental decision problem (and the corresponding hypothesis space) would allow observers not only to anticipate its likely course of information-directed behavior, but also to estimate what signals from their side could resolve her problem.

References

Aboody, R., Zhou, C., & Jara-Ettinger, J. (2021). In Pursuit of Knowledge: Preschoolers Expect Agents to Weigh Information Gain and Information Cost When Deciding Whether to Explore. *Child Development*.

Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour, 1*(4), 598.

Baldwin, D., Loucks, J., & Sabbagh, M. (2008). Pragmatics of Human Action. In T. F. Shipley & J. M. Zacks (Eds.), *Understanding Events: From Perception to Action*. OUP.

Bennett, J. (1976). *Linguistic Behaviour.* CUP.

Butterfill, S. A., & Apperly, I. A. (2013). How to construct a minimal theory of mind. *Mind & Language, 28*(5), 606–637.

Chater, N., Crocker, M. J., & Pickering, M. J. (1998). The rational analysis of inquiry. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition*. OUP.

Coenen, A., Nelson, J. D., & Gureckis, T. M. (2019). Asking the right questions about the psychology of human inquiry: Nine open challenges. *Psychonomic Bulletin & Review, 26*(5), 1548–1587.

Csibra, G. (2008). Goal attribution to inanimate agents by 6.5-month-old infants. *Cognition, 107*(2), 705–717.

Csibra, G., Hernik, M., Mascaro, O., Tatone, D., & Lengyel, M. (2016). Statistical treatment of looking-time data. *Developmental Psychology, 52*(4), 521.

Dalmaijer, E.S., Mathôt, S., & Van der Stigchel, S. (2014). PyGaze: an open-source, cross-platform toolbox for minimal-effort programming of eye tracking experiments. *Behavior Research Methods, 46*, 913–921.

Davidson, D. (1963). Actions, Reasons, and Causes. *The Journal of Philosophy, 60*(23), 685–700.

Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: the naïve theory of rational action. *Trends in Cognitive Sciences, 7*(7), 287–292.

Gervain, J., Bonatti, L., & Filippin, L. (2009). *The PsyCode Software Manual.* http://psy.ns.sissa.it/

Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences, 20*(11), 818–829.

Gottlieb, J., & Oudeyer, P.-Y. (2018). Towards a neuroscience of active sampling and curiosity. *Nature Reviews Neuroscience, 19*(12), 758–770.

Holmqvist, K. (2017). Common predictors of accuracy, precision and data loss in 12 eye-trackers. In *The 7th Scandinavian Workshop on Eye Tracking*.

Jackson, I., & Sirois, S. (2009). Infant cognition: going full factorial with pupil dilation. *Developmental Science, 12*(4), 670–679.

Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The Naïve Utility Calculus: Computational Principles Underlying Commonsense Psychology. *Trends in Cognitive Sciences, 20*(8), 589–604.

Kominsky, J. F. (2019) PyHab: Open-Source Real Time Infant Gaze Coding and Stimulus Presentation Software. *Infant Behavior & Development, 54,* 114–119.

Liszkowski, U., Carpenter, M., & Tomasello, M. (2008). Twelve-month-olds communicate helpfully and appropriately for knowledgeable and ignorant partners. *Cognition, 108*(3), 732–739.

Liu, S., & Spelke, E. S. (2017). Six-month-old infants expect agents to minimize the cost of their actions. *Cognition, 160*, 35–42.

Liu, S., Ullman, T. D., Tenenbaum, J. B., & Spelke, E. S. (2017). Ten-month-old infants infer the value of goals from the costs of actions. *Science, 358*(6366), 1038–1041.

Luo, Y., & Baillargeon, R. (2007). Do 12.5-month-old infants consider what objects others can see when interpreting their actions? *Cognition, 105*(3), 489–512.

Luo, Y., & Johnson, S. C. (2009). Recognizing the role of perception in action at 6 months. *Developmental Science, 12*(1), 142–149.

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG-and MEG-data. *Journal of Neuroscience Methods, 164*(1), 177–190.

Moll, H., & Tomasello, M. (2007). Cooperation and human cognition: the Vygotskian intelligence hypothesis. *Philos. Trans. R. Soc. Lond., B, Biol. Sci., 362*(1480), 639–648.

Oakes, L. M., Baumgartner, H. A., Barrett, F. S., Messenger, I. M., & Luck, S. J. (2013). Developmental changes in visual short-term memory in infancy: Evidence from eye-tracking. *Frontiers in Psychology, 4,* 697.

Peirce, J. W., Gray, J. R., Simpson, S., MacAskill, M. R., Höchenberger, R., Sogo, H., Kastman, E., Lindeløv, J. (2019). PsychoPy2: experiments in behavior made easy. *Behavior Research Methods*.

R Core Team (2013). *R: A language and environment for statistical computing.* https://www.r-project.org/

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review, 16*(2), 225–237.

Shannon, C. E. (1948). A Mathematical Theory of Communication. *The Bell System Technical Journal, 27*(3), 379–423.

Sperber, D., & Wilson, D. (2002). Pragmatics, Modularity and Mind-reading. *Mind & Language, 17*(1&2), 3–23.

Surian, L., Caldi, S., & Sperber, D. (2007). Attribution of beliefs by 13-month-old infants. *Psychological Science, 18*(7), 580–586.

Tauzin, T., & Gergely, G. (2018). Communicative mind-reading in preverbal infants. *Scientific Reports*, *8*(1), 1–9.

van der Wel, P., & van Steenbergen, H. (2018). Pupil dilation as an index of effort in cognitive control tasks: A review. *Psychonomic Bulletin & Review, 25*(6), 2005–2015.

Vouloumanos, A., Onishi, K. H., & Pogue, A. (2012). Twelve-month-old infants recognize that speech can communicate unobservable intentions. *Proceedings of the National Academy of Sciences, 109*(32), 12933–12937.

Whiten, A., & Byrne, R. W. (1997). *Machiavellian Intelligence II: Extensions and Evaluations*. CUP.

Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition, 69*(1), 1–34.