



HAL
open science

Actions Knowledge in Sight: Toddlers Plan Efficient Epistemic Actions by Anticipating Learning Gains

Marie Aguirre, Mélanie Brun, Auriane Couderc, Anne Reboul, Philomène Senez,
Olivier Mascaro

► **To cite this version:**

Marie Aguirre, Mélanie Brun, Auriane Couderc, Anne Reboul, Philomène Senez, et al.. Actions Knowledge in Sight: Toddlers Plan Efficient Epistemic Actions by Anticipating Learning Gains. *Cognitive Science*, 2022, 46 (2), <10.1111/cogs.13103>. <hal-03624303>

HAL Id: hal-03624303

<https://hal.science/hal-03624303v1>

Submitted on 30 Mar 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

**Knowledge in Sight: Toddlers Plan Efficient Epistemic Actions by
Anticipating Learning Gains**

Marie Aguirre^{1*}, Mélanie Brun¹, Auriane Couderc¹, Anne Reboul², Philomène Senez¹, Olivier Mascaró¹

¹ Université de Paris, CNRS, Integrative Neuroscience and Cognition Center, F-75006 Paris, France

² Laboratory of Cognitive Psychology, UMR 7290, CNRS and Aix-Marseille University, Marseille, France

Author Note

The authors have no conflicts to disclose.

Acknowledgements

We express our gratitude towards the daycare centers and all the families that participated in this project, as well as towards Shelby Wicklacz. We also wish to thank Nicole McNeil, Jonathan Kominsky, and two anonymous reviewers for their comments on previous versions of this manuscript. This research was supported by a grant from the Agence Nationale pour la Recherche (FoundTrust, ANR-21-CE28-0017-01) and by a grant from the Fyssen foundation (Infants' Representation of Informativeness) to OM. This study contributes to the IdEx Université de Paris ANR-18-IDEX-0001.

* Corresponding author: Marie Aguirre, CNRS/Université de Paris, Integrative Neuroscience and Cognition Center UMR 8002, 45 rue des Saints Pères, 75014, Paris, France. Contact: marie.aguirre77@gmail.com

Ethics Approval Statement

This research followed the guidelines of the Declaration of Helsinki and it was approved by the local board of daycare centers and by an independent ethical committee for

TODDLERS PLAN EFFICIENT EPISTEMIC ACTIONS

bio-medical research (CPP Sud-Est II, IRB: 00009118). All participants' parents/guardians gave written informed consent prior to inclusion in the studies.

Abstract

Anticipating the learning consequences of actions is crucial to plan efficient information-seeking. Such a capacity is needed for learners to determine which actions are most likely to result in learning. Here, we tested the early ontogeny of the human capacity to anticipate the amount of learning gained from seeing. In Study 1, we tested infants' capacity to anticipate the availability of sight. Fourteen-month-old infants ($N = 72$) were invited to search for a toy hidden inside a container. The participants were faster to attempt at opening a shutter when this action allowed them to see inside the container. Moreover, this effect was specifically observed when seeing inside the container was potentially useful to the participants' goals. Thus, infants anticipated the availability of sight, and they calibrated their information-seeking behaviors accordingly. In Studies 2-3, we tested toddlers' capacity to anticipate whether data would be cognitively useful for their goals. Two-and-a-half-year-olds ($N = 72$) had to locate a target character hidden among distractors. The participants flipped the characters more often, and were comparatively faster to initiate this action when it yielded access to visual data allowing them to locate the target. Thus, toddlers planned their information-seeking behaviors by anticipating the cognitive utility of sight. In contrast, toddlers did not calibrate their behaviors to the cognitive usefulness of auditory data. These results suggest that cognitive models of learning guide toddlers' search for information. The early developmental onset of the capacity to anticipate future learning gains is crucial for active learning.

Keywords: active learning, naïve epistemology, metacognition, perception, informativeness.

**Knowledge in Sight: Toddlers Plan Efficient Epistemic Actions by
Anticipating Learning Gains**

1 Background

Humans are hungry for information. They are attracted by novelty, and engage in independent exploration and in innovative hypothesis testing (Berlyne, 1966; Kidd & Hayden, 2015; Pelz & Kidd, 2020; Pisula, 2009; Wilson, 2000). Such behaviors shape societies, institutions, and daily lives for a good reason: they support learning (Kang et al., 2009). Information seeking is perhaps one of the most general functions of cognitive mechanisms. One can increase and update one's knowledge simply by processing incoming stimuli, without actively searching for useful information. This strategy may yield learning, up to a point. Yet, this type of passive learning implies a huge cost of opportunity. Indeed, in many cases, an organism can learn a lot more by actively searching for useful information in its environment. One way to plan information seeking in a flexible and efficient manner consists in representing the process of learning itself, in order to select actions optimizing learning. Here, we test the early ontogeny of this capacity.

1.1 Anticipating Information Gains Is Crucial to Plan Information-Seeking

Behaviors whose apparent function is to gather information are observed in a wide range of species, from simple invertebrates to humans (e.g., Blanchard et al., 2015; Butler, 1954; Degen et al., 2015; Glickman & Sroges, 1966; Kidd & Hayden, 2015; Vergassola et al., 2007). Yet, the

function of searching for information in an optimal fashion is likely to be implemented by very different cognitive mechanisms depending on the specie in which it is observed. Efficient active learning can often be achieved through evolved fixed behaviors triggered by the lack of specific information (Carruthers, 2018). These information-seeking heuristics do not require learners to anticipate if and when an action will result in learning. For instance, *C. Elegans*, a nematode with only a few hundred neurons, has been found to alternate patterns of movements in a way that optimizes information gains about the location of food (Calhoun et al., 2014). While these foraging mechanisms are efficient, it is likely that they involve no representation of learning, or of ignorance, on behalf of the nematodes. *C. Elegans'* foraging mechanisms are also probably quite inflexible: They have evolved to operate in a specific environment, and to solve a very specific problem (locating food). It is doubtful that these foraging mechanisms would allow *C. Elegans* to address novel questions such as, for instance, determining why food is abundant at a particular location or how it arrived there.

In contrast, humans can flexibly discover novel ways to search for information by representing and anticipating the learning outcomes of their actions. This capacity is crucial for searching for information and for teaching others. It also supports creative hypothesis testing, a cornerstone of human learning and science (e.g., Bass et al., 2019; Bridgers et al., 2020; Schulz, 2012; Shafto et al., 2012, 2014). Such a flexible and creative information seeking can be achieved by representing and anticipating the process of learning itself. By anticipating the potential learning gains resulting from future actions, a cognitive system can select the actions

that are most likely to yield knowledge. Thus, anticipating future learning gains is crucial to plan information search behaviors creatively. Here, we investigate the early ontogeny of this capacity to anticipate the learning consequences of actions, focusing on the precocious ontogeny of humans' sensitivity to the (i) availability and (ii) cognitive utility of sight.

1.2 Sight as an Intuitive Model of Learning

Sight is associated with knowledge in many cultures (Bloch, 2008), and since Plato and Aristotle, it has been a central model and metaphor for learning in Western philosophy (Synnott, 1992). The importance of sight in naïve and formal epistemology has to do, we suspect, with the central role of vision in human learning, and in humans' core representation of knowledge. In a seminal study by Call and Carpenter (2001), 2.5-year-olds had to find a reward hidden in one of three opaque tubes. In some trials they could see in which tube the reward was hidden (seen condition) and in other trials they could not see it (unseen condition). Most children crouched to look into the tubes before choosing one more frequently in the unseen condition than in the seen condition (see also Neldner et al., 2015). In addition to human children, all four great ape species and some species of monkeys engage in efficient information-seeking behaviors by trying to see (Beran et al., 2013; Call, 2010; Call & Carpenter, 2001; Marsh & MacDonald, 2012; Rosati & Santos, 2016). In short, children use efficient behaviors to look for visual information, a capacity that is ontogenetically and phylogenetically ancient, and perhaps shared in part with non-human primates. We capitalize on this evidence to test young children's capacity to plan their search for visual information and to anticipate its learning outcomes.

1.3 Cognitive Models of Learning

Infants are sensitive to the amount of learning gained from sight from an early age. They look longer at stimuli that are unexpected (Stahl & Feigenson, 2015), of intermediate complexity (Kidd et al., 2012), or yield maximal learning (Poli et al., 2020). Thus, infants can use the properties of visual data that are already accessible to them to determine whether they should continue looking at a scene (exploit) or should start looking elsewhere (explore). However, in many cases, learners need to assess the learning consequences of future perceptual inputs, before they can access them. This capacity is crucial for planning goal-directed information search, e.g., when deciding what to do in order to access a specific piece of information that one is missing.

The planification of goal-directed search for information is likely to rely on a cognitive model of learning allowing individuals to anticipate the epistemic consequences of actions. This model should (i) be sensitive to the availability of data (e.g., while sight is a basis for learning, it is not always available, for instance, when vision is obstructed by opaque objects), and (ii) aim to maximize the benefits and minimize the costs of processing data by collecting pieces of evidence with high, if not optimal, cognitive utility (e.g., Oaksford & Chater, 1994; Sperber & Wilson, 2001; Steyvers et al., 2003). Thus, we investigated infants' and toddlers' capacity to anticipate the availability and cognitive utility of sight.

1.4 Anticipating the Availability of Data

Tracking the availability of sight implies detecting the situational factors that specifically enable or impede seeing. For example, sight is influenced by opaque materials, whereas hearing is not. Thus, a sensitivity to the way opaque material might obstruct sight is instrumental to anticipate the availability of visual data. By five years of age, children adjust

their information search behaviors to what constraints seeing. In one study, preschoolers had to locate a sticker hidden under an opaque or transparent cup. The participants could observe by peeking while an experimenter placed the sticker for the trial. By five years of age, the children peeked for longer in the opaque than the transparent condition (Iwasaki et al., 2020).

Very little is known about the development of humans' capacity to anticipate the availability of sight prior to five years of age. Yet, infancy research suggests that humans' sensitivity to what constraints seeing emerges early. Young children, infancy onwards, are sensitive to what others can see, and use this ability appropriately during social interactions (Choi et al., 2018; Liszkowski et al., 2007; Luo & Baillargeon, 2007; Moll & Tomasello, 2006; O'Neill, 1996; Sodian et al., 2007; Southgate et al., 2007). For example, Choi, Mou, and Luo (2018) reported that three-month-old infants differentially processed the reaching actions of an agent depending on what she could see. One-year-old infants begin to take into account opaque barriers when following gaze (D'Entremont & Morgan, 2006; Meltzoff & Brooks, 2008; Moll & Tomasello, 2004). By two years of age, toddlers manipulate others' visual access by hiding objects (Flavell et al., 1978; Mascaro et al., 2017), or making them visible (Mascaro et al., 2019). Furthermore, toddlers develop some sensitivity to the disabling role of opaque materials for their own sight, a capacity that they use to determine what is visible to others (Király et al., 2018; Meltzoff & Brooks, 2008; Senju et al., 2011). In short, infants and toddlers track what others can see and are sensitive to the opacity of materials. Study 1 capitalizes on these phenomena to investigate infants' sensitivity to the constraints that specifically regulate access to visual data. We probe infants' capacity (i) to anticipate the availability of sight (depending on the opacity of a window), and (ii) to adjust their information-seeking behaviors accordingly.

1.5 Anticipating the Cognitive Utility of Data

Although sight is a central basis for learning, seeing does not always lead to knowing. For example, sight is insufficient for discriminating between two distinct individuals who look identical. Thus, the appropriate use of sight in learning must be sensitive to its cognitive utility in a given context. As vision is a central source of learning in primate taxa, acknowledging that sight is not always a relevant source of data is not trivial. In fact, children and adults often overestimate the knowledge that results from gaining visual access to an object (Robinson et al., 1997; Wang et al., 2014).

Yet, previous studies on the ontogeny of learning actions suggest that by four years of age, children anticipate the relevance of data gathered from specific perceptual sources. Four-year-old preschoolers adjust their reliance on specific perceptual sources according to what they want to learn (Fitneva et al., 2013; O'Neill & Chong, 2001; Pillow, 1993; Robinson et al., 2008), and this capacity develops further during preschool years (O'Neill et al., 1992; Perner & Ruffman, 1995; Robinson et al., 1997). To illustrate, when given the choice between looking and asking someone to learn about a character's properties, four-year-olds are more likely to choose looking to learn about the visible properties (e.g., a character's hair color) than to learn about the invisible properties (e.g., what makes a character sick) (Fitneva et al., 2013). Thus, four-year-olds are more likely to seek visual information when it is relevant to fill their knowledge gaps.

Studies 2 and 3 build upon this evidence to test comparable abilities in much younger participants, focusing on their capacity to track the cognitive utility of a visual input. We focused on informativeness, a dimension of cognitive utility, defined as the capacity of a stimulus to reduce uncertainty about a set of hypotheses (Frank & Goodman, 2012; Steyvers

et al., 2003), uncertainty being quantified using the standard entropy measure from information theory (Shannon, 1948).

1.6 Operationalization Principle

Studies show that children are more likely to engage in information-seeking behavior when they are ignorant rather than knowledgeable about the location of target object that they have to find (e.g., Call & Carpenter, 2001; Ruggeri et al., 2019). However, it has been suggested that in such experimental paradigms, the participants might be simply retrieving the target object directly when they know where it is, without exploring any other option or engaging in any other behavior. In contrast, when they ignore where the target object is, the participants might be more likely to engage in other behaviors than trying to retrieve the target, including (i) fixed behavioral routines or (ii) general exploratory behaviors (for variants of this argument, see Carruthers, 2008; Crystal & Foote, 2011; Hampton, 2009; Marsh, 2014; Perner, 2012).

Thus, since we wanted to focus on flexible learning, in our experimental conditions, the participants were *always* ignorant about the same piece of information. We modulated the amount of learning benefits resulting from the exact same action across conditions. We measured the latency with which our participants performed this action to have an estimate of our participants' capacity to anticipate the learning consequences of their actions. Thus, if our participants use a fixed information-seeking heuristic, they should behave similarly in all conditions. In contrast, if our participants plan their information search by anticipating the learning consequences of their actions, they should be faster to engage in a given behavior when it yields learning benefits. To probe our participants' anticipation of the learning consequences of their actions, we collected measures of children's latency to engage in a

specific action, depending on its anticipated informativeness. Measures of latency are appropriate to probe cognition from infancy onwards, and they have been successfully used to collect data about children's sensitivity to informational uncertainty by 12 months of age (Kim et al., 2020). By collecting measures of latency to engage in relevant information-search behaviors, we could investigate our participants' capacity to assess the learning consequences of future perceptual inputs, before they could access them. This measure of anticipation enables us to cast light on young children's capacity to represent future learning gains.

Study 1 investigated infants' capacity to anticipate the availability of sight. The participants enrolled in the experimental group were invited to search for a toy hidden inside a box. In the transparent condition, the participants could see inside the box by opening a shutter covering a transparent window. In the opaque condition, the window behind the shutter was opaque. Thus, in the opaque condition, opening the shutter did not allow infants to see inside the box. We measured the latency with which the participants tried to open the shutter. This measure allowed us to probe infants' capacity to anticipate the learning consequences of opening the shutter, before they completed this action. If infants anticipate the future availability of visual data resulting from opening the shutter, they should be faster to perform this action when it allows them to see inside the box (in the transparent condition) than when it does not (in the opaque condition). In Study 1, we also tested infants on a control condition in which no toy was hidden inside the box during the test trial. This condition allowed us to evaluate whether infants would still prefer to look inside the box when information about its content was irrelevant for their goals. This control condition also served to rule out the possibility that infants might be faster at attempting to open the shutter in the transparent condition than in the opaque condition just because of a visual preference

for transparent over opaque windows. Since in the control condition seeing inside the box was irrelevant to the infants' goals, we expected that infants would no longer be faster at attempting to open the shutter when the window was transparent (rather than opaque).

Studies 2 and 3 investigated toddlers' capacity to anticipate the cognitive utility of sight. The participants had to locate a target character hidden among others; we assumed that toddlers should to engage in actions on the characters more often and comparatively faster when these actions yielded access to visual data that was sufficiently informative to identify the target.

2 Study 1

2.1 Method

2.1.1 Participants

Study 1 required the participants to search for an object hidden in a container. Therefore, we chose to test fourteen-month-olds because by this age infants' capacity to adjust their search behaviors to their beliefs about the presence of one or a few objects hidden in boxes is well established (Cacchione et al., 2013; Feigenson & Carey, 2003). Seventy-two fourteen-month-old infants were enrolled either in the experimental group (transparent condition: $n = 18$; $M = 13.8$ months, $range = 13-15$ months, $SD = 0.86$; opaque condition: $n = 18$; $M = 13.8$ months, $range = 13-15$ months, $SD = 0.68$) or in the control group (transparent condition: $n = 18$; $M = 13.7$ months, $range = 13-15$ months, $SD = 0.70$; opaque condition: $n = 18$; $M = 13.4$ months, $range = 12-14$ months, $SD = 0.62$). Each participant was tested only once, in a single condition. Sixteen additional participants were excluded from the analysis for the following reasons: refusal to cooperate (3), unwillingness to finish the experiment (3), caregiver interference (1), and technical failure (9). For all the Studies (1-3), we recruited

participants by contacting a random sample of daycare centers in the /MASKED FOR BLIND REVIEW/ area. Daycares were selected to be within driving distance from our laboratory. For each daycare whose board validated the study, we invited families to participate in the project by sending an information letter to all the children falling within our target age ranges.

We used samples as large as possible, given the recruitment opportunities. A compromise power analysis performed using G*Power (Faul et al., 2007) revealed that the resulting sample sizes ($n = 18$ per group) yielded an implied power equal to .83 for between-group comparisons using Mann-Whitney U test ($d = .8$, $\alpha = .05$).

2.1.2 Materials and set-up

Infants were tested in a quiet room and were accompanied by a caregiver who was instructed not to influence the participants' choices at any time-point during the test. The participants faced the experimenter across a table on which the testing materials were placed. The participants had to search for a small plush toy (representing a cat) placed in a black box ($31 \times 26 \times 13$ cm). The same plush toy was used during all of the test trials. The front face of the box had an opening (14.5×6.5 cm) covered by blue spandex material with a horizontal slit across its width. The back of the box had a second opening that was not visible to the participant and was covered with black fabric. The experimenter used this second opening to remove objects from the box, unbeknownst to the infants (the box was adapted from Feigenson & Carey, 2003; Starkey, 1992). On the top of the box, there was a "shutter" that could be opened with a handle. Behind the shutter, there was a window made of PVC. In the transparent condition, the PVC was transparent, making it possible for the participants to see the contents of the box through the window. In the opaque condition, the PVC was opaque, making it impossible to see the contents of the box (see Figure 1). A camera (temporal

resolution = 25 frames per second) recorded the participants' behavior. The camera was positioned behind the participants, slightly above their head and on their side, in order to record any contact between the participants' hand and the shutter's handle.

Figure 1

Familiarization. The experiment began with a familiarization phase, which was identical in all conditions. The experimenter showed the box to the participant, while saying, "*I brought a box with me.*" Then, she opened the shutter located on the top of the box, while saying, "*Look, I can open the window.*" As a result, the participant could see the PVC window (and whether it was transparent or opaque, depending on the opacity condition). Next, the experimenter encouraged the participant to open the shutter by saying, "*Can you open the window?*" The experimenter waited till the participant grabbed the handle and opened the shutter. Next, the experimenter encouraged the participant to reach inside the box and said, "*Look, I can put my hand inside the box.*" while reaching inside the box through the front opening. Then, the experimenter removed her hand and asked, "*Can you put your hand inside the box?*" The experimenter waited till the participant reached inside the box through the front opening. Next, the experimenter showed the toy to the infant while saying, "*Look, I brought a little toy! It is a cat. Do you want to pet it?*" The experimenter allowed the infant to manipulate the toy for a few seconds before taking it back, and announcing, "*Look, I can put it inside the box.*" Then, the experimenter placed the toy inside the box through the front opening and said to the infant, "*Can you give me the cat?*" The experimenter waited till the participant reached for the cat in the box. During the familiarization phase, the caregivers were allowed to encourage the participants to manipulate the box and to reach for the cat when they were shy.

Test phase. The familiarization phase was followed by a test phase that differed between the groups (control vs. experimental). At the beginning of each test trial, in the experimental group, the experimenter took the box away from the participant, and she inserted the toy into the front opening of the box with her right hand, in full view of the child. Immediately after that, she took the toy out of the box through the secret back opening with her left hand, and she left the toy hidden on her knees, under the table. She encouraged the participant to look for the toy by saying, “*Where is the cat? Can you give me the cat?*” while pushing the box towards the infant, and taking her right hand out of it. The test trial started once the box was positioned, and the experimenter’s hand was out of it. During the trial, the experimenter waited for 30 seconds while the participant was left free to search inside the box or to manipulate the shutter. If the participant did not interact with the box until approximately 10 seconds after trial onset, the experimenter repeated the prompt sentence. After 30 seconds, the trial ended. In the experimental condition, to transition to the next test trial, the experimenter placed the toy back in the box through the back opening, unbeknownst to the participants. She reached inside the box through the front opening, retrieved the toy, and showed it to the participant, while saying, “*Look, the cat was there!*” before proceeding to the next test trial. There were three consecutive test trials, each following the same procedure.

In order to assess infants’ baseline behaviors when they were not looking for information about the box’s contents, we tested a second group of participants in a control condition. In the latter, the test trials followed the same procedure as in the experimental group, except that the experimenter did not place the toy inside the box during the test trial, and asked, “*Did you see my beautiful box?*” (instead of saying, “*Where is the cat? Can you give me the cat?*”). Thus, during each test trial in the control condition, the experimenter

simply placed the box in front of the participant, and said the prompt sentence, “*Did you see my beautiful box?*” before pushing the box towards them. The test trial started from the moment the box was positioned in front of the participant. Moreover, to transition from one test trial to the next in the control condition, the experimenter took the box away from the infant while saying, “*Can I take it back for a moment?*” before proceeding to the next test trial.

2.1.3 Coding

For each test trial, the videotapes were coded offline frame by frame for 30 consecutive seconds. To measure the participant’s anticipation of the information gained by opening the shutter, we coded the delay between the beginning of each test trial and the first time the participant touched the shutter’s handle with any of her fingers, when this contact subsequently led to opening the shutter. We did not measure whether the participants reached for the handle at all, because we expected that infants would interact with the handle at some point in all conditions. In the opaque condition, opening the shutter did not yield access to any visual information about the contents of the box. In contrast, in the transparent condition, opening the shutter allowed the participant to see inside the box. Thus, we expected participants to be faster to touch the shutter’s handle in the transparent condition than in the opaque condition. When the participants did not interact with the box at all during a test trial (i.e., when they did not touch the handle at all and did not reach inside the box either), we coded the trial as missing data since in those cases the infant showed no interest in participating in the task (12 trials out of 216).

For all the Studies reported in this paper, the data were coded first by a primary coder. A second coder, unaware of the hypotheses of the study, coded 50% of the videos. The

agreement between the coders was high (average $\rho = .98$, range = .96-1, all $ps < .001$). The statistical analyses were performed on the data from the primary coder for Studies 1-3.

2.1.4 Analysis

All statistical analyses reported in this paper were two-tailed. As our data did not fulfill the assumptions of normality and homoscedasticity required for standard parametric analyses, we conducted omnibus analyses on continuous data by running robust mixed model ANOVAs implementing Johansen's general formulation of the Welch-James's test with approximate degrees of freedom, trimming of data (per 0.2), and the use of a bootstrapping methodology to better control for type I-error (Erceg-Hurn & Mirosevich, 2008; Keselman et al., 2003; Villacorta, 2017; see the Supplementary Materials for details). For all other analyses, we used non-parametric tests. Our analyses were performed with R (v. 3.6.1; R Core Team, 2020), and the R package `welchADF` (defaults number of bootstrapping: 999; v. 0.3.3; Villacorta, 2017).

2.2 Results and Discussion

A robust mixed-model ANOVA using the Welch-James ADF procedure on mean delay to reach for the handle with group (control vs. experimental) and condition (opaque vs. transparent) as between-subject factors, and trial (1, 2, or 3) as a within-subject factor revealed an interaction between group and condition ($F(1, 18.78) = 5.19, p = .037$), indicating that the effect of the windows' opacity differed across groups (control vs. experimental). The robust full factorial ANOVA revealed no other significant effect, in particular, no effect of trial ($F(2, 15.18) = 1.00, p = .41$). Thus, we performed our subsequent analyses on the average value of the delay to reach for the shutter's handle across the three test trials, computed for each participant.

Planned post-hoc analyses revealed that in the experimental group, the participants' average delay to reach for the shutter's handle was significantly shorter in the transparent condition ($M = 7.63$, $SD = 7.86$, $M_{dn} = 4.51$) than in the opaque condition ($M = 15.30$, $SD = 7.85$, $M_{dn} = 14.60$; $U = 32$, $p = .002$, Mann-Whitney U test; Figure 1). Thus, in the experimental group of Study 1, infants were faster at attempting to open a shutter when this action allowed them to gain access to visual data about an object. Since infants' reaching responses were appropriately modulated by the availability of visual information in their environment, their behavior cannot be explained by the use of a fixed information-search strategy.

Figure 2

Furthermore, in the control group, the opacity of the window had no significant effect on the average delay to reach for the shutters' handle (transparent condition: $M = 9.42$, $SD = 9.47$, $M_{dn} = 5.76$; opaque condition: $M = 10.20$, $SD = 9.49$, $M_{dn} = 6.94$; $U = 82$, $p = .72$, Mann-Whitney U test; Figure 2). This result confirms that the behavior of the participants assigned to the experimental group cannot be explained by low-level factors, such as a preference to look at transparent rather than opaque materials.

In short, the results from Study 1 suggest that infants adjust their information-search behaviors by anticipating the availability of visual access. In Study 2, we investigated whether young children are sensitive to the fact that seeing does not always lead to knowing. To this end, we manipulated whether seeing yielded cognitive benefits. The participants had to find a target card hidden among distractors. The target and distractors had the same shape and plain flip side, and had symbols on their reverse side. Crucially, in the informative condition, the symbol on the target differed from that on the distractors. In contrast, in the non-informative condition, the target and the distractors all had the same symbols on their

reverse side (see Figure 3). In the test trials in all conditions, the target and distractors were shuffled with their plain flip side visible, so that the participants no longer knew where the target was. We expected that when searching for the target, the participants would try to see the symbols more often (by flipping the cards) in the informative condition —where this action yielded cognitive benefits— than in the non-informative condition. Indeed, in the informative condition, children could gain information about the target’s location by flipping cards. In contrast, in the non-informative condition, flipping cards provided no information at all about the location of the target, since all cards were visually identical on their symbol side (for a detailed step-by-step computation of expected information gains per Condition in Studies 2 and 3, see Table S1 in the supplementary materials).

Furthermore, we also collected a measure of toddlers’ latency to act on the cards, as a way to evaluate their capacity to anticipate the learning consequences of their actions. We expected that upon being asked to locate the target, the participants would wait for a longer time without knowing what to do in the non-informative condition, since in this case, they had no way of discovering the location of the target. In short, we assumed that if toddlers anticipate the cognitive utility of sight, they should be comparatively faster to attempt at flipping cards in the informative condition than in the non-informative condition.

3 Study 2

3.1 Method

3.1.1 Participants

Study 2 required the participants to memorize the identity of a target character among several others, and to track its displacement while cues of its identity were not visible. Thus, we tested 2.5-year-old toddlers because prior to that age, the capacity to track the invisible

displacement of objects may still be fragile (Call, 2001; Collier-Baker & Suddendorf, 2006). Participants were divided into two groups of eighteen toddlers (informative condition: $M = 28.7$ months, $range = 22-35$ months, $SD = 3.83$; non-informative condition: 28.2 months, $range = 22-35$ months, $SD = 4.60$). Each participant was tested only once, in a single condition. Three additional participants were excluded from the analysis because of fussiness (1), total lack of responsiveness (1), and technical failure (1).

3.1.2 Materials and set-up

The participants were tested in a quiet room at their daycare center. They sat in front of the experimenter, across a child-size table. A familiar caregiver (from the participants' daycare center) was present during the entire experiment. The caregiver was instructed not to influence the participants' choices at any time-point. The participants had to find a character named "Charlie" (henceforth, the "target"), one of four bear-shaped wooden cards of identical shape and size (about $8 \times 9.5 \times 0.5$ cm). There were symbols on the reverse side of each card. We manipulated the informativeness of seeing the symbols across conditions by changing their distribution. In the informative condition, the symbol on the target differed from that on all the other cards. In all other respects, the cards were perceptually identical. In the non-informative condition, all the cards had the same symbol on their reverse side (thus, they were all completely identical, see Figure 3). Therefore, it was possible to identify the target by looking at the symbol on its reverse side in the informative condition, but not in the non-informative condition. We used two different symbols: a red square and a pair of stars (one blue, one yellow). For both conditions, the symbol on the target was the square for half of the participants and the pair of stars for the other half. The symbol on the other characters varied accordingly.

Figure 3

During each warm-up and test trial, the cards were placed on a cardboard tray (64 × 27 cm) kept on the child-size table. A rectangular cardboard box (10 x 15 x 3 cm) was used as a “house” in which the participants had to place the target. A camera (temporal resolution = 30 frames per second) recorded the participants’ behavior.

Presentation of the game. At the beginning of the session, the experimenter placed the cardboard box on the right side of the table and the cardboard tray in front of her on the table. She announced, “Look, I brought small bears,” while placing the cards in a row on the cardboard tray with their symbol side visible. In the informative condition, the experimenter placed the card with a symbol different from the others at the right end of the row (all positions are given from the viewpoint of the experimenter). In the non-informative condition, the procedure was the same as in the informative condition, except that the card placed at the right end of the row was visually identical to the others. The experimenter then told each character’s name to the child, by pointing successively towards each of them (moving in the row from left to right) while saying, “This is Peter. This is John. This is Marc. And this is Charlie.” After telling the characters’ names, the experimenter said, “We are going to play ‘find Charlie,’ okay?” Next, the warm-up trials started.

Warm-up trials. At the beginning of each warm-up trial, the experimenter asked, “*Where is Charlie?*” before pushing the cardboard tray toward the participants to encourage them to select one card. She added “*Can you put him in his house?*” while holding the cardboard box and pointing toward it. If the participants did not place any card in the box after approximately 15 seconds, the experimenter prompted them again by asking, “*Where is Charlie? Can you put him in his house?*” Once the participants placed the correct card in the box, the experimenter congratulated them before pulling back the cardboard tray and repositioning the cards for the next trial. When the participants placed the wrong card in the

box, they were corrected by the experimenter who said (in a neutral tone of voice), “*That’s not Charlie...*” She then placed the wrong card back to its initial location, and pointed to the correct card, while saying, “*This one here is Charlie!*” before asking again, “*Where is Charlie? Can you put him in his house?*” Two consecutive warm-up trials were conducted. At the end of the first warm-up trial, the experimenter took the target out of the box, and placed it in the second position from the left in the row, before starting the second warm-up trial. The two warm-up trials were followed by a baseline trial.

Baseline trial. The baseline trial unfolded as the warm-up trials, except that when positioning Charlie on the cardboard tray, the experimenter placed the card referred to as “Charlie” symbol-side down, in the third position from the left. The other cards remained symbol-side up, with their symbol visible to the participant, thus making it possible for the participants to locate the target. This baseline trial served to measure the participants’ baseline latency to reach for the character.

Test. At the beginning of each test trial, the experimenter placed all the cards facing down on the cardboard tray, so that the symbols were no longer visible. The experimenter said, “*Let’s mix them up!*” Next, she rearranged the cards by mixing them up quickly on the cardboard tray, thus making it impossible for the participants to track the spatial position of the target. Then, the experimenter positioned the cards in a row on the tray, and asked the participants to locate the target by saying, “*And now, where is Charlie? Can you put him in his house?*” while pushing the tray toward the participant to encourage them to search. The trial ended once the participants had placed one card in the cardboard box. Four consecutive test trials were conducted, without any feedback from the experimenter to participants on their performance.

3.1.3 Coding and analysis

We coded the number of cards the participants flipped before placing one in the cardboard box for each of the test trials. In case a participant flipped the same card multiple times, we coded only a single flip, so that the maximum flipped cards were four. In case a participant grasped a card and directly placed it in the box without flipping it, we coded 0 for the number of characters flipped. In order to assess our participants' anticipation of the cognitive utility of sight, we coded the "grasping latency" of the participants by computing the time elapsed from the moment the experimenter finished saying the prompt (i.e., when she pronounced the last syllable of the sentence "*And now, where is Charlie?*"), till the participant grasped one of the cards (i.e., touched it simultaneously with the thumb and any of the other four fingers). We measured the grasping latency for the baseline trial (to have a baseline measure of the participants' motor speed) and for the four test trials. Next, we computed the grasping latency ratios (GLR) to estimate the relative speeding up (or slowing down) of participants during the test trials compared to the baseline. This ratio was computed for each test trial and each participant by dividing the grasping latency during the test trial by the grasping latency during the baseline trial. Thus, an average GLR below 1 (respectively above 1) indicates that participants are faster (respectively slower) to grasp a card during test trials than during the baseline trial. By performing our analyses on the GLR, a baseline-corrected measure, we reduced the influence of inter-individual differences in grasping speed on our results. In the non-informative condition, flipping the cards to see the symbol on their backside resulted in no information gain. In contrast, in the informative condition, the target could be identified by seeing the symbol on the cards' backside. Thus, we expected that if toddlers anticipate the cognitive utility of seeing symbols in Study 2, their GLR should be lower in the informative condition than in the non-informative condition. In the informative

condition, we also measured whether the participants found the target card. For each trial, the participants received a score of 1 when the first card they placed in the cardboard box was the target, and 0 otherwise. We planned to treat as missing data the data from trials in which the participants would not put any card inside the box. In fact, this never happened (0 trial out of 144).

3.2 Results and Discussion

First, we conducted a preliminary analysis to validate that the participants succeeded in finding the target character in the informative condition. Since there were four characters to choose from, chance predicted a mean success ratio of 0.25 in each trial. Across the four test trials in the informative condition, children's mean success ratio in finding the target character was significantly higher than that predicted by chance (0.25) ($M = 0.78$, $SD = 0.34$, $M_{dn} = 1.00$, $Z = 133.5$, $p < .001$, Wilcoxon signed-rank test). In a complementary analysis, we assessed the efficiency of toddlers' information search (see the Supplementary Materials for details). For most of the trials of the informative condition (50 out of 72), children searched for the target in an efficient manner: they flipped card successively, and stopped once they found the target card. Expectedly, in the informative condition, the participants' mean success ratio in finding the target was positively correlated with the number of trials in which they searched for the target in an efficient manner ($\rho = .89$, $p < .001$, Spearman's rank correlation).

Next, we assessed whether the participants' information search strategies differed across conditions. We first analyzed the participants' number of flips. A Friedman test revealed that trial number had no effect on the number of flips ($\chi^2(3) = 1.98$, $p = .577$, Kendall's $W = .02$). Subsequently, we computed the average number of flips per participant

across the four trials, and performed our analyses on this average score. The participants' average number of flipped cards was significantly higher in the informative condition ($M = 1.82$, $SD = 0.86$, $M_{dn} = 2.00$) than in the non-informative condition ($M = 0.91$, $SD = 0.45$, $M_{dn} = 1.00$, $U = 275$, $p < .001$, Mann-Whitney U test). Complementary analyses confirmed that there was no significant main effect of trial number, and no significant interaction between trial number and condition on children's number of flips (see the Supplementary Materials). Thus, the participants flipped cards more often when this action was informative for finding the target than when it was non-informative.

Second, we analyzed the participants' grasping latency ratio (GLR). A robust mixed-model full factorial ANOVA using the Welch-James ADF procedure on mean GLR with condition (informative vs. non-informative) as between-subject factors and Trial (1-4) as a within-subject factor revealed a main effect of condition ($F(1, 20.96) = 8.21$, $p = .008$), indicating that the participants' mean GLR was significantly lower in the informative condition than in the non-informative condition in this condition. The robust ANOVA also revealed a main effect of trial ($F(3, 17.57) = 5.37$, $p = .011$) and an interaction between condition and trial ($F(3, 17.57) = 3.34$, $p = .036$). Planned comparisons revealed that the participants' average GLR was significantly below 1 in the informative condition ($M = 0.69$, $SD = 0.46$, $M_{dn} = 0.69$, $Z = 31$, $p = .016$, Wilcoxon signed-rank test), indicating that the participants were faster to grasp a card during the test trials than during the baseline trials. In contrast, the participants' average GLR did not differ significantly from 1 in the non-informative condition ($M = 1.17$, $SD = 0.54$, $M_{dn} = 1.13$, $Z = 112$, $p = .265$, Wilcoxon signed-rank test).

Since the robust ANOVA on the average GLR revealed a main effect of trial and an interaction between trial and condition, we separately analyzed the participants' GLR data for

each trial. These analyses confirmed the effect of condition. As Figure 4 shows, the mean GLR was significantly lower in the informative condition than in the non-informative condition in the first three test trials. In the final test trial, the mean GLR did not differ significantly across conditions (informative vs. non-informative), possibly because the speed at which the participants grasped a card reached the ceiling by the end of the experiment.

Figure 4

In a few trials of Study 2, the participants grasped a card and directly placed it in the box without flipping it. As a result, the status of grasping latency as a measure of learning expectation during these trials was ambiguous. Thus, we reanalyzed our results for the GLR after excluding data from these trials (thus considering only the grasping gestures that led to flipping a card). These analyses confirmed our key results by showing an effect of condition on the participants' average GLR, with mean GLR significantly below that predicted by chance in the informative condition, and not significantly different from chance in the non-informative condition (see Supplemental Analyses in the Supplemental Materials).

The GLR data also rule out a potential alternative interpretation of our results. In Study 2, toddlers engaged in an action more often (flipping a card) when this action was informative. In itself, this pattern of behavior is consistent with two interpretations. The first interpretation, that we favor, is that toddlers evaluated the learning consequences of flipping a card, thus explaining why they performed this action more often when it yielded learning benefits. An alternative interpretation of this result would be that toddlers explored the cards without anticipating the learning consequences of their actions. They may for instance have performed a set of unspecific exploratory behaviors, including flipping cards, and stopped once they found a card that looked like the target. This alternative hypothesis can account for the effect of condition (informative vs. non-informative) on the number of cards that children

flipped. Indeed, in the non-informative condition, toddlers were guaranteed to discover a card looking like the target (without flipping any card). In contrast, in the informative condition, flipping at least one card—or more—was necessary to discover one that looked like the target (since in this condition, there was only one chance out of four to flip first the target card).

The participants' GLRs allow us to judge between these two hypotheses. The view that toddlers explore cards without anticipating the learning consequences of their action predicts that they should be equally fast to act on cards in all conditions. In contrast, we observed that toddlers' average GLR was consistently lower in the informative condition than in the non-informative condition. Thus, toddlers were comparatively faster to attempt at flipping the characters when this action yielded access to visual data allowing them to locate the target. This effect of condition on children's average GLR confirms that toddlers anticipated the possible learning consequences of their actions.

Our results suggest that the toddlers anticipated the informativeness of seeing the symbol located at the back of the cards. They did not use a purely confirmatory strategy (e.g., collecting evidence confirming their initial hypothesis). If they had done so, the participants' GLRs should have been the smallest in the non-informative condition, in which case the symbol located at the back of the card was guaranteed to confirm their hypothesis. Instead, our participants used an information gain strategy, such that their GLR to reach for a target was lower when they had a way to reduce their uncertainty about its location (by flipping it) than when they did not. In Study 2, we manipulated the informativeness of a given source of data (seeing). In Study 3, we sought to conceptually replicate and extend the generalizability of these findings by testing whether toddlers can select the most informative source of data

when choosing between two possible actions: flipping a character to see a symbol, or squeezing it to hear the sound it produced.

Study 3 also addresses a methodological issue. The results of Study 2 suggested that toddlers were sensitive to the informativeness of seeing symbols, with one caveat. In the informative condition there were more different kinds of cards to play with than in the non-informative condition. One anonymous reviewer suggested that perhaps, as a result of this difference, the game might be more enticing to children in the informative condition than in the non-informative condition, thus leading to faster reaching in the informative condition. This methodological issue is addressed in Study 3. In this case, the distribution of symbols was strictly identical across all conditions.

4 Study 3

4.1 Method

4.1.1 Participants

In Study 3, the participants were divided into two groups of eighteen toddlers (symbol condition: $M = 29.1$ months, $range = 23-34$ months, $SD = 2.36$; sound condition: $M = 30.6$ months, $range = 24-37$ months, $SD = 3.55$). Each participant was tested once, in a single condition. Seven additional participants were removed from the analysis because of fussiness (1) or refusing to play the game (6).

4.1.2 Materials and set-up

The materials and set-up were identical to those in Study 2, with the following exception. Instead of wooden cards, cushions ($10 \times 10 \times 5.5$ cm) with googly eyes on one side and a symbol sticker on the reverse side served as characters in the game (see Figure 5). All four characters were of identical shape and size, and were visually identical with the eye-

side facing up. One of the characters could be identified by pressing it. When squeezed, it emitted a loud squeaking noise, while the other three characters produced no sound. A second character could be identified by looking at the symbol placed on its reverse side. This character was the only one with a symbol different from the other characters. For half of the participants, the distinctive symbol placed on the reverse side of this second character was a red circle, while visually identical pairs of triangles were on the reverse side of each of the other characters; for the other half of the participants, we used the opposite pattern. The last two characters were completely identical; they made no sound when pressed, and had the same symbol on their reverse side. A rectangular cardboard box served as a “house” (hereafter referred to as the house-box; 13cm wide x 18cm deep x 7.5cm high). Furthermore, Study 3 also required the use of an opaque paper grocery bag.

Figure 5

Presentation of the game. The participants had to find a target character named “*Baptiste*” from among four cushion-shaped characters. At the beginning of the session, the experimenter placed the house-box on her right, at the far end of the board and said, “*Look, I brought little toys to play with.*” Then, she presented the first character to the child by lifting it, holding it eye-side towards the child while saying, “*You see, there is this one here.*” Next, she placed the character eye-side down on the board, thereby revealing the symbol placed on its backside, before pressing on the character to demonstrate whether it squeaked or made no sound. The experimenter repeated this procedure to present the second, third, and fourth characters, lining them up on the board one by one from her right to her left. When presenting the fourth and last character, the experimenter also added, “*This one is named Baptiste*” before placing it in line with the others and pressing on the character. In the sound condition, the target character, which was referred to as “Baptiste” was the only one that made a

squeaking sound when pressed. In the symbol condition, the character that was referred to as “Baptiste” was the only one to have a distinctive symbol placed on its reverse side. Apart from this difference, the procedures for the two conditions were identical. After presenting the cushions, the experimenter told the participant, “*We are going to play a game in which you have to find Baptiste, okay?*” and the warm-up trials started.

Warm-up trials. At the beginning of the first warm-up trial, the experimenter asked, “*Where is Baptiste now? Can you put him in his house?*” Next, she pushed the board and the house-box toward the child to indicate that it was the participant’s turn to place a toy into the box. When the participants did not choose the correct cushion, the experimenter corrected them by saying, “*That’s not Baptiste, this one here is Baptiste*” while pointing at the target character before pressing on it, and she repeated the prompt questions. Once the participants chose the target, they were congratulated, and the experimenter brought the materials back to her side of the table. Next, the experimenter positioned the characters on the board, and proceeded with the next warm-up trial. The second and third warm-up trials proceeded just like the first, with two exceptions. First, the position of the target character changed. For the second warm-up trial, the experimenter placed the target symbol side up at the right end of the row of characters. For the third warm-up trial, the experimenter placed the target second from the left end of the row, this time with its eye-side up, and its symbol side down. Second, in order to help the participants remember which character produced a sound when squeezed, the experimenter pressed sequentially on each cushion from left to right at the very beginning of the second and third warm-up trials. After the three warm-up trials, the participants proceeded to the test phase.

Test phase. At the beginning of each test trial, the experimenter placed the four characters into the grocery bag and said, “*Let’s mix them up!*” Next, the experimenter shook

the bag and removed the characters two at a time and placed them on the board in front of the participant, all with their eyes upward (making it impossible for the child to know which character was the target). Next, to help the participants understand that the characters were not in the same position as in the past trials, the experimenter quickly mixed them up, before placing them in a row on the board. All the characters now appeared identical from the viewpoint of the child. Discovering where the target was could only be accomplished by pressing on the characters (in the sound condition), or by flipping them to see the symbols (in the symbol condition). Next, the experimenter asked the prompt questions, “*Where is Baptiste now? Can you put him in his house?*” and she slid the board toward the child. The trial ended when the participant placed one of the characters in the house box, or 1 minute after the experimenter asked the prompt questions in case the participant did not interact at all with any of the cushions. There were four consecutive test trials, without any feedback from the experimenter to participants on their performance.

4.1.3 Coding and analysis

We coded the following four measures for each test trial: (i) whether the participants succeeded in finding the target character “Baptiste” (coded as 1 when the character the participants placed first in the cardboard box was the target, and 0 otherwise); (ii) the number of characters that the child squeezed; (iii) the number of characters that the child flipped (if a participant performed the same action, that is, flipping or squeezing a toy, on the same cushion multiple times, it was coded only once); and (iv) an estimate of the relative latency with which the participants engaged in flipping or squeezing a toy. We found it difficult to accurately code the exact latency for squeezing a toy. Thus, instead, we coded the first action that the participants performed on the toys (flipping or squeezing a character) for each trial in which children either flipped or squeezed a character. We used this measure of children’s

first action to test whether toddlers anticipated the learning consequences of their future behaviors at the planning stage.

For the trials in which the participant did not interact with any cushion, we coded 0 for the success in finding the target character, and 0 for the number of characters flipped and squeezed. When the participants did not place any toy inside the box before the end of the trial, we coded the data from the trial as missing data (3 trial out of 144).

4.2 Results and Discussion

Friedman tests revealed no main effect of trial on the participants' mean ratio of success in finding the target, the number of flips, and the number of squeezes (all p s > .659). Complementary analyses confirmed that there was no significant main effect of trial number, and no significant interaction between trial number and condition on any of our dependent variables (see the Supplementary Materials). Thus, for each of our measures, the participants' scores were averaged across the four trials. We performed subsequent statistical analyses on these average scores.

In a preliminary analysis, we assessed the participants' success in finding the target toy. Since there were four test trials, chance predicted an average success ratio of 0.25. Participants' mean ratio of success in finding the target tended to be higher than predicted by chance (0.25) in the symbol condition ($M = 0.72$, $SD = 0.36$, $M_{dn} = 1.00$, $Z = 146$, $p < .001$, Wilcoxon signed-rank test), but not in the sound condition ($M = 0.42$, $SD = 0.35$, $M_{dn} = 0.50$, $Z = 90$, $p = .079$, Wilcoxon signed-rank test). Thus, the participants tended to succeed in finding the target character in the experiment.

In a complementary analysis, we assessed the efficiency of participants' search strategies in the symbol and the sound conditions (see the Supplementary Materials for more

details). In the symbol condition, efficient searches — flipping the characters successively and stopping the search after flipping the target— were the modal response (44 trials out of 72). In the sound condition, efficient searches —squeezing the characters successively and stopping the search after squeezing the target— were observed (14 trials out of 72), but they were less frequent than in the sound condition. As in Study 2, there was a positive correlation between the participants' mean ratio of success in finding the target, and the number of trials in which they used an efficient search strategy ($\rho = .89$, $p < .001$, Spearman's rank correlation).

Next, we assessed whether the participants' strategy for finding the target differed across conditions. As shown in Figure 6, during the test trials, the participants flipped significantly more toys in the symbol condition ($M = 1.92$, $SD = 1.00$, $M_{dn} = 2.25$) than in the sound condition ($M = 0.42$, $SD = 0.73$, $M_{dn} = 0.25$, $U = 282$, $p < .001$, Mann-Whitney U test). Thus, the participants flipped toys more often when this action was more informative to discover the location of the target. Conversely, the effect of condition on the participants' tendency to squeeze toys were not statistically significant (symbol condition: $M = 0.47$, $SD = 0.66$, $M_{dn} = 0.00$; sound condition: $M = 0.82$, $SD = 0.84$, $M_{dn} = 0.50$, $U = 119.5$, $p = .164$, Mann-Whitney U test). Analyses focusing on the participants' first actions revealed that children flipped a toy first significantly more often in the symbol condition ($M = 0.93$, $SD = 0.17$, $M_{dn} = 1$) than in the sound condition ($M = 0.43$, $SD = 0.48$, $M_{dn} = 0.13$, $U = 210$, $p = .002$, Mann-Whitney U test, see Figure 6). Therefore, the participants were more likely to perform first an action relevant to discovering the location of the target.

Figure 6

Thus, the results of Study 3 confirm that toddlers adjust their search behaviors depending on the informativeness of visual data. Children were more likely to first flip a toy

when this action resulted in accessing visual information that was sufficiently informative to locate the target (in the symbol condition), than when it was not (in the sound condition). The evidence suggesting that the participants anticipated the amount of information gained from squeezing the toy was less clear (although the data trended in the expected direction). Our data do not allow us to pinpoint the exact sources of children's difficulty in the sound condition. We speculate that children found it difficult to establish that squeezing the toy was relevant to locating the target. This issue need not be originating from a general difficulty to process the informativeness of auditory information. Rather, it may be specific to our experimental set-up. For instance, children may have found it harder to track the informativeness of squeezing the toys because they could not be entirely certain that the target toy was the only one producing a sound when squeezed (since the causal mechanism producing the sound within the toys was not directly observable). Nevertheless, the results of Study 3 confirm those of Study 2 in suggesting that children can anticipate the informativeness of the data gained from visual access.

5 General Discussion

Mounting evidence suggests that infants monitor the learning benefits resulting from receiving a piece of visual information (Kidd et al., 2012; Poli et al., 2020; Stahl & Feigenson, 2015). Our results indicate that by toddlerhood onward, humans also rely on a model of their own learning from sight to anticipate the epistemic consequences of their actions. This early developing model of learning takes into account both the availability and cognitive utility of sight.

In Study 1, 14-month-old infants were faster to perform an action (opening a shutter) when it allowed them to see an object inside a box, than when it did not. This result adds to

previous evidence suggesting that infants monitor what others can and cannot see (e.g., Choi et al., 2018; Liskowski et al., 2007; Luo & Baillargeon, 2007; Moll & Tomasello, 2006; O'Neill, 1996; Sodian et al., 2007; Southgate et al., 2007). Our results suggest that infants can use their capacity to represent visual access to guide their own search for information. This capacity implies that infants form an epistemic goal (seeing the object inside the box), and select the most appropriate action to achieve it, taking into account the constraints of the situation (in our study, the opacity or transparency of the window). Thus, infants rely on information about the availability of visual data to plan their information seeking behaviors.

In Studies 2 and 3, we found that 2.5-year-old toddlers adjust their information search to the cognitive utility of seeing. In these experiments, the participants flipped characters in order to see a symbol more often when this action was informative than when it was not. In Study 2, toddlers' Grasping Latency Ratio before flipping a character was significantly lower when this action was more informative. Similarly, in Study 3, toddlers were more likely to flip a character before performing an alternative action (squeezing a character) when flipping characters was informative. These results suggest that toddlers anticipate the cognitive gains resulting from sight (in our experiments, from seeing a symbol) when planning their search for data. Children's capacity to select and assess learning actions and teaching based on their informativeness has been evidenced in studies of preschoolers' exploratory play (Cook et al., 2011; Ruggeri et al., 2019; Schulz & Bonawitz, 2007; Van Schijndel et al., 2015), in their assessment of informants (Gweon et al., 2014; Gweon & Asaba, 2018), and in their formulation of questions (Legare et al., 2013; Ruggeri et al., 2017). We demonstrated comparable capacities in the perceptual domain in much younger participants. Instead of expecting that seeing necessarily leads to knowing, toddlers adjusted their behaviors to the capacity of a visual input to reduce their own uncertainty. Thus, well before they can

explicitly talk about sources of knowledge, toddlers do not conceive of seeing as a purely behavioral event (such as merely building an unobstructed line of sight to an object). Instead, they are sensitive to the amount of relevant information carried by a given visual input, and they adjust their information search accordingly. To clarify: we are not claiming that visual data are fundamentally different from data coming from other modalities with respect to their informativeness. However, whether infants and toddlers' capacity to anticipate the learning consequences of future perceptual inputs generalizes to other modalities than vision is a matter of future empirical research.

Our data also provide information about the early development of the capacity to represent hypotheses. Several authors argue that humans' flexible learning is likely to be supported by the capacity to represent hypotheses from infancy onwards (e.g., Cesana-Arlotti et al., 2018, 2020; Goddu et al., 2021; Gweon, 2021; Schulz, 2012; Stahl & Feigenson, 2015). In contrast, other authors claim that before late preschool age, children might lack the cognitive resources to mark representations as merely possible, thus making it impossible for them to represent and test hypotheses (Leahy & Carey, 2020). According to this second view, when evidence is compatible with several mutually exclusive hypotheses, children under four years of age simply pick one of them, and behave as if it were true. In some circumstances, such a process could result in behaviors that, for an external observer, may appear to be information search. However, such behaviors would not involve any genuine representation of possibilities (for details about how this type of argument might account for impressive demonstrations of infants' flexible learning, such as the data from Stahl and Feigenson (2015), see the Supplementary Materials).

Our data suggest that toddlers' information search cannot be reduced to a simple strategy such as generating a guess and acting on it. Had toddlers used such a strategy, they should have behaved exactly in the same manner in all the conditions of Studies 2 and 3 (i.e., they should have merely picked one toy — based on their guess — and assumed that it was the target without engaging in any information search at all). Instead, children adequately generated the evidence that was relevant to discriminate between the precise hypotheses they needed to assess. For instance, in Studies 2 and 3, children were comparatively faster to flip first toys to reveal symbols on their backs when this information was needed to locate a target toy, than when this information was insufficient to locate the target. In order to be sensitive to the future learning gains of flipping a toy in our experiment, toddlers needed to anticipate that their flipping action could result in one of two different possible outcomes (e.g., either finding a circle or finding a pair of triangles at the back of the toy). In other words, they needed to represent two distinct possibilities. Thus, our data dovetail with recent studies suggesting that children might be able to represent hypotheses from an early age (e.g., Cesana-Arlotti et al., 2020; Goddu et al., 2021).

Our data also contribute to debates on young children's sensitivity to their uncertainty. Young children and toddlers often behave differently when they are ignorant or uncertain than when they are knowledgeable about a piece of information (e.g., Call & Carpenter, 2001; Coughlin et al., 2015; Goupil et al., 2016; Kim et al., 2020; Ruggeri et al., 2019; for a review, see Goupil & Kouider, 2019). These results suggest that young children are sensitive to their own uncertainty. Yet, there are debates about the nature of humans' precocious sensitivity to uncertainty. For instance, it has been argued that full-blown representations of the content

of what one is uncertain or ignorant about would not emerge before 4 years of age (Kloo et al., 2017; Leahy & Carey, 2020; Perner, 2012; Rohwer et al., 2012). According to these views, prior to four years of age, children might monitor correlates of their uncertainty (e.g., an elevated heart-beat, hesitating when making a choice, the vividness of a representation, and so on), without representing the specific hypotheses that they are uncertain about.

Our data suggest that from toddlerhood onwards, humans do more than monitor general correlates of their uncertainty. In all of our studies, if children's information seeking behaviors had been merely triggered by general correlates of uncertainty, children would have behaved identically in all of our test conditions (since the participants were equally ignorant about the exact same piece of information in all cases). Instead, children planned appropriate actions to resolve their ignorance by anticipating their own learning. Such anticipations and context-sensitive planning are consistent with the view that toddlers can represent what they are uncertain about (e.g., the shape of a symbol located at the back of a toy).

In its most elaborate form, anticipating the informativeness of a learning action may seem like a very complex task. It may imply, for example, tracking all the possible perceptual events resulting from a learning action, and anticipating the overall reduction of uncertainty resulting from all possible outcomes (as the ideal learner depicted in Figure 3, panel B). While we do not rule out that toddlers may use such a complex mechanism, they may also rely on simpler procedures that remain effective whilst having a lower cognitive cost. For example, they may simply consider a single or a few hypotheses for the future perceptual event resulting from a learning action (e.g., discovering one specific symbol on a character's

backside), and assess the resulting reduction in their uncertainty. These computations can be performed by simply monitoring one's uncertainty (Coughlin et al., 2015; Geurten & Bastin, 2019; Goupil et al., 2016; Kim et al., 2020), and representing future perceptual events (Siegel et al., in press).

In short, even if many species can search for specific pieces of information efficiently, humans stand out in their ability to plan learning actions in creative and flexible ways. The capacity to discover and adjust information-seeking strategies in a contextually sensitive manner requires assessing the learning outcomes of future actions or events. We offer evidence of the early ontogenetic roots of this capacity. Our data indicate that before their third birthday, toddlers plan their epistemic actions by anticipating whether they will result in learning, and adjust them to (i) epistemic constraints and (ii) the cognitive utility of a piece of evidence. Our data dovetail with a growing body of evidence suggesting that young children engage in sophisticated forms of active learning (e.g., for reviews, see Begus & Bonawitz, 2020; Schulz, 2012; Twomey & Westermann, 2018). Importantly, our Studies—just like many other studies on active learning—have tested children from predominately wealthy, urban, educated and industrialized backgrounds. Whether the early development of active learning generalizes to children from other socio-cultural backgrounds is an important question for future research.

Data of all Studies and analysis scripts are accessible on an open repository (URL : https://osf.io/9jpmv/?view_only=7e6b6a4b36714fcea6a537f8bc21b557).

6 References

- Bass, I., Gopnik, A., Hanson, M., Ramarajan, D., Shafto, P., Wellman, H., & Bonawitz, E. (2019). Children's developing theory of mind and pedagogical evidence selection. *Developmental Psychology*, 55(2), 286-302. <https://doi.org/10.1037/dev0000642>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1). <https://doi.org/10.18637/jss.v067.i01>
- Begus, K., & Bonawitz, E. (2020). The rhythm of learning: Theta oscillations as an index of active learning in infancy. *Developmental Cognitive Neuroscience*, 45, 100810. <https://doi.org/10.1016/j.dcn.2020.100810>
- Beran, M. J., Smith, J. D., & Perdue, B. M. (2013). Language-Trained Chimpanzees (Pan troglodytes) Name What They Have Seen but Look First at What They Have Not Seen. *Psychological Science*, 24(5), 660-666. <https://doi.org/10.1177/0956797612458936>
- Berlyne, D. E. (1966). Curiosity and Exploration. *Science*, 153(3731), 25-33.
- Blanchard, T. C., Hayden, B. Y., & Bromberg-Martin, E. S. (2015). Orbitofrontal Cortex Uses Distinct Codes for Different Choice Attributes in Decisions Motivated by Curiosity. *Neuron*, 85(3), 602-614. <https://doi.org/10.1016/j.neuron.2014.12.050>
- Bloch, M. (2008). Truth and sight: Generalizing without universalizing. *Journal of the Royal Anthropological Institute*, 14(s1), S22-S32. <https://doi.org/10.1111/j.1467-9655.2008.00490.x>
- Bridgers, S., Jara-Ettinger, J., & Gweon, H. (2020). Young children consider the expected utility of others' learning to decide what to teach. *Nature Human Behaviour*, 4(2), 144-152. <https://doi.org/10.1038/s41562-019-0748-6>
- Butler, R. A. (1954). Curiosity in Monkeys. *Scientific American*, 190(2), 70-75.

- Cacchione, T., Schaub, S., & Rakoczy, H. (2013). Fourteen-month-old infants infer the continuous identity of objects on the basis of nonvisible causal properties. *Developmental Psychology*, 49(7), 1325-1329. <https://doi.org/10.1037/a0029746>
- Calhoun, A. J., Chalasani, S. H., & Sharpee, T. O. (2014). Maximally informative foraging by *Caenorhabditis elegans*. *eLife*, 3, e04220. <https://doi.org/10.7554/eLife.04220>
- Call, J. (2001). Object permanence in orangutans (*Pongo pygmaeus*), chimpanzees (*Pan troglodytes*), and children (*Homo sapiens*). *Journal of Comparative Psychology*, 115(2), 159-171. <https://doi.org/10.1037/0735-7036.115.2.159>
- Call, J. (2010). Do apes know that they could be wrong? *Animal Cognition*, 13(5), 689-700. <https://doi.org/10.1007/s10071-010-0317-x>
- Call, J., & Carpenter, M. (2001). Do apes and children know what they have seen? *Animal Cognition*, 3(4), 207-220. <https://doi.org/10.1007/s100710100078>
- Carruthers, P. (2008). Meta-cognition in animals: A skeptical look. *Mind & Language*, 23(1), 58-89. <https://doi.org/10.1111/j.1468-0017.2007.00329.x>
- Carruthers, P. (2018). Basic questions. *Mind & Language*, 33(2), 130-147. <https://doi.org/10.1111/mila.12167>
- Cesana-Arlotti, N., Kovács, Á. M., & Téglás, E. (2020). Infants recruit logic to learn about the social world. *Nature Communications*, 11(1), 5999. <https://doi.org/10.1038/s41467-020-19734-5>
- Cesana-Arlotti, N., Martín, A., Téglás, E., Vorobyova, L., Cetnarski, R., & Bonatti, L. L. (2018). Precursors of logical reasoning in preverbal human infants. *Science*, 359(6381), 1263-1266. <https://doi.org/10.1126/science.aao3539>

- Choi, Y., Mou, Y., & Luo, Y. (2018). How do 3-month-old infants attribute preferences to a human agent? *Journal of Experimental Child Psychology*, *172*, 96-106. <https://doi.org/10.1016/j.jecp.2018.03.004>
- Collier-Baker, E., & Suddendorf, T. (2006). Do Chimpanzees (Pan troglodytes) and 2-Year-Old Children (Homo sapiens) Understand Double Invisible Displacement? *Journal of Comparative Psychology*, *120*(2), 89-97. <https://doi.org/10.1037/0735-7036.120.2.89>
- Cook, C., Goodman, N. D., & Schulz, L. E. (2011). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, *120*(3), 341-349. <https://doi.org/10.1016/j.cognition.2011.03.003>
- Coughlin, C., Hembacher, E., Lyons, K. E., & Ghetti, S. (2015). Introspection on uncertainty and judicious help-seeking during the preschool years. *Developmental Science*, *18*(6), 957-971. <https://doi.org/10.1111/desc.12271>
- Crystal, J. D., & Foote, A. L. (2011). Evaluating information-seeking approaches to metacognition. *Current Zoology*, *57*(4), 531-542. <https://doi.org/10.1093/czoolo/57.4.531>
- Degen, J., Kirbach, A., Reiter, L., Lehmann, K., Norton, P., Storms, M., Koblöfsky, M., Winter, S., Georgieva, P. B., Nguyen, H., Chamkhi, H., Greggers, U., & Menzel, R. (2015). Exploratory behaviour of honeybees during orientation flights. *Animal Behaviour*, *102*, 45-57. <https://doi.org/10.1016/j.anbehav.2014.12.030>
- D'Entremont, B., & Morgan, R. (2006). Experience with visual barriers and its effects on subsequent gaze-following in 12- to 13-month-olds. *British Journal of Developmental Psychology*, *24*(3), 465-475. <https://doi.org/10.1348/026151005X51248>
- Erceg-Hurn, D. M., & Mirosevich, V. M. (2008). Modern robust statistical methods : An easy way to maximize the accuracy and power of your research. *American Psychologist*, *63*(7), 591-601. <https://doi.org/10.1037/0003-066X.63.7.591>

- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3 : A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175-191. <https://doi.org/10.3758/BF03193146>
- Feigenson, L., & Carey, S. (2003). Tracking individuals via object-files : Evidence from infants' manual search. *Developmental Science*, 6(5), 568-584. <https://doi.org/10.1111/1467-7687.00313>
- Fitneva, S. A., Lam, N. H. L., & Dunfield, K. A. (2013). The development of children's information gathering : To look or to ask? *Developmental Psychology*, 49(3), 533-542. <https://doi.org/10.1037/a0031326>
- Flavell, J. H., Shipstead, S. G., & Croft, K. (1978). Young Children's Knowledge about Visual Perception: Hiding Objects from Others. *Child Development*, 49(4), 1208-1211. JSTOR. <https://doi.org/10.2307/1128761>
- Frank, M. C., & Goodman, N. D. (2012). Predicting Pragmatic Reasoning in Language Games. *Science*, 336(6084), 998-998. <https://doi.org/10.1126/science.1218633>
- Geurten, M., & Bastin, C. (2019). Behaviors speak louder than explicit reports : Implicit metacognition in 2.5-year-old children. *Developmental Science*, 22(2). <https://doi.org/10.1111/desc.12742>
- Glickman, S. E., & Sroges, R. W. (1966). Curiosity in Zoo Animals. *Behaviour*, 26(1-2), 151-187. <https://doi.org/10.1163/156853966X00074>
- Goddu, M. K., Sullivan, J. N., & Walker, C. M. (2021). Toddlers learn and flexibly apply multiple possibilities. *Child Development*, cdev.13668. <https://doi.org/10.1111/cdev.13668>
- Goupil, L., & Kouider, S. (2019). Developing a Reflective Mind : From Core Metacognition to Explicit Self-Reflection. *Current Directions in Psychological Science*, 28(4), 403-408. <https://doi.org/10.1177/0963721419848672>

- Goupil, L., Romand-Monnier, M., & Kouider, S. (2016). Infants ask for help when they know they don't know. *Proceedings of the National Academy of Sciences*, *113*(13), 3492-3496. <https://doi.org/10.1073/pnas.1515129113>
- Gweon, H. (2021). Inferential social learning : Cognitive foundations of human social learning and teaching. *Trends in Cognitive Sciences*, *25*(10), 896-910. <https://doi.org/10.1016/j.tics.2021.07.008>
- Gweon, H., & Asaba, M. (2018). Order Matters: Children's Evaluation of Underinformative Teachers Depends on Context. *Child Development*, *89*(3), e278-e292. <https://doi.org/10.1111/cdev.12825>
- Gweon, H., Pelton, H., Konopka, J. A., & Schulz, L. E. (2014). Sins of omission : Children selectively explore when teachers are under-informative. *Cognition*, *132*(3), 335-341. <https://doi.org/10.1016/j.cognition.2014.04.013>
- Hampton, R. R. (2009). Multiple demonstrations of metacognition in nonhumans: Converging evidence or multiple mechanisms? *Comparative cognition & behavior reviews*, *4*, 17-28.
- Iwasaki, S., Kuroshima, H., Arahori, M., & Fujita, K. (2020). Prospective information-seeking in human children (Homo sapiens): When to seek and what to seek. *Journal of Comparative Psychology*, *134*(3), 341-348. <https://doi.org/10.1037/com0000217>
- Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T., & Camerer, C. F. (2009). The Wick in the Candle of Learning : Epistemic Curiosity Activates Reward Circuitry and Enhances Memory. *Psychological Science*, *20*(8), 963-973.
- Keselman, H. J., Wilcox, R. R., & Lix, L. M. (2003). A generally robust approach to hypothesis testing in independent and correlated groups designs. *Psychophysiology*, *40*(4), 586-596. <https://doi.org/10.1111/1469-8986.00060>

- Kidd, C., & Hayden, B. Y. (2015). The Psychology and Neuroscience of Curiosity. *Neuron*, 88(3), 449-460. <https://doi.org/10.1016/j.neuron.2015.09.010>
- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The Goldilocks Effect: Human Infants Allocate Attention to Visual Sequences That Are Neither Too Simple Nor Too Complex. *PLOS ONE*, 7(5), e36399. <https://doi.org/10.1371/journal.pone.0036399>
- Kim, S., Sodian, B., & Proust, J. (2020). 12- and 24-Month-Old Infants' Search Behavior Under Informational Uncertainty. *Frontiers in Psychology*, 11. <https://doi.org/10.3389/fpsyg.2020.00566>
- Király, I., Oláh, K., Csibra, G., & Kovács, Á. M. (2018). Retrospective attribution of false beliefs in 3-year-old children. *Proceedings of the National Academy of Sciences*, 115(45), 11477-11482. <https://doi.org/10.1073/pnas.1803505115>
- Kloo, D., Rohwer, M., & Perner, J. (2017). Direct and indirect admission of ignorance by children. *Journal of Experimental Child Psychology*, 159, 279-295. <https://doi.org/10.1016/j.jecp.2017.02.014>
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software*, 82(13). <https://doi.org/10.18637/jss.v082.i13>
- Leahy, B. P., & Carey, S. E. (2020). The Acquisition of Modal Concepts. *Trends in Cognitive Sciences*, 24(1), 65-78. <https://doi.org/10.1016/j.tics.2019.11.004>
- Legare, C. H., Mills, C. M., Souza, A. L., Plummer, L. E., & Yasskin, R. (2013). The use of questions as problem-solving strategies during early childhood. *Journal of Experimental Child Psychology*, 114(1), 63-76. <https://doi.org/10.1016/j.jecp.2012.07.002>

- Liszkowski, U., Carpenter, M., & Tomasello, M. (2007). Pointing out new news, old news, and absent referents at 12 months of age. *Developmental Science*, *10*(2), F1-F7. <https://doi.org/10.1111/j.1467-7687.2006.00552.x>
- 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. <https://doi.org/10.1016/j.cognition.2006.10.007>
- Marsh, H. L. (2014). Metacognitive-like information seeking in lion-tailed macaques: A generalized search response after all? *Animal Cognition*, *17*(6), 1313-1328. <https://doi.org/10.1007/s10071-014-0767-7>
- Marsh, H. L., & MacDonald, S. E. (2012). Information seeking by orangutans: A generalized search strategy? *Animal Cognition*, *15*(3), 293-304. <https://doi.org/10.1007/s10071-011-0453-y>
- Mascaro, O., Aguirre, M., Brun, M., Couderc, A., & Mercier, H. (2019). Nonverbal rhetoric: 2- to 4-year-old children select relevant evidence when trying to influence others. *Developmental Psychology*, *55*(10), 2039-2047. <https://doi.org/10.1037/dev0000779>
- Mascaro, O., Morin, O., & Sperber, D. (2017). Optimistic expectations about communication explain children's difficulties in hiding, lying, and mistrusting liars. *Journal of Child Language*, *44*(5), 1041-1064. <https://doi.org/10.1017/S0305000916000350>
- Meltzoff, A. N., & Brooks, R. (2008). Self-experience as a mechanism for learning about others: A training study in social cognition. *Developmental Psychology*, *44*(5), 1257-1265. <https://doi.org/10.1037/a0012888>
- Moll, H., & Tomasello, M. (2004). 12- and 18-month-old infants follow gaze to spaces behind barriers. *Developmental Science*, *7*(1), F1-F9. <https://doi.org/10.1111/j.1467-7687.2004.00315.x>

- Moll, H., & Tomasello, M. (2006). Level 1 perspective-taking at 24 months of age. *British Journal of Developmental Psychology*, 24(3), 603-613. <https://doi.org/10.1348/026151005X55370>
- Neldner, K., Collier-Baker, E., & Nielsen, M. (2015). Chimpanzees (*Pan troglodytes*) and human children (*Homo sapiens*) know when they are ignorant about the location of food. *Animal Cognition*, 18(3), 683-699. <https://doi.org/10.1007/s10071-015-0836-6>
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101(4), 608-631. <https://doi.org/10.1037/0033-295X.101.4.608>
- O'Neill, D. K. (1996). Two-Year-Old Children's Sensitivity to a Parent's Knowledge State When Making Requests. *Child Development*, 67(2), 659-677. <https://doi.org/10.1111/j.1467-8624.1996.tb01758.x>
- O'Neill, D. K., Astington, J. W., & Flavell, J. H. (1992). Young Children's Understanding of the Role That Sensory Experiences Play in Knowledge Acquisition. *Child Development*, 63(2), 474-490. <https://doi.org/10.1111/j.1467-8624.1992.tb01641.x>
- O'Neill, D. K., & Chong, S. (2001). Preschool Children's Difficulty Understanding the Types of Information Obtained through the Five Senses. *Child Development*, 72(3), 803-815. <https://doi.org/10.1111/1467-8624.00316>
- Pelz, M., & Kidd, C. (2020). The elaboration of exploratory play. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 375(1803), 20190503. <https://doi.org/10.1098/rstb.2019.0503>
- Perner, J. (2012). MiniMeta: In search of minimal criteria for metacognition. In M. J. Beran, J. Brandl, J. Perner, & J. Proust (Éds.), *Foundations of Metacognition* (p. 94-116). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199646739.003.0007>

- Perner, J., & Ruffman, T. (1995). Episodic Memory and Autonoetic Consciousness: Developmental Evidence and a Theory of Childhood Amnesia. *Journal of Experimental Child Psychology*, 59(3), 516-548. <https://doi.org/10.1006/jecp.1995.1024>
- Pillow, B. H. (1993). Preschool children's understanding of the relationship between modality of perceptual access and knowledge of perceptual properties. *British Journal of Developmental Psychology*, 11. <https://doi.org/10.1111/j.2044-835X.1993.tb00610.x>
- Pisula, W. (2009). *Curiosity and information seeking in animal and human behavior*. Brown Walker Press.
- Poli, F., Serino, G., Mars, R. B., & Hunnius, S. (2020). *Infants tailor their attention to maximize learning* [Data set]. Radboud Data Repository. <https://repository.ubn.ru.nl/handle/2066/221851>
- R Core Team. (2020). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. <https://www.R-project.org/>
- Robinson, E. J., Haigh, S. N., & Pendle, J. E. C. (2008). Children's working understanding of the knowledge gained from seeing and feeling. *Developmental Science*, 11(2), 299-305. <https://doi.org/10.1111/j.1467-7687.2008.00676.x>
- Robinson, E. J., Thomas, G. V., Parton, A., & Nye, R. (1997). Children's overestimation of the knowledge to be gained from seeing. *British Journal of Developmental Psychology*, 15(3), 257-273. <https://doi.org/10.1111/j.2044-835X.1997.tb00520.x>
- Rohwer, M., Kloo, D., & Perner, J. (2012). Escape From Metaignorance: How Children Develop an Understanding of Their Own Lack of Knowledge: Overestimation of Own Knowledge in Young Children. *Child Development*, 83(6), 1869-1883. <https://doi.org/10.1111/j.1467-8624.2012.01830.x>

- Rosati, A. G., & Santos, L. R. (2016). Spontaneous Metacognition in Rhesus Monkeys. *Psychological Science, 27*(9), 1181-1191. <https://doi.org/10.1177/0956797616653737>
- Ruggeri, A., Sim, Z. L., & Xu, F. (2017). “Why is Toma late to school again?” Preschoolers identify the most informative questions. *Developmental Psychology, 53*(9), 1620-1632. <https://doi.org/10.1037/dev0000340>
- Ruggeri, A., Swaboda, N., Sim, Z. L., & Gopnik, A. (2019). Shake it baby, but only when needed : Preschoolers adapt their exploratory strategies to the information structure of the task. *Cognition, 193*, 104013. <https://doi.org/10.1016/j.cognition.2019.104013>
- Schulz, L. (2012). The origins of inquiry : Inductive inference and exploration in early childhood. *Trends in Cognitive Sciences, 16*(7), 382-389. <https://doi.org/10.1016/j.tics.2012.06.004>
- Schulz, L., & Bonawitz, E. B. (2007). Serious fun : Preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology, 43*(4), 1045-1050. <https://doi.org/10.1037/0012-1649.43.4.1045>
- Senju, A., Southgate, V., Snape, C., Leonard, M., & Csibra, G. (2011). Do 18-Month-Olds Really Attribute Mental States to Others? : A Critical Test. *Psychological Science, 22*(7), 878-880. <https://doi.org/10.1177/0956797611411584>
- Shafto, P., Goodman, N. D., & Frank, M. C. (2012). Learning From Others : The Consequences of Psychological Reasoning for Human Learning. *Perspectives on Psychological Science, 7*(4), 341-351. <https://doi.org/10.1177/1745691612448481>
- Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology, 71*, 55-89. <https://doi.org/10.1016/j.cogpsych.2013.12.004>
- Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal, 27*(3), 379-423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>

- Siegel, M. H., Magid, R. W., Pelz, M., Tenenbaum, J. B., & Schulz, L. E. (in press). Intuitive psychophysics: Children's exploratory play tracks the discriminability of hypotheses. *Nature Communications*, in press.
- Sodian, B., Thoermer, C., & Metz, U. (2007). Now I see it but you don't: 14-month-olds can represent another person's visual perspective. *Developmental Science*, 10(2), 199-204. <https://doi.org/10.1111/j.1467-7687.2007.00580.x>
- Southgate, V., Senju, A., & Csibra, G. (2007). Action Anticipation Through Attribution of False Belief by 2-Year-Olds. *Psychological Science*, 18(7), 587-592. <https://doi.org/10.1111/j.1467-9280.2007.01944.x>
- Sperber, D., & Wilson, D. (2001). *Relevance: Communication and cognition* (2nd ed). Blackwell Publishers.
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, 348(6230), 91-94. <https://doi.org/10.1126/science.aaa3799>
- Starkey, P. (1992). The early development of numerical reasoning. *Cognition*, 43(2), 93-126. [https://doi.org/10.1016/0010-0277\(92\)90034-F](https://doi.org/10.1016/0010-0277(92)90034-F)
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27(3), 453-489. https://doi.org/10.1207/s15516709cog2703_6
- Synnott, A. (1992). The eye and I: A sociology of sight. *International Journal of Politics, Culture, and Society*, 5(4), 617-636. <https://doi.org/10.1007/BF01419559>
- Twomey, K. E., & Westermann, G. (2018). Curiosity-based learning in infants: A neurocomputational approach. *Developmental Science*, 21(4), e12629. <https://doi.org/10.1111/desc.12629>

- Van Schijndel, T. J. P., Visser, I., van Bers, B. M. C. W., & Raijmakers, M. E. J. (2015). Preschoolers perform more informative experiments after observing theory-violating evidence. *Journal of Experimental Child Psychology*, *131*, 104-119. <https://doi.org/10.1016/j.jecp.2014.11.008>
- Vergassola, M., Villermaux, E., & Shraiman, B. I. (2007). 'Infotaxis' as a strategy for searching without gradients. *Nature*, *445*(7126), 406-409. <https://doi.org/10.1038/nature05464>
- Villacorta, P., J. (2017). The welchADF Package for Robust Hypothesis Testing in Unbalanced Multivariate Mixed Models with Heteroscedastic and Non-normal Data. *The R Journal*, *9*(2), 309. <https://doi.org/10.32614/RJ-2017-049>
- Wang, J. J., Miletich, D. D., Ramsey, R., & Samson, D. (2014). Adults see vision to be more informative than it is. *Quarterly Journal of Experimental Psychology*, *67*(12), 2279-2292. <https://doi.org/10.1080/17470218.2014.915331>
- Wilson, T. (2000). Human Information Behavior. *Informing Science: The International Journal of an Emerging Transdiscipline*, *3*, 049-056. <https://doi.org/10.28945/576>


7 Figures and Legends

Figure 1. Box Used in Study 1, Per Condition (Informative vs. Non-informative).




Figure 2. Mean Delay to Touch the Shutter's Handle (SEM) Per Group (Experimental vs. Control), and Per Condition (Transparent vs. Opaque). Stars represent p-values for comparisons between conditions by Mann-Whitney U tests.

ns : not significant, ** : $p < .01$.

TODDLERS PLAN EFFICIENT EPISTEMIC ACTIONS

Condition	Study 2		Study 3	
	Informative	Uninformative	Symbol	Sound
$E(I_{\theta})$ - Expected information gain (reduction of uncertainty) of a first flip	-.82	0	-.82	-.11
$E(I_{\theta})$ - Expected information gain (reduction of uncertainty) of a first squeeze	-	-	-.11	-.82

C

TODDLERS PLAN EFFICIENT EPISTEMIC ACTIONS

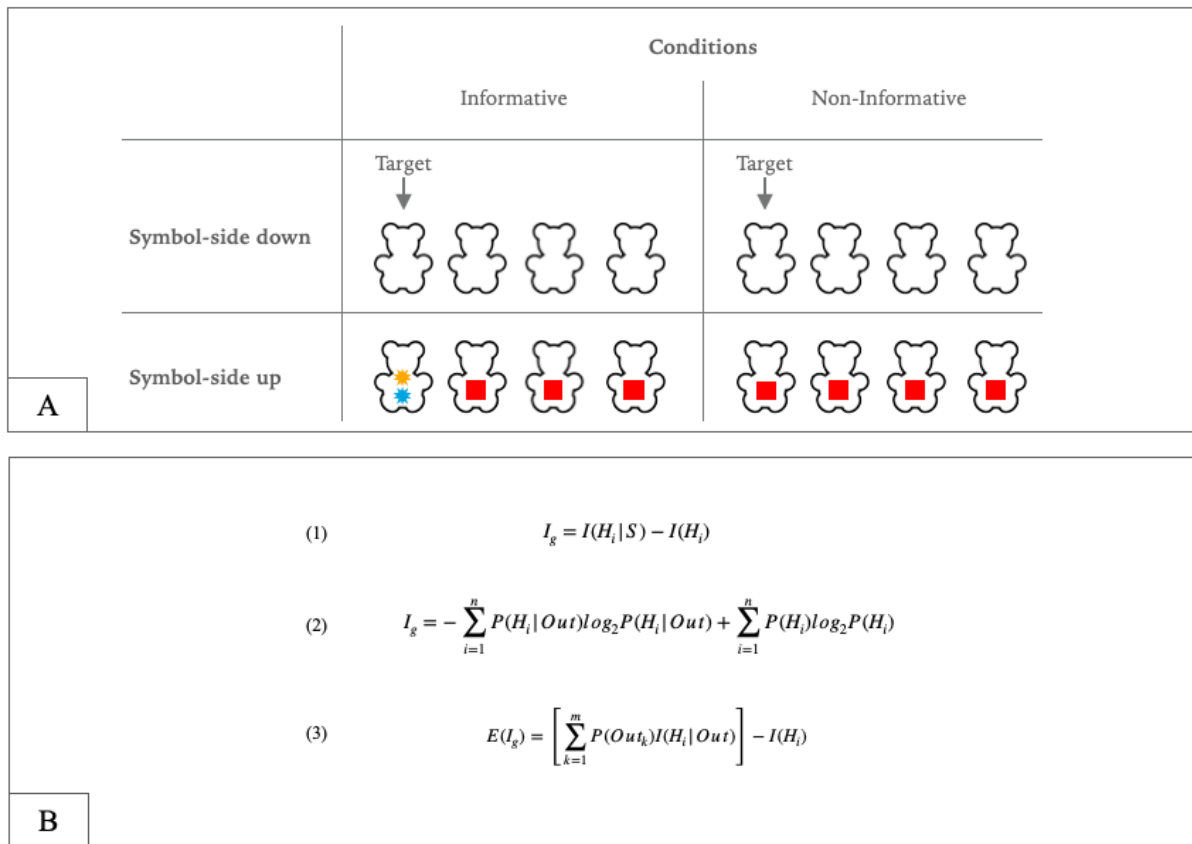


Figure 3. Panel A: Differences Between the Two Conditions of Study 2. The bear shaped cards were all visually identical when their symbol-side was down. When their symbol-side was up, the target card differed visually from the others in the informative condition only. The arrows indicate the target card that the participants had to find. Panel B: A Formal Model of Expected Information Gains (Adapted from Oaksford & Chater, 1994). In this model, the learners aim to reduce their uncertainty I over a set of hypotheses H_i (i.e., $I(H_i)$). After observing the outcome Out of a given action (e.g., seeing a specific symbol from flipping a card), the learners revise their uncertainty to $I(H_i|Out)$. The information gain (I_g) resulting from observing a specific outcome Out is the reduction of uncertainty (1), where uncertainty is quantified by the standard entropy measure from information theory (2). The learner does not know what the outcome of given action will be before engaging in it (e.g., she does not

know which symbol will be at the back of a card before flipping it). Thus, the expected information gain $E(I_g)$ is computed taking into account all the k possible outcomes that may be observed as a result of a given action (e.g., in Study 2's informative condition, two outcomes may result from flipping a card : observing either the unique symbol found on the target, or the symbol that is found on all the distractor cards). $E(I_g)$ is the uncertainty after performing a given action, weighted by the probability of each specific outcome, minus the prior uncertainty (3). Panel C: Expected information gains of actions computed by the model depending on condition. In the non-informative condition of the Study 2, the expected information gain of flipping a card is null (given that flipping a card has no impact on one's uncertainty about target's location). In the informative condition of Study 2, the expected information gain of flipping a first card is higher than in the non-informative condition, although it is not maximal. Indeed, in this case there is a one in four chance of flipping the target card — an outcome that would reduce children's uncertainty maximally; there is a three in four chance of flipping another card — an outcome that would only reduce children's uncertainty a little (since it would merely exclude one of the possible locations of the target). In Study 3, the expected information gains of the first action differ across conditions. In the symbol condition flipping a character is more informative than squeezing one. The opposite is true in the sound condition. The expected information gain of first flipping a character is not null in the sound condition because it may result in discovering the character with the unique symbol (which does not squeak), and thus, in excluding one possibility about the potential location of the target. Similarly, the expected information gain of first squeezing a character is not null in the symbol condition because it may result in discovering the character that squeaks (which, in this condition, is not the target).

TODDLERS PLAN EFFICIENT EPISTEMIC ACTIONS




Figure 4. Mean Grasping Latency Ratios for Each Trial (SEM) Per Condition (Informative vs. Non-informative). Stars represent p-values for comparisons between conditions by Mann-Whitney U tests.

ns : not significant, * : $p < .05$, ** : $p < .01$.

TODDLERS PLAN EFFICIENT EPISTEMIC ACTIONS









	Conditions	
	Symbol	Sound
Symbol-side down	Target 	Target 
Symbol-side up		
Sound when squeezed	X X X 	X X X 

Figure 5. Differences Between the Two Conditions of Study 3. The characters were all visually identical when placed symbol side-down, in both conditions. In the symbol condition the target character was the only one with a symbol different from the others. In the sound condition, the target character was the only one producing a loud noise when squeezed. The arrows indicate the target characters.



B

Figure 6. Panel A: Mean Proportion of the First action Type Over Trials (SEM; either flip or squeeze) per Condition (sound and symbol). Panel B: Mean Number of Flips and Squeezes per Trials (SEM) and per Condition. Stars represent p-values for comparisons between conditions by Mann-Whitney U tests.

** : $p < .01$; *** : $p < .001$.