



HHS Public Access

Author manuscript

Neuron. Author manuscript; available in PMC 2016 November 04.

Published in final edited form as:

Neuron. 2015 November 4; 88(3): 449–460. doi:10.1016/j.neuron.2015.09.010.

The psychology and neuroscience of curiosity

Celeste Kidd and Benjamin Y. Hayden*

Department of Brain and Cognitive Sciences and Center for Visual Science, University of Rochester, Rochester, NY

SUMMARY

Curiosity is a basic element of our cognition, yet its biological function, mechanisms, and neural underpinning remain poorly understood. It is nonetheless a motivator for learning, influential in decision-making, and crucial for healthy development. One factor limiting our understanding of it is the lack of a widely agreed upon delineation of what is and is not curiosity; another factor is the dearth of standardized laboratory tasks that manipulate curiosity in the lab. Despite these barriers, recent years have seen a major growth of interest in both the neuroscience and psychology of curiosity. In this Perspective, we advocate for the importance of the field, provide a selective overview of its current state, and describe tasks that are used to study curiosity and information-seeking. We propose that, rather than worry about defining curiosity, it is more helpful to consider the motivations for information-seeking behavior and to study it in its ethological context.

Keywords

Curiosity; information-seeking; learning; Goldilocks effect; uncertainty

BACKGROUND

Curiosity is such a basic component of our natures that we are nearly oblivious to its pervasiveness in our lives. Consider, though, how much of our time we spend seeking and consuming information, whether listening to the news or music, browsing the internet, reading books or magazines, watching TV, movies, and sports, or otherwise engaging in activities not directly related to eating, reproduction, and basic survival. Our insatiable demand for information drives a much of the global economy and, on a micro-scale, motivates learning and drives patterns of foraging in animals. Its diminution is a symptom of depression, and its overexpression contributes to distractibility, a symptom of disorders such as attention-deficit/hyperactivity disorder. Curiosity is thought of as the noblest of human drives, and is just as often as it is denigrated as dangerous (as in the expression “curiosity killed the cat”). And despite its link with the most abstract human thoughts, some rudimentary forms of it can be observed even in the humble worm *C. elegans*.

*Corresponding author: Benjamin Hayden, Department of Brain and Cognitive Sciences, University of Rochester, Rochester NY 14620, benhayden@gmail.com.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Despite its pervasiveness, we lack even the most basic integrative theory of the basis, mechanisms, and purpose of curiosity. Nonetheless, as a psychological phenomenon, curiosity—and the desire for information more broadly—has attracted the interest of the biggest names in the history of psychology (e.g., James, 1913; Pavlov, 1927; Skinner, 1938). Despite this interest, only recently have psychologists and neuroscientists begun widespread and coordinated efforts to unlock its mysteries (e.g., Gottlieb et al., 2013; Gruber, Gelman & Ranganath, 2014; Kang et al., 2009). The present Perspective aims summarize this recent research, motivate new interested in the problem, and to tentatively propose a framework for future studies of the neuroscience and psychology of curiosity.

DEFINITION AND TAXONOMY OF CURIOSITY

One factor that has hindered the development of a formal study of curiosity is the lack of a single widely accepted definition of the term. In particular, many observers think that curiosity is a special type of the broader category of information-seeking. But carving out a formal distinction between the curiosity and information-seeking has proven difficult. As a consequence, much research that is directly relevant to the problem of curiosity does not use the term curiosity and instead focuses on what are considered to be distinct phenomena. These phenomena include, for example, play, exploration, reinforcement learning, latent learning, neophilia, and self-reported desire for information (e.g., Deci, 1975; Gruber, Gelman & Ranganath, 2014; Jirout & Klahr, 2012; Kang et al., 2009; Sutton & Barto, 1998; Tolman & Gleitman, 1949). Conversely, studies that do use the term curiosity range quite broadly in topic area. In laboratory studies, the term curiosity itself is broad enough to encompass both the desire for answers to trivia questions and the strategic deployment of gaze in free viewing (Gottlieb et al., 2013).

We consider this diversity of definitions to be both characteristic of a nascent field and healthy. Here we consider some classic views with an aim towards helping us think about how to study curiosity in the future.

Classic descriptions of curiosity

Philosopher and psychologist William James (1899) called curiosity “the impulse towards better cognition,” meaning that it is the desire to understand what you know that you do not. He noted that, in children, it drives them towards objects of novel, sensational qualities—that which is “bright, vivid, startling”. This early definition of curiosity, he said, later gives way to a “higher, more intellectual form”—an impulse towards more complete scientific and philosophic knowledge. Psychologist-educators G. Stanley Hall and Theodate L. Smith (1903) pioneered some of the earliest experimental work on the development of curiosity by collecting questionnaires and child biographies from mothers on the development of interest and curiosity. From these data, they describe children’s progression through four stages of development, starting with “passive staring” as early as the second week of life, on through “curiosity proper” at around the fifth month.

The history of studies of animal curiosity is nearly as long as the history of the study of human curiosity. Ivan Pavlov, for example, wrote about the spontaneous orienting behavior in dogs to novel stimuli (which he called the “What-is-it?” reflex) as a form of curiosity

(Pavlov, 1927). In the mid 20th century, exploratory behavior in animals began to fascinate psychologists, in part because of the challenge of integrating it into strict behaviorist approaches (e.g. Tolman, 1948). Some behaviorists counted curiosity as a basic drive, effectively giving up on providing a direct cause (e.g. Pavlov, 1927). This stratagem proved useful even as behaviorism declined in popularity. For example, this view was held by Harry Harlow—the psychologist best known for demonstrating that infant rhesus monkeys prefer the company of a soft, surrogate mother over a bare wire mother. Harlow referred to curiosity as a basic drive in and of itself—a “manipulatory motive”—that drives organisms to engage in puzzle-solving behavior that involved no tangible reward (e.g., Harlow, Blazek, & McClearn, 1956; Harlow, Harlow, & Meyer, 1950; Harlow & McClearn, 1954).

Psychologist Daniel Berlyne is among the most important figures in the 20th century study of curiosity. He distinguished between the types of curiosity most commonly exhibited by human and non-humans along two dimensions: perceptual versus epistemic, and specific versus diversive (Berlyne, 1954). Perceptual curiosity refers to the driving force that motivates organisms to seek out novel stimuli, which diminishes with continued exposure. It is the primary driver of exploratory behavior in non-human animals and potentially also human infants, as well as a possible driving force of human adults’ exploration. Opposite perceptual curiosity was epistemic curiosity, which Berlyne described as a drive aimed “not only at obtaining access to information-bearing stimulation, capable of dispelling uncertainties of the moment, but also at acquiring knowledge”. He described epistemic curiosity as applying predominantly to humans, thus distinguishing the curiosity of humans from that of other species (Berlyne, 1966).

The second dimension of curiosity that Berlyne described informational specificity. Specific curiosity referred to desire for a particular piece of information, while diversive curiosity referred to a general desire for perceptual or cognitive stimulation (e.g., in the case of boredom). For example, monkeys robustly exhibit specific curiosity when solving mechanical puzzles, even without food or any other extrinsic incentive (e.g., Davis, Settlege, & Harlow, 1950; Harlow, Harlow, & Meyer, 1950; Harlow, 1950). However, rats exhibit diversive curiosity when, devoid of any explicit task, they robustly prefer to explore unfamiliar sections of a maze (e.g., Dember, 1956; Hughes, 1968; Kivy, Earl, & Walker, 1956). Both specific and diversive curiosity were described as species-general information-seeking behaviors.

Contemporary views of curiosity

A common contemporary view of curiosity is that it is a special form of information-seeking distinguished by the fact that it is internally motivated (Loewenstein, 1994; Oudeyer & Kaplan, 2007). By this view, curiosity is strictly an *intrinsic* drive, while information-seeking refers more generally to a drive that can be either intrinsic or extrinsic. An example of an extrinsic type of information-seeking is paying a nominal price to know the outcome of a gamble before choosing it in order to make a more profitable choice. In other words, contexts in which agents seek information for immediately strategic reasons are not considered curiosity in the strict sense. While this definition is intuitively appealing (and

most consistent with the use of the term curiosity in everyday speech), it is accompanied by some problems.

For example, it is often difficult for an external observer to know whether a decision-maker is motivated intrinsically or extrinsically. Animals and preverbal children, for example, cannot tell us why they do what they do, and may labor under biased theories about the structure of their environment or other unknown cognitive constraints. Consider a child choosing between a safe door and a risky one (Butler, 1953). If the child chooses the risky option, should we call her curious or just risk-seeking? Or consider a rhesus monkey who performs a color discrimination task to obtain the opportunity to visually explore their environment. Perhaps the monkey is laboring under the assumption that the view of the environment offers some actionable information, and we should put him in the same place on the curiosity spectrum as the child (whatever that place is). To make things more complicated, perhaps the monkey has decided—or even experienced selective pressure—to favor a policy of information-seeking in most contexts. It would be a challenging philosophical problem to classify this behavior as true or ersatz curiosity by the intrinsic definition.

Thus, for the moment, we favor the rough and ready formulation of curiosity as a drive state for information. Decision-makers can be thought of as wanting information for several overlapping reasons just as they want food, water, and other basic goods. This drive may be internal or external, conscious or unconscious, or slowly evolved, or some mixture of the above. We hope that future work will provide a solid taxonomy of different factors that constitute our umbrella term.

Instead of figuring out the taxonomy, we advocate a different approach: we suggest that it is helpful to think about curiosity in the context of Tinbergen's Four Questions. Named after Dutch biologist Nikolaas Tinbergen, these questions are designed to provide four complementary scientific perspectives on any particular type of behavior (Tinbergen, 1963). These questions in turn offer four vantage points from which we can describe a behavior or a broad class of behaviors, even if its boundaries are not yet fully delineated. In this spirit, our Perspective will discuss current work on curiosity as seen through the lens of Tinbergen's Four Questions, here simplified to one word each: (1) function, (2) evolution, (3) mechanism, and (4) development.

THE FUNCTION OF CURIOSITY

Although information is intangible, it has real value to any organism with the capacity to make use of it. The benefits may accrue immediately or in the future; the delayed benefits require a learning system. Not surprisingly then, the most popular theory about the function of curiosity is to motivate learning. George Loewenstein (1994) described curiosity as “a cognitive induced deprivation that arises from the perception of a gap in knowledge and understanding.” Loewenstein's information gap theory holds that curiosity functions like other drive states, such as hunger, which motivates eating. Building on this theory, Loewenstein suggests that a small amount of information serves as a priming dose, which greatly increases curiosity. Consumption of information is rewarding but, eventually, when

enough information is consumed, satiation occurs and information serves to reduce further curiosity.

Loewenstein's idea is supported by a recent study by Kang and colleagues (Figure 1B, Kang et al., 2009). They found that curiosity about the answer to a trivia question is a U-shaped function of confidence about knowing that answer. Decision-makers were least curious when they had no clue about the answer and if they were extremely confident; they were most curious when they had some idea about the answer, but lacked confidence. In these circumstances, compulsion to know the answer was so great that they were even willing to pay for the information even though curiosity could have been sated for free after the session. (The neural findings of this study are discussed below.)

Kang and colleagues also found that curiosity enhances learning, consistent with the theory that the primary function of curiosity is to facilitate learning. This idea also motivated O'Keefe and Nadel's thinking about the factors that promote spatial learning in rodents (O'Keefe and Nadel, 1978). This idea is also popular in the education literature (e.g., Day, 1971; Engel, 2011, 2015; Gray, 2013), and has been for quite some time, as evidenced by attempts by education researchers to develop scales to quantify children's degree of curiosity, both generally and in specific learning materials (e.g., Harty & Beall, 1984; Jirout & Klahr, 2012; Pelz, Yung, & Kidd, 2015; Penney & McCann, 1964). One potential benefit of such research would be to improve education. More recently, the role of curiosity in enhancing learning is gaining adherents in cognitive science (see Gureckis & Markant, 2012, for a review). The idea is that allowing a learner to indulge their curiosity allows them to focus their effort on useful information that they do not yet possess. Further, there is a growing body of evidence suggesting that curiosity enables even infant learners to play an active role in optimizing their learning experiences (Oudeyer & Smith, in press). This work suggests that allowing a learner to expose the information they require themselves—which would be inaccessible via passive observation—may further benefit the learner by enhancing the encoding and retention of the new information.

THE EVOLUTION OF CURIOSITY

Information allows for better choices, more efficient search, more sophisticated comparisons, and better identification of conspecifics. Acquiring information, of course, is the primary evolutionary purpose of the sense organs, and has been a major driver of evolution for hundreds of millions of years. Complex organisms actively control their sense organs to maximize intake of information. For example, we choose our visual fixations strategically to learn about the things that are important to us in the context (Yarbus 1956; Gottlieb et al., 2012, 2013, 2014). Given its important role, it is not surprising that our visual search is highly efficient. It is nearly optimal when compared to an 'ideal searcher' that uses precise statistics of the visual scene to maximize search efficiency (Najemnik & Geisler 2005). Moreover, the strong base of information we have about the visual system makes it an appealing target for studies of curiosity (Gottlieb et al., 2013, 2014). Just as eye movements can be highly informative, our overt behaviors, including choice, can provide evidence for and against specific theories about how we seek information, which can in turn

help us understand the root causes of evolution. In this section we discuss the spectrum of basic information-seeking behaviors.

Elementary information-seeking

Even very simple organisms trade off information for rewards. While their information-seeking behavior is not typically categorized as curiosity, the simplicity of their neural systems makes them ideally suited for studies that may provide its foundation. For example, *C. elegans* is a roundworm whose nervous system contains 302 neurons and that actively forages for food, mostly bacteria. When placed on a new patch (such as a petri dish in a lab), it first explores locally (for about 15 minutes), then abruptly adjusts strategies and makes large, directed movements in a new direction (Calhoun, Chalasani, & Sharpee, 2014). This search strategy is more sophisticated and beneficial than simply moving towards food scents (or guesses about where food may be); instead, it provides better long-term payoff because it provides information as well. It maximizes a conjoint variable that includes both expected reward and information about the reward. This behavior, while computationally difficult, is not too difficult for worms. A small network of three neurons can plausibly implement it. Other organisms that have simple information-seeking behavior include crabs (Zeil, 1998), bees (Gould 1986; Dyer, 1991) ants (Wehner et al., 2002), and moths (Vergasolla et al., 2007). Information gained from such organisms can help us to understand how simple networks can perform information-seeking.

Information-tradeoff tasks

In primates (including humans), one convenient way to study information-seeking is the *k-arm bandit task* (Gittins & Jones, 1974, Figure 2). In this task, decision-makers are faced with a series of choices between stochastic rewards (Whittle, 1988). The optimal strategy requires adjudication between exploration (sampling to improve knowledge, and thus future choices) and exploitation (choosing known best options). Sampling typically gives lower immediate payoff but can provide information that improves choices in the future, leading to greater overall performance. Humans and monkeys can do quite well at this task (Daw et al., 2006; Pearson et al., 2009). One particular advantage of such tasks is that they allow for sophisticated formal models of information tradeoffs; this level of rigor is often absent in conventional curiosity studies (Averbeck, 2015).

Daw and colleagues showed that humans performing a 4-arm bandit task choose options probabilistically based on expected values of the options (a “softmax” policy, Daw, 2006). This probabilistic element causes them to occasionally explore other possibilities, leading them to better overall choices. The frontopolar cortex and intraparietal sulcus are significantly more active during exploration, whereas striatum and ventromedial prefrontal cortex (vmPFC) are more active during exploitative choices (Daw et al., 2006). These are canonical reward areas, thus these results link curiosity to the reward system (a theme that we will return to). They proposed that the activation of higher-level prefrontal regions during exploration indicates a control mechanism overriding the exploitative tendency.

In a similar task, neurons in the posterior cingulate cortex (PCC) have greater tonic firing rates on exploratory trials than on exploitative trials (even after controlling for reward

expectation, Pearson et al., 2009, Figure 2). Firing rates also predict adjustments from exploitative to exploratory strategy and vice versa. These results highlight the contribution of the PCC, a critical yet mostly mysterious hub of the reward system, in both the transition to exploration and in its maintenance (Pearson et al., 2011). PCC is linked to both reward and regulation of learning, thus underscoring the possible linkage between these processes and curiosity (Heilbronner & Platt, 2013; Hayden et al., 2008). PCC responses are also driven by the salience of an option, a factor that relates directly to its ability to motivate interest, rather than reward value per se (Heilbronner et al., 2013). The precuneus, a region adjacent to, and closely interconnected with, the PCC, was also associated with curiosity in one study: it is enlarged in capuchins that are particularly curious (Phillips, Subiaul, & Sherwood, 2012).

Above and beyond the strategic benefit of exploration, we have a tendency to seek out new and unfamiliar options, which may offer more information than familiar ones. The bandit task can be modified to measure this tendency (Wittmann et al., 2008). In one case, subjects chose between four different images on each trial; the identity of the images was arbitrary and served to distinguish the options. The value of each image was stable but stochastic, so sampling was required to learn its value. Some images were familiar, others were novel; however, image novelty had no special meaning in the context of the task. Nonetheless, subjects were more likely to choose novel images (that is, they motivated exploratory choices). This bias towards choosing novel images was mathematically expressible as a novelty bonus (Gittins & Jones, 1974). Interestingly, this novelty bonus increased the expected reward for the novel images (as measured by an increase in reward prediction error (RPE) signal in ventral striatum). These results support the idea that novelty-seeking reflects an injection into choice of motivation provided by the brain's reward systems.

Bandit tasks can also be used to measure the effect of strategic context on information-seeking. For example, if the information relates to future events that may not happen, it ought to be discounted. Thus, the horizon (the number of trials available to search the environment before it changes dramatically) matters (Wilson et al., 2014; see also Averbeck, 2015). Humans can adjust appropriately to changes in horizon: with longer horizons, subjects were more likely to choose an exploratory strategy than an exploitative one. Together, these results highlight the power and flexibility of bandit tasks as a way of studying information-seeking in a rigorous and highly quantifiable way.

Temporal resolution of uncertainty tasks

What about when the drive for information has no clear benefit? One convenient way to study this is to take advantage of the preference for immediate information about the outcome of a risky choice (Kreps & Porteus, 1978; Lieberman et al., 1997; Luhmann et al., 2008; Prokasy, 1956; Wyckoff, 1952). In a *temporal resolution of uncertainty* task, monkeys choose between two gambles with identical probabilities (50/50) and identical payoffs (a large or a small squirt of juice delayed by 2.25 seconds, Bromberg-Martin & Hikosaka, 2009). The only difference between the two gambles is that one offers immediate information about win vs. loss (that is, immediate temporal resolution of uncertainty) while in the other the information is delayed. The reward is delayed in both cases, so preference

for sooner reward would not affect choice. Despite the brevity of the delay, monkeys reliably choose the option with the immediate resolution of uncertainty (the informative option, Bromberg-Martin & Hikosaka, 2009, 2011; Blanchard et al., 2015). This preference for earlier temporal resolution of uncertainty is not strategic because the information cannot improve choices. Thus, these tasks satisfy a stricter notion of curiosity.

We modified this task to quantify the value of information by titrating the values of the rewards (Blanchard, Hayden, & Bromberg-Martin, 2015, Figure 3). In the *curiosity tradeoff task*, by determining the indifference point between informative and uninformative options, we found that the value of information about a reward is about 25% of the value of the reward itself—surprisingly high. This finding indicates that monkeys choose information even when it has a measurable cost. In addition, the value of information increases with the stakes. In other words, monkeys will pay more for information about a high stakes gamble than for information about a low stakes gamble. These results are similar to some recent findings observed in pigeons (Stagner and Zentall, 2010). Pigeons will choose a risky option that provides an average of 2 pellets over one that provides an average of 3 pellets as long as the one that provides 2 also provides what they call a discriminative cue—meaning a cue that reliably predicted whether a reward would come (see also Gipson et al., 2009).

Zentall and colleagues *did* make the link between their risk-seeking pigeons and human gamblers (Zentall & Stagner, 2011). This link is potentially important: curiosity is often mooted as an explanation for risk-seeking behavior (Bromberg-Martin & Hikosaka, 2009). Rhesus monkeys, for example, are often risk-seeking in laboratory tasks (Blanchard & Hayden, 2014; Heilbronner & Hayden 2013; Monosov & Hikosaka, 2012; O’Neill & Schultz, 2010; So & Stuphorn, 2012; Strait et al., 2014 and 2015). Risky choices provide information about the status of uncertain stimuli in the world, so animals may naturally seek such information. We trained monkeys to perform a gambling task in which both the location and value of a preferred high-variance option are uncertain; knowing the location of that option allowed the monkeys to perform better in the future, but knowing its value was irrelevant (Hayden, Pearson, & Platt, 2009). We found that, following choices of the low variance (and thus non-preferred) option, when it was too late to change anything, monkeys will spontaneously shift gaze to its position, suggesting they want to know information about it.

These findings demonstrate the power of the desire for temporal resolution of uncertainty as a motivator for choice, and thus as a potential tool for the study of information-seeking. This phenomenon is particularly useful because the information sought is demonstrably useless, making it a good potential model for more basic and fundamental (i.e. non-strategic) forms of information seeking than the bandit task. It is also, like the bandit task, one that works well in animals (meaning behavior is reliable and stable across large numbers of trials), so it has potential utility in circuit-level studies.

THE (NEURAL) MECHANISMS OF CURIOSITY

Tinbergen’s third question is about the proximate mechanism of a behavior. The mechanism of any behavior is in device that produces it—the brain.

As noted above, Kang and colleagues used a curiosity induction task to test Loewenstein's hypothesis that curiosity reflects an information gap (Loewenstein, 1994). Human subjects read trivia questions and rated their feelings of curiosity while undergoing fMRI (Kang et al., 2009). Brain activity in the caudate nucleus and inferior frontal gyrus (IFG) was associated with self-reported curiosity. These structures are activated by anticipation of many types of rewards, so these results suggest that curiosity elicits an anticipation of a reward state—consistent with Loewenstein's theory (Delgado et al., 2000, 2003, 2008; De Quervain et al., 2004; Fehr & Camerer, 2007; King-Casas et al., 2005; Rilling et al., 2002). Puzzlingly, the nucleus accumbens, which is one of the most reliably activated structures for reward anticipation, was not activated (Knutson et al., 2001). When the answer was revealed, activations generally were found in structures associated with learning and memory, such as parahippocampal gyrus and hippocampus. Again this is a bit puzzling, because classic structures that respond to receipt of reward were not particularly activated. In any case, the learning effect was particularly strong on trials on which subjects' guesses were incorrect—the trials on which learning was greatest.

Jepma et al. (2012) showed subjects blurry photos with ambiguous contents that piqued their curiosity; curiosity activated the anterior cingulate cortex and anterior insula - regions sensitive to aversive conditions (but to many other things too); resolution of curiosity activated striatal reward circuits. Like Kang and colleagues, they found that resolution of curiosity activated learning structures and also drove learning. However, the differences between the two studies were larger than the similarities. In the Jepma study, curiosity is a fundamentally aversive state, while in the Kang study it is pleasurable. Specifically, curiosity is seen as a lack of something wanted (information) and thus unpleasant, and this unpleasantness motivates information, which will alleviate it.

Gruber and colleagues (2014) measured brain activity while subjects answered trivia questions and rated their curiosity for each question. They were also shown interleaved photographs of neutral, unknown faces which acted as a probe for learning. When tested later, subjects recalled the faces shown in high curiosity trials better than faces shown on low curiosity trials. Thus, the curiosity state led to better learning, even for the things people weren't curious about. Curiosity drove activity in both midbrain (implying the dopaminergic regions) and nucleus accumbens; memory was correlated with midbrain and hippocampal activity. These results suggest that, although curiosity reflects intrinsic motivation, it is mediated by the same mechanisms as extrinsically motivated rewards.

Single unit recordings from the temporal resolution of uncertainty task further support this overlap. In this task, dopamine neuron activity (DA) is enhanced by the prospect of both a possible reward and early information. Dopamine neurons provide a key learning and motivation signal that is critical for many types of reward-related cognition (Redgrave & Gurney, 2006; Bromberg-Martin, Matsumoto & Hikosaka, 2010; Schultz & Dickinson, 2000). The phasic dopamine response is thought to serve as a general reward prediction error—indicating rewards or reward prospects of any type that are greater than expected (Schultz et al., 1997). Information is not a primary reward (as juice or water would be in this context), but is a more indirect kind of reward. The fact that dopamine neurons signal both primary and informational reward suggests that the dopamine response reflects an

integration of multiple reward components to generate an abstract reward response. This finding further suggests that dopamine responses not associated with rewards—such as surprising and aversive events—may reflect the value that information provides (Horvitz, 2000; Matsumoto & Hikosaka, 2009; Redgrave & Gurney, 2006;).

These results suggest that, to subcortical reward structures, informational value is treated the same as any other valued good. To further test this idea, the authors asked whether midbrain neurons encode information prediction error (Bromberg-Martin & Hikosaka, 2011). While the positive RPE is carried by DA neurons, its inverse, the negative RPE, is carried by neurons in the lateral habenula (LHb). They made use of this fact in task in which there was an option to choose a stochastically informative gamble, meaning it would provide (50/50 chance) valid or invalid information about the upcoming reward. They found that neurons in the LHb encode the unexpected occurrence of information and the unexpected denial of information—just as they do with basic rewards (water and juice).

Where does the domain-general curiosity signal come from? It has recently been proposed that the dopamine reward signal is constructed out of input signals originating in the orbitofrontal cortex (OFC), which in turn receives input from basic sensory and association structures (Öngür & Price, 2000; Schoenbaum et al., 2011; Takahashi et al., 2011; Rushworth et al., 2011). If OFC is an input to the evaluation system, then it should carry information about the reward value of curiosity but may not carry a single general reward signal. In other words, OFC may serve as a kind of workshop that represents elements of reward that can guide choice, but not a single domain general value signal. In the curiosity tradeoff task (see above and Figure 3), OFC neurons encode both the stakes of the gamble, and also the information value of the options (Blanchard, Hayden, & Bromberg-Martin, 2015). But it doesn't integrate them into a single value signal. Thus, at least within this one task, curiosity is computed separately from other factors that influence value and combined at a specific point (or points) in the pathway between the OFC and the DA nuclei.

THE DEVELOPMENT OF CURIOSITY

The fourth of Tinbergen's questions concerns development of a behavior. Curiosity has been central to the study of infant and child attention and learning, and a major focus in research on early education for decades (e.g., Berlyne, 1978; Dember & Earl, 1957; Kinney & Kagan, 1976; Sokolov, 1960). The world of infants is full of potential sources for learning, but they possess limited information-processing resources. Thus, infants must solve what is known as *the sampling problem*: their attentional mechanisms must select a subset of material from everything available in their environments in order to make learning tractable. Furthermore, they must sample in a way that ensures that learning is efficient, which is tricky considering the fact that what material is most useful changes as the infant gains more knowledge.

Infants enter the world with some simple, low-level heuristics for guiding their attention towards certain informative features of the world. Haith (1980) argued that these organizing principles for visual behavior are fundamentally stimulus-driven. For example, infants' gaze is pulled towards areas of high contrast, which is useful for detecting objects and perceiving

their shapes (e.g., Salapatek & Kessen, 1966), and motion onset, which is useful for detecting animacy (e.g., Aslin & Shea, 1990). Infants also have an innate bias to orient towards faces (e.g., Farroni et al., 2005; Johnson et al., 1991), which relay both social information and cues that guide language learning (e.g., Baldwin, 1993). While this desire for information is surely intrinsic, whether or not these low-level mechanisms that guide infants' early attentional behavior could be explained with curiosity depends on the chosen definition. If curiosity requires an explicit mental representation of the need for new information, these low-level heuristics do not qualify. However, a broader definition, which sees curiosity as any mechanism that guides an organism towards new information, regardless of mental substrate, they certainly do. Regardless of how you classify them, these attentional biases get the infant started down the road of knowledge acquisition.

Externally driven motivation is not sufficient. Learners also must adapt to changing needs as they build up and modify their mental representations of the world. Many early researchers posited that novelty was the primary stimulus feature of relevance for infants (e.g., Sokolov, 1960). Infants prefer novel stimuli in many paradigms, such as those used by Fantz (1964), the high-amplitude sucking procedure (Siqueland & DeLucia, 1969), and head-turn preference procedure (Kemler Nelson et al., 1995). Novelty preference is also seen in habituation procedures, in which infants' attention to a recurring stimulus decreases with lengthened exposure. Novelty theories, however, cannot account for infants' attested familiarity preferences, such as their affinity for their native languages and familiar faces (e.g., Bushnell et al., 1989; DeCasper & Spence, 1986).

Later theories sought to unify infants' novelty and familiarity preferences by explaining them in terms of infants' changing knowledge states. In other words, an infant's interest in a particular stimulus was theorized to be determined by that infant's particular mental status. For example, as infants attempt to encode various features of a visual stimulus, the efficiency or depth of this encoding process determines their subsequent preferences. Infants were theorized to exhibit a preference for stimuli that were partially—but not fully—encoded into memory (e.g., Dember & Earl, 1957; Hunter & Ames, 1988; Kinney & Kagan, 1976; Roder, Bushnell, & Sasseville, 2000; Rose et al., 1982; Wagner, S.H., & Sakovitsjkk, 1986). This idea recalls the fact that we are curious for things that we are moderately certain of (Kang et al., 2009).

Among these theories was Kinney and Kagan's *moderate discrepancy hypothesis*, which suggested that infants preferentially attend to stimuli that were "optimally discrepant," meaning those that were just the right amount of distinguishable from mental representations that the infant already possessed (Kinney & Kagan, 1976). Under Dember and Early's *theory of choice/preference*, learners seek stimuli that match their preferred level of complexity, which increases over time as they build up mental representations and acquire more knowledge (Dember & Earl, 1957). Berlyne, similarly, noted that complexity-driven preferences could represent an optimal strategy for learning (Berlyne, 1960). Such *processing-based theories* of curiosity predict that learners will exhibit a U-shaped pattern of preference for stimulus complexity, where complexity is defined in terms of the learner's current set of mental representations. The theories predict that learners will preferentially select stimuli of an intermediate level of complexity—material that is neither overly simple

(already encoded into memory) nor overly complex (too disparate from existing representations already encoded into memory).

Recent infant research supports these accounts (e.g., Kidd, Piantadosi, & Aslin, 2012, 2014, Figure 4). We showed 7- and 8-month-old infants visual event sequences of varying complexity, as measured by an idealized learning model, and measured points at which infants' attention drifted (as indicated by looks away from the display). We found that infants' probability of looking away was greatest to events of either very low information content (highly predictable) or very high information content (highly surprising). This attentional strategy holds in multiple types of visual displays (Kidd, Piantadosi, & Aslin, 2012), for auditory stimuli (Kidd, Piantadosi, & Aslin, 2014), and even within individual infants (Piantadosi, Kidd, & Aslin, 2014). These results suggest that infants implicitly decide to direct attention in order to maintain intermediate rates of information absorption. This attentional strategy likely prevents them from wasting cognitive resources on overly predictable or overly complex events, thus helping to maximize their learning potential.

Related findings show that children structure their play in a way that reduces uncertainty and allows them to discover causal structures in the world (e.g., Schulz & Bonawitz, 2007). This work is in line with earlier theories of Jean Piaget (1930) that asserted that the purpose of curiosity and play was to “construct knowledge” through interactions with the world. If curiosity aims to reduce uncertainty in the world, we would expect learners to exhibit increased curiosity to stimuli in the world that they do not understand. In fact, this is a behavior that is well attested in recent developmental psychology studies, such as work by Bonawitz and colleagues (Bonawitz, van Schijndel, Friel, & Schulz, 2012) that demonstrates that children prefer to play with toys that violate their expectations. Children also exhibit increased curiosity outside of pedagogical contexts, in the absence of explicitly given explanations (Bonawitz et al., 2011). In an experiment in which Bonawitz and colleagues gave children a novel toy to explore, either prefaced or not with partial instruction of how the toy works, children played for longer and discovered more of the toys' functions in the non-pedagogical conditions.

In line with the idea that the function of curiosity is to reduce uncertainty, children exhibit increased interest in situations with high degrees of uncertainty, such as preferentially playing with toys whose underlying mechanisms are not yet understood. Perhaps even more impressively, Schulz and Bonawitz (2007) found that children preferentially engaged with toys that allowed them to deconfound potential causal variables underlying toys' inner workings. In these experiments, Schulz and Bonawitz had children play with toys consisting of boxes and levers. In both the confounded and unconfounded conditions, the researcher would help a child play with a red box with two levers. In the confounded condition, the researcher and the child each pressed down on a lever at the same time and, in response, two small puppets (a chick and a pom-pom) popped out of the top of the red box (Figure 5). The puppets' location—dead center—was not informative about which of the two levers caused each one to rise. In the unconfounded conditions, the researcher and child took turns pressing down on their respective levers one at a time *or* the researcher demonstrated each lever independently; thus, in both cases, it was clear which lever controlled each puppet. After this demonstration, the researcher uncovered a second, yellow box. After the

demonstration and yellow-box reveal, children were left alone and instructed to play in the researcher's absence for 60 seconds. During this period, children in the confounded condition preferentially explored the demonstrated red box over the novel yellow one.

The idea that children structure their play in a way that is sensitive to information gain is further bolstered by a recent study by Cook and colleagues (Cook, Goodman, & Schulz, 2011). They manipulated the ambiguity of various causal variables for a toy box that played music when certain—but not all—beads were placed on top of it. A researcher initially demonstrated how the box worked by placing a pair of connected beads on top, thereby making it ambiguous which of the two beads was causally responsible for the music playing. Children were effective at both selecting and designing informative interventions to figure out the underlying causal structure when it was unclear from the demonstration. When given ambiguous evidence, children tested individual beads when possible and—even more impressively—when bead pair was permanently stuck connected together, children held them such that only one side was touching the box in order to isolate the effect of that particular bead on the box.

This hypothesis-testing behavior is now widely attested in the developmental psychology literature. Children appear to structure their play in order to deconfound variables when causal mechanisms at play in the world are unclear (e.g., Denison et al., 2013; Gopnik, Meltzoff, & Kuhl, 1999; Gweon et al., 2014; Schulz, Gopnik, & Glymour, 2007; van Schijndel et al., 2015), and also make efficient use of information that they encounter in the world to learn correct causal structures (e.g., Gopnik & Schulz, 2007; Gopnik et al., 2001). These findings are important because they highlight the fact that children's curiosity appears specifically well suited to teaching them about the causal structure of the world. Thus, these strategic information-seeking behaviors in young children are far more sophisticated than the simple attentional heuristics that characterize early infant attention.

CONCLUSION

Curiosity has long fascinated laymen and scholars alike, but remains poorly understood as a psychological phenomenon. We argue that one factor impeding our understanding has been too much focus on delineating what is and is not curiosity. Another has been too much emphasis on taxonomy. These divide-then-conquer approaches are premature because they do not rely on empirical data. Perhaps the plethora of definitions and schemes attests more to differences in scholars' intuitions than to differences in their data. Thus we recommend that the definition stage follow a relatively solid characterization of curiosity, defined as broadly as possible. For this reason, we are reluctant to commit to a strict definition now. This approach has risks, of course. It means that there will be a variety of studies using similar terms to describe different phenomena, and different terms to describe the same phenomena, which can be confusing. Nonetheless, we think the benefits of open-mindedness outweigh the costs.

Broadening the scope of inquiry has several advantages. First, it allows us to study information-seeking in non-humans, including monkeys, rats, and even roundworms. Animal techniques allow for a granular view of mechanism, allows a greater range of

manipulations, and allows cross-species comparisons. Second, it allows us to temporarily put aside speculation about decision-makers' motivations and focus on other questions. Third, by refusing to isolate curiosity from other cognitive processes, we can make bridges with other phenomena, especially reward and learning. Finally, it lets us take advantage of powerful new tasks invented in the past decade for studying the cognitive neuroscience of information-seeking.

Tinbergen's Four Questions are designed to provide a way to explain the causes of any behavior. This approach already provides a convenient framework for considering the knowledge we have so far. In the domain of *function*, it seems clear that curiosity serves to motivate acquisition of knowledge and learning. In the domain of *evolution*, it seems that curiosity can tentatively be said to improve performance, yielding fitness benefits to organisms with it, and is likely to be an evolved trait. In the domain of *mechanisms*, it seems that the drive for information augments internal representations of value, thus biasing decision-makers towards informative options and actions. It also seems that curiosity activates learning systems in the brain. In the domain of *development*, we can infer that curiosity is critical for learning and that it reflects both external features and internal representations of own knowledge.

In the future we hope to see answers to some of these questions:

- In what ways does curiosity resemble other basic drive states? How does it differ? To what extent is curiosity fundamentally different from drives like hunger and thirst?
- What is the most useful taxonomy of curiosity? How well does Berlyne's categorization hold up? What factors unite distinct forms of curiosity?
- How is curiosity controlled? What factors govern curiosity, and how does the brain integrate these factors into decision-making to produce decisions? To what extent is curiosity context-dependent?
- To what extent does curiosity in nematodes overlap (if at all) with curiosity in children? How useful is it to think of curiosity as being a single construct across a broad range of taxa?
- Does our continuing curiosity in adulthood serve a purpose or is it vestigial? Does continued curiosity serve to maintain cognitive abilities throughout adulthood?
- What is the link between curiosity and learning?
- Why and how is curiosity affected by diseases like depression and ADHD? Can sensitive measures of curiosity be used to predict and measure cognitive decline, senility, and Alzheimers' Disease?

We can already sketch out rough guesses about how some of these questions will be answered. For example, we anticipate that, although useful in the past, Berlyne's categories will be replaced with other, differently-formulated subtypes, and that these newer ones will be motivated by new neural and developmental data. We suspect that curiosity serves a similar purpose in adulthood as it does in childhood, albeit in perhaps a more refined way.

Even as adults we need to continue to adjust our understanding of the world. Finally, we are optimistic that scientists will eventually uncover a consistent set of principles that characterize curiosity across a wide range of taxa.

Acknowledgments

This research was supported by a R01 (DA038615) to BYH. We thank Sarah Heilbronner, Steve Piantadosi, Shraddha Shah, Maya Wang, Habiba Azab, and Maddie Pelz for helpful comments.

REFERENCES

- Aslin RN, Shea SL. Velocity thresholds in human infants: implications for the perception of motion. *Developmental Psychology*. 1990; 26:589–598.
- Averbeck B. Theory of choice in bandit, information sampling and foraging tasks. *PLoS Computational Biology*. 2015; 11(3):e1004164. [PubMed: 25815510]
- Baldwin DA. Infants' ability to consult the speaker for clues to word reference. *Journal of Child Language*. 1993; 20(2):395–418. [PubMed: 8376476]
- Berlyne DE. A theory of human curiosity. *British Journal of Psychology*. 1954; 45(3):180–191. [PubMed: 13190171]
- Berlyne DE. The influence of complexity and novelty in visual figures on orienting responses. *Journal of Experimental Psychology*. 1958; 55:289–296. [PubMed: 13513951]
- Berlyne, DE. Conflict, arousal, and curiosity. New York: McGraw-Hill; 1960.
- Berlyne DE. Curiosity and exploration. *Science*. 1966; 153:25–33. [PubMed: 5328120]
- Berlyne DE. Curiosity and learning. *Motivation and Emotion*. 1978; 2(2):97–175.
- Blanchard TC, Hayden BY. Neurons in dorsal anterior cingulate cortex signal postdecisional variables in a foraging task. *The Journal of Neuroscience*. 2014; 34(2):646–655. [PubMed: 24403162]
- Blanchard TC, Hayden BY, Bromberg-Martin ES. Orbitofrontal Cortex Uses Distinct Codes for Different Choice Attributes in Decisions Motivated by Curiosity. *Neuron*. 2015; 85(3):1–13. [PubMed: 25569341]
- Bonawitz EB, Shafto P, Gweon H, Goodman N, Spelke E, Schulz LE. The double-edged sword of pedagogy: Teaching limits children's spontaneous exploration and discovery. *Cognition*. 2011; 120(3):322–330. [PubMed: 21216395]
- Bonawitz EB, van Schijndel T, Friel D, Schulz L. Balancing theories and evidence in children's exploration, explanations, and learning. *Cognitive Psychology*. 2012; 64(4):215–234. [PubMed: 22365179]
- Bromberg-Martin ES, Hikosaka O. Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron*. 2009; 63(1):119–126. [PubMed: 19607797]
- Bromberg-Martin ES, Matsumoto M, Hikosaka O. Dopamine in motivational control: rewarding, aversive, and alerting. *Neuron*. 2010; 68(5):815–834. [PubMed: 21144997]
- Bromberg-Martin ES, Hikosaka O. Lateral habenula neurons signal errors in the prediction of reward information. *Nature Neuroscience*. 2011; 14(9):1209–1216. [PubMed: 21857659]
- Bushnell EW, Sai F, Mullin JT. Neonatal recognition of the mother's face. *British Journal of Developmental Psychology*. 1989; 7:3–15.
- Butler RA. Discrimination learning by rhesus monkeys to visual-exploration motivation. *Journal of Comparative and Physiological Psychology*. 1953; 46(2):95. [PubMed: 13044866]
- Calhoun AJ, Chalasani SH, Sharpee TO. Maximally informative foraging by *Caenorhabditis elegans*. *Elife*. 2014; 3:e04220.
- Cook C, Goodman ND, Schulz LE. Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*. 2011; 120:341–349. [PubMed: 21561605]
- Davis RT, Settlage PH, Harlow HF. Performance of Normal and Brain-Operated Monkeys on Mechanical Puzzles with and without Food Incentive. *Journal of Genetic Psychology*. 1950; 77:305–311. [PubMed: 14803690]

- Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ. Cortical substrates for exploratory decisions in humans. *Nature*. 2006; 441(7095):876–879. [PubMed: 16778890]
- Day, HI. The measurement of specific curiosity. In: Day, HI.; Berlyne, DE.; Hunt, DE., editors. *Intrinsic motivation: A new direction in education*. New York: Holt, Rinehart & Winston; 1971.
- DeCasper AJ, Spence MJ. Prenatal maternal speech influences newborns' perception of speech sounds. *Infant Behavioral Development*. 1986; 9:133–150.
- Deci, EL. *Intrinsic Motivation*. New York: Plenum; 1975.
- Delgado MR, Nystrom LE, Fissell C, Noll DC, Fiez JA. Tracking the hemodynamic responses to reward and punishment in the striatum. *Journal of Neurophysiology*. 2000; 84(6):3072–3077. [PubMed: 11110834]
- Delgado MR, Locke HM, Stenger VA, Fiez JA. Dorsal striatum responses to reward and punishment: effects of valence and magnitude manipulations. *Cognitive, Affective, & Behavioral Neuroscience*. 2003; 3(1):27–38.
- Delgado MR, Schotter A, Ozbay EY, Phelps EA. Understanding overbidding: using the neural circuitry of reward to design economic auctions. *Science*. 2008; 321(5897):1849–1852. [PubMed: 18818362]
- Dember WN, Earl RW. Analysis of Exploratory, Manipulatory, and Curiosity Behaviors. *Psychological Review*. 1957; 64:91–96. [PubMed: 13420283]
- Denison S, Bonawitz EB, Gopnik A, Griffiths T. Rational Variability in Children's Causal Inferences: The Sampling Hypothesis. *Cognition*. 2013; 126(2):285–300. [PubMed: 23200511]
- De Quervain DJF, Fischbacher U, Treyer V, Schellhammer M, Schnyder U, Buck A, Fehr E. The neural basis of altruistic punishment. *Science*. 2004; 305(5688):1254. [PubMed: 15333831]
- Engel S. Children's Need to Know: Curiosity in Schools. *Harvard Educational Review*. 2011; 81(4): 625–645.
- Engel, S. *The Hungry Mind: The Origins of Curiosity in Childhood*. Cambridge, MA: Harvard University Press; 2015.
- Farroni T, Johnson MH, Menon E, Zulian L, Faraguna D, Csibra G. Newborns' preference for face-relevant stimuli: Effects of contrast polarity. *Proceedings of the National Academy of Sciences*. 2005; 102(47):17245–17250.
- Fehr E, Camerer CF. Social neuroeconomics: the neural circuitry of social preferences. *Trends in Cognitive Sciences*. 2007; 11(10):419–427. [PubMed: 17913566]
- Genovesio A, Tsujimoto S, Navarra G, Falcone R, Wise SP. Autonomous encoding of irrelevant goals and outcomes by prefrontal cortex neurons. *The Journal of Neuroscience*. 2014; 34(5):1970–1978. [PubMed: 24478376]
- Gipson CD, Alessandri JJ, Miller HC, Zentall TR. Preference for 50% reinforcement over 75% reinforcement by pigeons. *Learning and Behavior*. 2009; 37(4):289–298. [PubMed: 19815925]
- Gittins, J.; Jones, D. *Progress in statistics*. Vol. 2. Amsterdam: North-Holland; 1974. p. 9
- Gottlieb J. Attention, learning, and the value of information. *Neuron*. 2012; 76(2):281–295. [PubMed: 23083732]
- Gopnik, A.; Meltzoff, AN.; Kuhl, PK. *The Scientist in the Crib: Minds, Brains, & How Children Learn*. New York: William Morrow & Co.; 1999.
- Gopnik, A.; Schulz, L., editors. *Causal Learning: Psychology, Philosophy, and Computation*. Oxford, UK: Oxford University Press; 2007.
- Gopnik A, Sobel DM, Schulz LE, Gylmour C. Causal Learning Mechanisms in Very Young Children: Two-, Three-, and Four-Year-Olds Infer Causal Relations from Patterns of Variation and Covariation. *Developmental Psychology*. 2001; 37(5):620–629. [PubMed: 11552758]
- Gottlieb J, Hayhoe M, Hikosaka O, Rangel A. Attention, reward, and information seeking. *Journal of Neuroscience*. 2014; 34(46):15497–15504. [PubMed: 25392517]
- Gottlieb J, Oudeyer P-Y, Lopes M, Baranes A. Information seeking, curiosity and attention: computational and neuronal mechanisms. *Trends in Cognitive Science*. 2013; 17(11):585–593.
- Gould JL. The locale map of honey bees: do insects have cognitive maps? *Science*. 1986; 232(4752): 861–863. [PubMed: 17755968]

- Gray, Peter. *Free to Learn: Why Unleashing the Instinct to Play Will Make Our Children Happier, More Self-Reliant, and Better Students for Life*. New York: Basic Books; 2013.
- Gruber MJ, Gelman BD, Ranganath C. States of Curiosity Modulate Hippocampus-Dependent Learning via the Dopaminergic Circuit. *Neuron*. 2014; 84(2):486–496. [PubMed: 25284006]
- Gureckis TM, Markant DB. Self-Directed Learning: A Cognitive and Computational Perspective. *Perspectives on Psychological Science*. 2012; 7(5):464–481. [PubMed: 26168504]
- Gweon H, Pelton H, Konopka JA, Schulz LE. Sins of omission: Children selectively explore when agents fail to tell the whole truth. *Cognition*. 2014; 132:335–341. [PubMed: 24873737]
- Haith, MM. *Rules That Babies Look By*. Mahwah, NJ: L. Erlbaum Associates, Inc.; 1980.
- Hall GS, Smith TL. Curiosity and Interest. *The Pedagogical Seminary*. 1903; 10(3):315–358.
- Harlow HF. Learning and satiation of response in intrinsically motivated complex puzzle performance by monkeys. *Journal of Computation and Physiological Psychology*. 1950; 43:289–294.
- Harlow HF, Harlow MK, Meyer DR. Learning motivated by a manipulation drive. *Journal of Experimental Psychology*. 1950; 40:228–234. [PubMed: 15415520]
- Harty H, Beall D. Toward the Development of a Children's Science Curiosity Measure. *Journal of Research in Science Teaching*. 1984; 21(4):425–436.
- Hayden BY, Platt ML. Neurons in anterior cingulate cortex multiplex information about reward and action. *The Journal of Neuroscience*. 2010; 30(9):3339–3346. [PubMed: 20203193]
- Hayden BY, Nair AC, McCoy AC, Platt ML. Posterior cingulate cortex mediates outcome-contingent allocation of behavior. *Neuron*. 2008; 60:19–25. [PubMed: 18940585]
- Heilbronner SR, Hayden BY. Contextual factors explain risk-seeking preferences in rhesus monkeys. *Frontiers in Neuroscience*. 2013; 7
- Heilbronner SR, Hayden BY, Platt ML. Decision salience signals in posterior cingulate cortex. *Frontiers in Decision Neuroscience*. 2011; 5 article 55.
- Heilbronner SR, Platt ML. Causal evidence of performance monitoring by posterior cingulate cortex during learning. *Neuron*. 2013; 80(6):1384–1391. [PubMed: 24360542]
- Horvitz JC. Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neuroscience*. 2000; 96(4):651–656. [PubMed: 10727783]
- Hughes RN. Behaviour of male and female rats with free choice of two environments differing in novelty. *Animal Behavior*. 1968; 16(1):92–96.
- Hunter MA, Ames EW. A multifactor model of infant preferences for novel and familiar stimuli. *Advances in Infancy Research*. 1988; 5:69–95.
- James, W. *Talks to Teachers on Psychology: And to Students on Some of Life's Ideals*. New York: Henry Holt & Company; 1899.
- James W. *The principles of psychology*. 1913; II
- Jepma M, Verdonschot RG, van Steenbergen H, Rombouts SARB, Nieuwenhuis S. Neural mechanisms underlying the induction and relief of perceptual curiosity. *Frontiers in Behavioral Neuroscience*. 2012; 6(5):1–9. [PubMed: 22279431]
- Jirout J, Klahr D. Children's Scientific Curiosity: In Search of an Operational Definition of an Elusive Concept. *Developmental Review*. 2012; 32(2):125–160.
- Johnson MH, Dziurawiec S, Ellis HD, Morton J. *Cognition*. 1991; 40:1–19. [PubMed: 1786670]
- Kang MJ, Hsu M, Krajbich IM, Loewenstein G, McClure SM, Wang JTY, Camerer CF. The wick in the candle of learning epistemic curiosity activates reward circuitry and enhances memory. *Psychological Science*. 2009; 20(8):963–973. [PubMed: 19619181]
- Kemler Nelson DG, Jusczyk PW, Mandel DR, Myers J, Turk A, Gerken L. The Head-Turn Preference Procedure for Testing Auditory Perception. *Infant Behavior and Development*. 1995; 18:111–116.
- Kidd C, Piantadosi ST, Aslin RN. The Goldilocks Effect in Infant Auditory Cognition. *Child Development*. 2014; 85(5):1795–1804. [PubMed: 24990627]
- Kidd C, Piantadosi ST, Aslin RN. The Goldilocks Effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLOS ONE*. 2012; 7(5):e36399. [PubMed: 22649492]

- King-Casas B, Tomlin D, Anen C, Camerer CF, Quartz SR, Montague PR. Getting to know you: reputation and trust in a two-person economic exchange. *Science*. 2005; 308(5718):78–83. [PubMed: 15802598]
- Kinney DK, Kagan J. Infant attention to auditory discrepancy. *Child Development*. 1976; 47:155–164. [PubMed: 954493]
- Knutson B, Adams CM, Fong GW, Hommer D. Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *Journal of Neuroscience*. 2001; 21(16):RC159. [PubMed: 11459880]
- Kreps DM, Porteus EL. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica: Journal of the Econometric Society*. 1978; 46(1):185–200.
- Lieberman DA, Cathro JS, Nichol K, Watson E. The role of S– in human observing behavior: Bad news is sometimes better than no news. *Learning and Motivation*. 1997; 28(1):20–42.
- Loewenstein G. The Psychology of Curiosity: A Review and Reinterpretation. *Psychological Bulletin*. 1994; 116(1):75–98.
- Luhmann CC, Chun MM, Yi DJ, Lee D, Wang XJ. Neural dissociation of delay and uncertainty in intertemporal choice. *The Journal of Neuroscience*. 2008; 28(53):14459–14466. [PubMed: 19118180]
- Matsumoto M, Hikosaka O. Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature*. 2009; 459(7248):837–841. [PubMed: 19448610]
- Monosov IE, Hikosaka O. Regionally distinct processing of rewards and punishments by the primate ventromedial prefrontal cortex. *The Journal of Neuroscience*. 2012; 32(30):10318–10330. [PubMed: 22836265]
- Najemnik J, Geisler WS. Optimal eye movement strategies in visual search. *Nature*. 2005; 434(7031):387–391. [PubMed: 15772663]
- O’Keefe, J.; Nadel, L. *The Hippocampus as a Cognitive Map*. Oxford, UK: Clarendon Press; 1978.
- O’Neill M, Schultz W. Coding of reward risk by orbitofrontal neurons is mostly distinct from coding of reward value. *Neuron*. 2010; 68(4):789–800. [PubMed: 21092866]
- Öngür D, Price JL. The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cerebral Cortex*. 2000; 10(3):206–219. [PubMed: 10731217]
- Oudeyer P-Y, Kaplan F. What is intrinsic motivation? A typology of computational approaches. *Frontiers in Neurorobotics*. 2007; 1:6. [PubMed: 18958277]
- Oudeyer P-Y, Smith L. How Evolution May Work through Curiosity-Driven Developmental Process. *Topics in Cognitive Science*. (In press).
- Padoa-Schioppa C. Neurobiology of economic choice: a good-based model. *Annual Review of Neuroscience*. 2011; 34:333.
- Pavlov, IP. *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex*. Oxford: Oxford University Press; 1927.
- Pearson JM, Hayden BY, Raghavachari S, Platt ML. Neurons in posterior cingulate cortex signal exploratory decisions in a dynamic multioption choice task. *Current Biology*. 2009; 19(18):1532–1537. [PubMed: 19733074]
- Pearson JM, Heilbronner SR, Barack DL, Hayden BY, Platt ML. Posterior cingulate cortex: adapting behavior to a changing world. *Trends Cognitive Science*. 2011; 15(4):143–151.
- Penney RK, McCann B. The Children’s Reactive Curiosity Scale. *Psychological Reports*. 1964; 15:323–334.
- Pelz M, Yung A, Kidd C. Quantifying Curiosity and Exploratory Play on Touchscreen Tablets. *Proceedings of the IDC 2015 Workshop on Digital Assessment and Promotion of Children’s Curiosity*. 2015
- Phillips KA, Subiaul F, Sherwood CC. Curious monkeys have increased gray matter density in the precuneus. *Neuroscience Letters*. 2012; 518(2):172–175. [PubMed: 22579821]
- Piantadosi ST, Kidd C, Aslin RN. Rich Analysis and Rational Models: Inferring individual behavior from infant looking data. *Developmental Science*. 2014; 17(3):321–337. [PubMed: 24750256]

- Prokasy WF Jr. The acquisition of observing responses in the absence of differential external reinforcement. *Journal of Comparative and Physiological Psychology*. 1956; 49(2):131. [PubMed: 13319523]
- Redgrave P, Gurney K. The short-latency dopamine signal: a role in discovering novel actions? *Nature Reviews Neuroscience*. 2006; 7(12):967–975. [PubMed: 17115078]
- Rilling JK, Gutman DA, Zeh TR, Pagnoni G, Berns GS, Kilts CD. A neural basis for social cooperation. *Neuron*. 2002; 35(2):395–405. [PubMed: 12160756]
- Roder BJ, Bushnell EW, Sasseville AM. Infants' preferences for familiarity and novelty during the course of visual processing. *Infancy*. 2000; 1:491–507.
- Rose SA, Gottfried AW, Melloy-Carminar P, Bridger WH. Familiarity and novelty preferences in infant recognition memory: Implications for information processing. *Developmental Psychology*. 1982; 18:704–713.
- Rushworth MFS, Noonan MP, Boorman ED, Walton ME, Behrens TE. Frontal Cortex and Reward-Guided Learning and Decision-Making. *Neuron*. 2011; 70(6):1054–1069. [PubMed: 21689594]
- Salapatek P, Kessen W. Visual scanning of triangles by the human newborn. *Journal of Experimental Child Psychology*. 1966; 3:1955–1967.
- Schoenbaum G, Takahashi Y, Liu TL, McDannald MA. Does the orbitofrontal cortex signal value? *Annals of the New York Academy of Sciences*. 2011; 1239(1):87–99. [PubMed: 22145878]
- Schulz L, Bonawitz EB. Serious fun: Preschoolers play more when evidence is confounded. *Developmental Psychology*. 2007; 43(4):1045–1050. [PubMed: 17605535]
- Schulz L, Gopnik A, Glymour C. Preschool children learn about causal structure from conditional interventions. *Developmental Science*. 2007; 10(3):322–332. [PubMed: 17444973]
- Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science*. 1997; 275(5306):1593–1599. [PubMed: 9054347]
- Schultz W, Dickinson A. Neuronal coding of prediction errors. *Annual Review of Neuroscience*. 2000; 23(1):473–500.
- Skinner BF. *The behavior of organisms: An experimental analysis*. 1938
- Sokolov, E. *Perception and the Conditioned Reflex*. Oxford, England: Pergamon; 1963.
- So N, Stuphorn V. Supplementary eye field encodes reward prediction error. *The Journal of Neuroscience*. 2012; 32(9):2950–2963. [PubMed: 22378869]
- Strait CE, Blanchard TC, Hayden BY. Reward value comparison via mutual inhibition in ventromedial prefrontal cortex. *Neuron*. 2014; 82(6):1357–1366. [PubMed: 24881835]
- Strait CE, Slezzer BJ, Hayden BY. Signatures of value comparison in ventral striatum neurons. *PLoS Biology*. 2015 (in press).
- Stagner JP, Zentall TR. Suboptimal choice behavior by pigeons. *Psychonomic Bulletin and Review*. 2010; 17(3):412–416. [PubMed: 20551367]
- Sutton, RS.; Barto, AG. *Introduction to reinforcement learning*. MIT Press; 1998.
- Takahashi YK, Roesch MR, Wilson RC, Toreson K, O'Donnell P, Niv Y, Schoenbaum G. Expectancy-related changes in firing of dopamine neurons depend on orbitofrontal cortex. *Nature Neuroscience*. 2011; 14(12):1590–1597. [PubMed: 22037501]
- Tinbergen, Nikolaas. *On Aims and Methods of Ethology*. *Zeitschrift für Tierpsychologie*. 1963; 20:410–433.
- Tolman EC. Cognitive maps in rats and men. *Psychological Review*. 1948; 55(4):189. [PubMed: 18870876]
- Tolman EC, Gleitman H. Studies in learning and motivation. *Journal of Experimental Psychology*. 1949; 39(6):810–819. [PubMed: 15398592]
- van Schijndel TJP, Visser I, van Bers BMCW, Raijmakers MEJ. Preschoolers perform more informative experiments after observing theory-violating evidence. *Journal of Experimental Child Psychology*. 2015; 131:104–119. [PubMed: 25544394]
- Vergassola M, Villermaux E, Shraiman BI. 'Infotaxis' as a strategy for searching without gradients. *Nature*. 2007; 445(7126):406–409. [PubMed: 17251974]
- Wagner SH, Sakovitsjkk LJ. A process analysis of infant visual and crossmodal recognition memory: Implications for an amodal code. *Advances in Infancy Research*. 1986; 4:195–217.

- Wehner R, Gallizzi K, Frei C, Vesely M. Calibration processes in desert ant navigation: vector courses and systematic search. *Journal of Comparative Physiology A*. 2002; 188(9):683–693.
- Whittle P. Restless bandits: Activity allocation in a changing world. *Journal of Applied Probability*. 1988:287–298.
- Wilson RC, Geana A, White JM, Ludvig EA, Cohen JD. Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*. 2014; 143(6):2074. [PubMed: 25347535]
- Wittmann BC, Daw ND, Seymour B, Dolan RJ. Striatal activity underlies novelty-based choice in humans. *Neuron*. 2008; 58(6):967–973. [PubMed: 18579085]
- Wyckoff LB Jr. The role of observing responses in discrimination learning. Part I. *Psychological Review*. 1952; 59(6):431. [PubMed: 13004146]
- Yarbus AL. The motion of the eye in the process of changing points of fixation. *Biofizika*. 1956; 1:76–78.
- Zeil J. Homing in fiddler crabs (*Uca lactea annulipes* and *Uca vomeris*: Ocypodidae). *Journal of Comparative Physiology A*. 1998; 183(3):367–377.
- Zentall TR, Stagner J. Maladaptive choice behaviour by pigeons: an animal analogue and possible mechanism for gambling (sub-optimal human decision-making behaviour). *Proceedings of the Royal Society B. Biological Sciences*. 2011; 278(1709):1203–1208. [PubMed: 20943686]

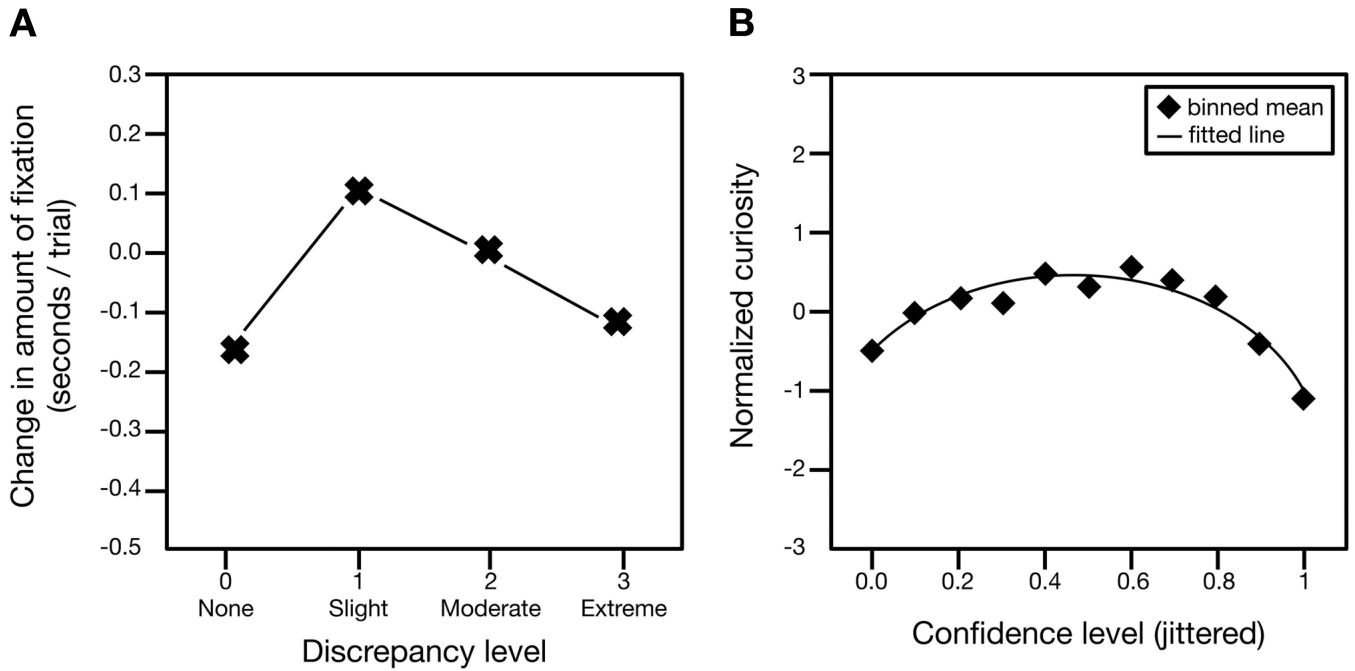


Figure 1.

A: Data from Kinney & Kagan (1976). Attention to auditory stimuli shows an inverted U-shaped pattern, with infants making the most fixations to auditory stimuli estimated to be moderately discrepant from the auditory stimuli for which infants already possessed mental representations. B: Data from Kang et al. (2009): Subjects were most curious about the answers to trivia questions for which they were moderately confident about their answers. This pattern suggests that subjects exhibited the greatest curiosity for information that was partially—but not fully—encoded.

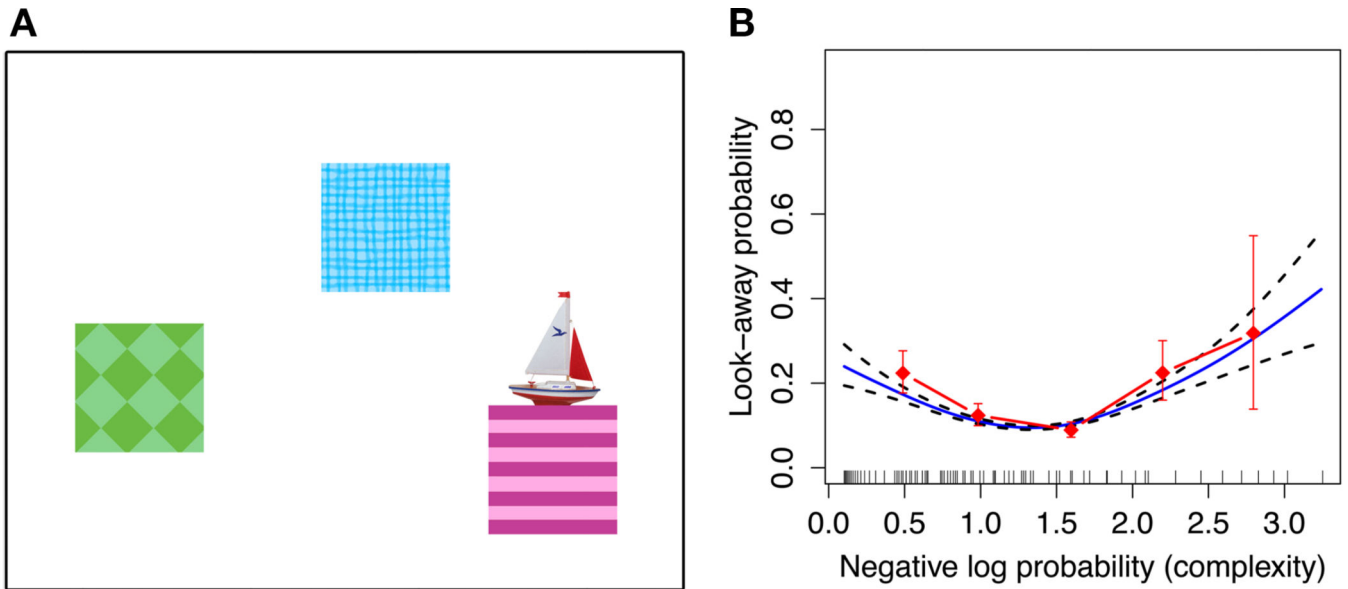


Figure 2.

A. In a four-arm restless bandit task, subjects choose on each trial from one of four targets. B. The value associated with each option changes in value (uncued) stochastically on each trial. Consequently, when the subject has identified the best target, there is a benefit to occasionally interspersing trials where an alternative is chosen (exploration) into the more common pattern of choosing the known best option (exploitation). For example, the subject may choose option A (red color) for several trials but would not know that blue (B) will soon overtake A in value without occasionally exploring other options. C. In this task, neurons in posterior cingulate cortex show higher tonic firing on explore trials than on exploit trials.




Figure 3.

A. In the curiosity tradeoff task, subjects choose between two gambles that vary in informativeness (cyan vs. magenta) and gamble stakes (the size of the white inset bar). On each trial, two gambles appear in sequence on a computer screen (indicated by a black rectangle); when both options appear, subjects shift gaze to one to select it. Then, following 2.25 seconds, they receive a juice reward. Following choice of the informative option, they receive a cue telling them whether they get the reward (50% chance); following choice of the uninformative option, subjects get not valid information. B. Two subjects both showed a preference for informative options (indicated by a left shift of the psychometric curve) over uninformative ones, despite the fact that this information provided no strategic benefit. C. In this task, OFC does not integrate value due to information (vertical axis) with value due to reward size (horizontal axis).

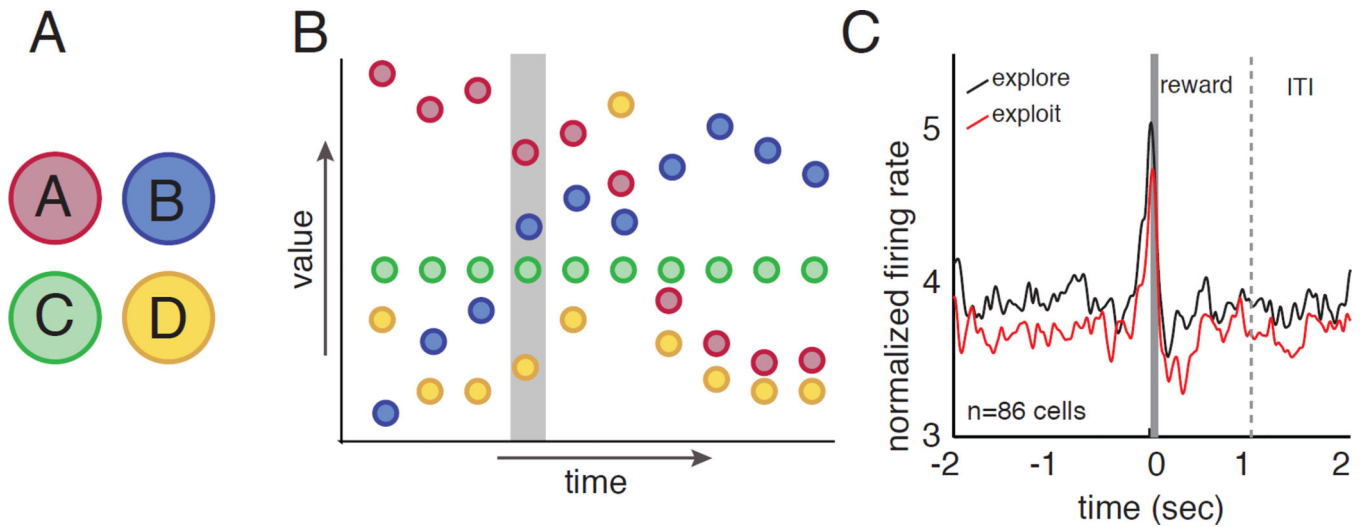
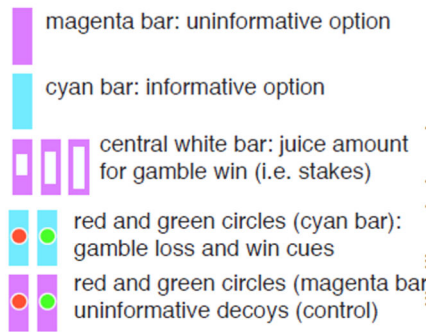


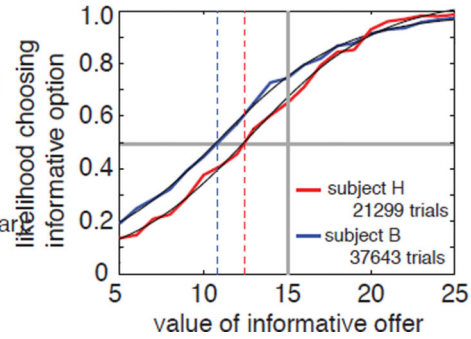
Figure 4.

A: Example display from Kidd, Piantadosi, & Aslin (2012). Each display featured 3 unique boxes hiding 3 unique objects that revealed themselves one at a time according to one of 32 sequences of varying complexity. The sequence continued until the infant looked away for 1-second. B: Infant look-away data plotted by complexity (information content) as estimated by an ideal observer model over the transitional probabilities. The U-shaped pattern indicates that infants were least likely to look away at events with intermediate information content. Infants probability of looking away was greatest to events of either very low information content (highly predictable) or very high information content (highly surprising), consistent with an attentional strategy that aims to maintain intermediate rates of information absorption.


A



B



C

**Figure 5.**

Experimental stimuli from Schulz & Bonawitz (2007). When both levers were pressed simultaneously, two puppets (a straw pompom and a chick) emerged from the center of the box. In this confounded case, the evidence was not informative about which of the two levers caused each puppet to rise. In the unconfounded conditions, one lever was pressed at a time, making it clear which lever caused each puppet to rise. During a free-play period following the toy's demonstration, children played more with the toy when the demonstrated evidence was confounded.