

The Psychology and Neuroscience of Curiosity

Celeste Kidd¹ and Benjamin Y. Hayden^{1,*}

¹Department of Brain and Cognitive Sciences and Center for Visual Science, University of Rochester, Rochester, NY 14620, USA

*Correspondence: benhayden@gmail.com

<http://dx.doi.org/10.1016/j.neuron.2015.09.010>

Curiosity is a basic element of our cognition, but 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.

Background

Curiosity is such a basic component of our nature 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 television, 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 (ADHD). Curiosity is thought of as the noblest of human drives but is commonly denigrated as dangerous (as in the expression “curiosity killed the cat”). Despite its link with the most abstract human thoughts, some rudimentary forms of it can be observed even in the humble worm *C. elegans*.

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 et al., 2014; Kang et al., 2009). This Perspective aims to summarize this recent research, motivate new interest in the problem, and 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 curiosity and infor-

mation-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 et al., 2014; Jirout and Klahr, 2012; Kang et al., 2009; Sutton and Barto, 1998; Tolman and 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 toward helping us think about how to study curiosity in the future.

Classic Descriptions of Curiosity

Philosopher and psychologist William James called curiosity “the impulse towards better cognition,” meaning that it is the desire to understand that which you do not (James, 1899). He noted that, in children, it drives them toward 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 toward more complete scientific and philosophic knowledge. Psychologist-educators G. Stanley Hall and Theodate L. Smith 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 (Hall and Smith, 1903). 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 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 et al., 1950, 1956; Harlow and McCleary, 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 in 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, therefore distinguishing the curiosity of humans from that of other species (Berlyne, 1966).

The second dimension of curiosity that Berlyne described is informational specificity. Specific curiosity referred to desire for a particular piece of information, whereas 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 et al., 1950; Harlow et al., 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 et al., 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 and Kaplan, 2007). By this view, curiosity is strictly an intrinsic drive, whereas 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 to make a more profitable choice. In other words, contexts in which agents seek information for immediate strategic reasons are not considered curiosity in the strict sense. Although 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 its 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.

Therefore, 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, 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 described curiosity as “a cognitive induced deprivation that arises from the perception of a gap in knowledge and understanding” (Loewenstein, 1994). 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 that 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 et al. (2009) (Figure 1B). They found that curiosity about the




Figure 1. Curiosity and Information-Seeking
(A) Data from Kinney and 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.

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 when they were extremely confident. They were most curious when they had some idea about the answer but lacked confidence. Under these circumstances, the compulsion to know the answer was so great that they were even willing to pay for the information even though curiosity could have been be satiated for free after the session. (The neural findings of this study are discussed below.)

Kang et al. (2009) 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 and Beall, 1984; Jirout and Klahr, 2012; Pelz et al., 2015; Penney and 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 and Markant, 2012, for a review). The idea is that allowing a learner to indulge his curiosity allows him to focus his effort on useful information that he does not yet possess. Furthermore, there is a growing body of evidence suggesting that curiosity enables even infant learners to play an active role in optimizing their learning experiences (P. Oudeyer, personal communication). This work suggests that allowing a learner to expose the information he requires himself—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 searches, 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 the intake of information. For example, we choose our visual fixa-

tions strategically to learn about the things that are important to us in the context (Yarbus 1956; Gottlieb, 2012;

Gottlieb et al., 2013, 2014). Given its important role, it is not surprising that our visual search is highly efficient. It is nearly optimal compared with an “ideal searcher” that uses precise statistics of the visual scene to maximize search efficiency (Najemnik and 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 reward. Although 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 min), then abruptly adjusts strategies and makes large, directed movements in a new direction (Calhoun et al., 2014). This search strategy is more sophisticated and beneficial than simply moving toward 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, although 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 (Vergassola et al., 2007). Information gained from such organisms can help us 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 and Jones, 1974; Figure 2). In this task, decision-makers are faced with a series of choices between stochastic reward (Whittle, 1988).




Figure 2. Bandit Task

(A) In a four-arm restless bandit task, subjects choose one of four targets in each trial.

(B) The value associated with each option changes stochastically in value (uncued) in 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 the posterior cingulate cortex show higher tonic firing in explore trials than in exploit trials. ITI, inter-trial interval.

The optimal strategy requires adjudication between exploration (sampling to improve knowledge and, therefore, future choices) and exploitation (choosing known best options). Sampling typically gives a 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 et al. (2006) have shown that humans performing a four-arm bandit task choose options probabilistically based on expected values of the options (a “softmax” policy). 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 the striatum and ventromedial prefrontal cortex (vmPFC) are more active during exploitative choices (Daw et al., 2006). These are canonical reward areas; therefore, these results link curiosity to the reward system (a theme to which we will return). 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 in exploratory trials than in exploitative trials (even after controlling for reward expectation; Pearson et al., 2009; Figure 2). Firing rates also predict adjustments from an exploitative to exploratory strategy and vice versa. These results highlight the contribution of the PCC, a critical but mostly mysterious hub of the reward system, in both the transition to exploration and in its maintenance (Pearson et al., 2011). The PCC is linked to both reward and regulation of learning, therefore underscoring the possible link between these processes and curiosity (Heilbronner and 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., 2011). 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 et al., 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 in 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 toward choosing novel images was mathematically expressible as a novelty bonus (Gittins and 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 the 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, then it ought to be discounted. Therefore, 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




Figure 3. Curiosity Tradeoff Task

(A) In the curiosity tradeoff task, subjects choose between two gambles that vary in informativeness (cyan versus magenta) and gamble stakes (the size of the white inset bar). In each trial, two gambles appear in sequence on a computer screen (indicated by a black rectangle). When both options appear, subjects shift their gaze to one to select it. Then, after 2.25 s, 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 because of information (vertical axis) with value because of reward size (horizontal axis).

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 (Krebs and 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 s) (Bromberg-Martin and Hikosaka, 2009). The only difference between the two gambles is that one offers immediate information about win versus loss (that is, immediate temporal resolution of uncertainty), whereas, 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 and Hikosaka, 2009, 2011; Blanchard et al., 2015). This preference for earlier temporal resolution of uncertainty is not strategic because the information cannot improve choices. Therefore, these tasks satisfy a stricter notion of curiosity.

We modified this task to quantify the value of information by titrating the values of the reward (Blanchard et al., 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 two pellets over one that provides an average of three pellets as long as the one that provides two pellets 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 Stagner (2011) did make the link between their risk-seeking pigeons and human gamblers. This link is potentially important. Curiosity is often mooted as an explanation for risk-seeking behavior (Bromberg-Martin and Hikosaka, 2009). Rhesus monkeys, for example, are often risk-seeking in laboratory tasks (Blanchard and Hayden, 2014; Heilbronner and Hayden 2013; Monosov and Hikosaka, 2012; O'Neill and Schultz, 2010; So and Stuphorn, 2012; Strait et al., 2014, 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 et al., 2009). We found that, following choices of the low variance (and, therefore, non-preferred) option, when it was too late to change anything, monkeys will spontaneously shift gaze to its position, suggesting that 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, therefore, 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 et al. (2009) 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 reward, 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 and 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 the 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 et al. (2009), 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 (Jepma et al., 2012), whereas, in the Kang study it is pleasurable (Kang et al., 2009). Specifically, curiosity is seen as a lack of something wanted (information) and, therefore, unpleasant, and this unpleasantness motivates information, which will alleviate it.

Gruber et al. (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 that 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. Therefore, the curiosity state led to better learning, even for the things people were not curious about. Curiosity drove activity in both the 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 reward.

Single-unit recordings from the temporal resolution of an 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 and Gurney, 2006; Bromberg-Martin et al., 2010; Schultz and Dickinson, 2000). The phasic dopamine response is thought to serve as a general reward prediction error—indicating reward 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 a more indirect kind of reward. The fact that dopamine neurons signal both the 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 a reward—such as

surprising and aversive events—may reflect the value that information provides (Horvitz, 2000; Matsumoto and Hikosaka, 2009; Redgrave and 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 errors (Bromberg-Martin and Hikosaka, 2011). Although 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 a 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 a basic reward (water and juice).

Where does the domain-general curiosity signal come from? It has recently been proposed that the dopamine reward signal is constructed of input signals originating in the orbitofrontal cortex (OFC), which, in turn, receives input from basic sensory and association structures (Ongür and 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 et al., 2015), but it does not integrate them into a single value signal. Therefore, 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 the 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 and Earl, 1957; Kinney and Kagan, 1976; Sokolov, 1963). The world of infants is full of potential sources for learning, but they possess limited information-processing resources. Therefore, 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 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 toward certain informative features of the world. Haith (1980) argued that these organizing principles for visual behavior are fundamentally stimulus-driven. For example, an infant's gaze is pulled toward areas of high contrast, which is useful for detecting objects and perceiving their shapes (e.g., Salapatek and Kessen, 1966), and motion onset, which is useful for detecting animacy (e.g., Aslin and Shea, 1990). Infants




Figure 4. Curiosity and Attention

(A) Example display from Kidd et al. (2012). Each display featured three unique boxes hiding three 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 s.

(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. The 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.

also have an innate bias to orient toward 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). Although 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, they do by a broader definition, which sees curiosity as any mechanism that guides an organism toward new information regardless of mental substrate. 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 must also 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, 1963). Infants prefer novel stimuli in many paradigms, such as those used by Fantz (1964), the high-amplitude sucking procedure (Siqueland and DeLucia, 1969), and the 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 and 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 and Earl, 1957; Hunter and Ames, 1988; Kinney and Kagan, 1976; Roder et al., 2000; Rose et al., 1982; Wagner and Sakovits, 1986). This idea recalls the fact that we are curious about things of which we are moderately certain (Kang et al., 2009).

Among these theories was Kinney and Kagan's moderate discrepancy hypothesis, which suggested that infants preferen-

tially attend to stimuli that were "optimally discrepant," meaning those that were just the right amount of distinguishable from mental representations the infant already possessed (Kinney and Kagan, 1976). Under Dember and Earl'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 and Earl, 1957). Similarly, Berlyne (1960) noted that complexity-driven preferences could represent an optimal strategy for learning. 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 et al., 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 the 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 et al., 2012), for auditory stimuli (Kidd et al., 2014), and even within individual infants (Piantadosi et al., 2014). These results suggest that infants implicitly decide to direct attention to maintain intermediate rates of information absorption. This attentional strategy likely prevents them from wasting cognitive resources on overly predictable or overly complex events, therefore 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 and Bonawitz, 2007). This work is in line with the earlier theories of Jean Piaget that asserted that the purpose of curiosity and play was to "construct knowledge" through interactions with the world (Piaget, 1945). If curiosity aims to reduce uncertainty in the world, then we would expect learners to exhibit increased curiosity to stimuli in the world that they do not understand. In fact, this is a behavior




Figure 5. Curiosity and Learning

Experimental stimuli from [Schulz and Bonawitz \(2007\)](#). When both levers were pressed simultaneously, two puppets (a straw pom-pom 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. Under 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.

that is well attested in recent developmental psychology studies, such as work by [Bonawitz et al., 2012](#)), which 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 et al. \(2011\)](#) gave children a novel toy to explore, either prefaced or not with partial instructions on how the toy works, children played for longer and discovered more of the toys' functions under 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 the toys' inner workings. In these experiments, [Schulz and Bonawitz \(2007\)](#) had children play with toys consisting of boxes and levers. Under both the confounded and unconfounded conditions, the researcher would help a child play with a red box with two levers. Under 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. Under 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. Therefore, 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 s. 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 et al., 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 the bead pair was permanently stuck together, children held it so that only one side was touching the box 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 to deconfound variables when causal mechanisms at play in the world are unclear (e.g., [Denison et al., 2013](#); [Gopnik et al., 1999](#); [Gweon et al., 2014](#); [Schulz et al., 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 and 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. Therefore, these strategic information-seeking behaviors in young children are far more sophisticated than the simple attentional heuristics that characterize early infant attention.

Conclusions

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. Therefore, we recommend that the definition stage should 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, a greater range of manipulations, and 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 the 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, therefore biasing decision-makers toward 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? Which factors unite distinct forms of curiosity?
- How is curiosity controlled? Which 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 Alzheimer's 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, perhaps, in 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 grant from the NIH, R01 (DA038615) (to B.Y.H.). We thank Sarah Heilbronner, Steve Piantadosi, Shraddha Shah, Maya Wang, Habiba Azab, and Maddie Pelz for helpful comments.

REFERENCES

- Aslin, R.N., and Shea, S.L. (1990). Velocity thresholds in human infants: implications for the perception of motion. *Dev. Psychol.* 26, 589–598.
- Averbeck, B.B. (2015). Theory of choice in bandit, information sampling and foraging tasks. *PLoS Comput. Biol.* 11, e1004164.
- Baldwin, D.A. (1993). Infants' ability to consult the speaker for clues to word reference. *J. Child Lang.* 20, 395–418.
- Berlyne, D.E. (1954). A theory of human curiosity. *Br. J. Psychol.* 45, 180–191.
- Berlyne, D.E. (1960). *Conflict, Arousal, and Curiosity* (McGraw-Hill).
- Berlyne, D.E. (1966). Curiosity and exploration. *Science* 153, 25–33.
- Berlyne, D.E. (1978). Curiosity and learning. *Motiv. Emot.* 2, 97–175.
- Blanchard, T.C., and Hayden, B.Y. (2014). Neurons in dorsal anterior cingulate cortex signal postdecisional variables in a foraging task. *J. Neurosci.* 34, 646–655.
- Blanchard, T.C., Hayden, B.Y., and Bromberg-Martin, E.S. (2015). Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron* 85, 602–614.
- Bonawitz, E.B., Shafto, P., Gweon, H., Goodman, N., Spelke, E., and Schulz, L.E. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition* 120, 322–330.
- Bonawitz, E.B., van Schijndel, T., Friel, D., and Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognit. Psychol.* 64, 215–234.
- Bromberg-Martin, E.S., and Hikosaka, O. (2009). Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron* 63, 119–126.
- Bromberg-Martin, E.S., and Hikosaka, O. (2011). Lateral habenula neurons signal errors in the prediction of reward information. *Nat. Neurosci.* 14, 1209–1216.
- Bromberg-Martin, E.S., Matsumoto, M., and Hikosaka, O. (2010). Dopamine in motivational control: rewarding, aversive, and alerting. *Neuron* 68, 815–834.
- Bushnell, E.W., Sai, F., and Mullin, J.T. (1989). Neonatal recognition of the mother's face. *Br. J. Dev. Psychol.* 7, 3–15.
- Butler, R.A. (1953). Discrimination learning by rhesus monkeys to visual-exploration motivation. *J. Comp. Physiol. Psychol.* 46, 95–98.
- Calhoun, A.J., Chalasani, S.H., and Sharpee, T.O. (2014). Maximally informative foraging by *Caenorhabditis elegans*. *eLife* 3, e04220.
- Cook, C., Goodman, N.D., and Schulz, L.E. (2011). Where science starts: spontaneous experiments in preschoolers' exploratory play. *Cognition* 120, 341–349.
- Davis, R.T., Settlage, P.H., and Harlow, H.F. (1950). Performance of normal and brain-operated monkeys on mechanical puzzles with and without food incentive. *J. Genet. Psychol.* 77, 305–311.
- Daw, N.D., O'Doherty, J.P., Dayan, P., Seymour, B., and Dolan, R.J. (2006). Cortical substrates for exploratory decisions in humans. *Nature* 441, 876–879.
- Day, H.I. (1971). The measurement of specific curiosity. In *Intrinsic Motivation: A New Direction in Education*, H.I. Day, D.E. Berlyne, and D.E. Hunt, eds. (Holt, Rinehart & Winston).
- de Quervain, D.J.F., Fischbacher, U., Treyer, V., Schellhammer, M., Schnyder, U., Buck, A., and Fehr, E. (2004). The neural basis of altruistic punishment. *Science* 305, 1254–1258.
- DeCasper, A.J., and Spence, M.J. (1986). Prenatal maternal speech influences newborns' perception of speech sounds. *Infant Behav. Dev.* 9, 133–150.
- Deci, E.L. (1975). *Intrinsic Motivation* (Plenum).
- Delgado, M.R., Nystrom, L.E., Fissell, C., Noll, D.C., and Fiez, J.A. (2000). Tracking the hemodynamic responses to reward and punishment in the striatum. *J. Neurophysiol.* 84, 3072–3077.

- Delgado, M.R., Locke, H.M., Stenger, V.A., and Fiez, J.A. (2003). Dorsal striatum responses to reward and punishment: effects of valence and magnitude manipulations. *Cogn. Affect. Behav. Neurosci.* 3, 27–38.
- Delgado, M.R., Schotter, A., Ozbay, E.Y., and Phelps, E.A. (2008). Understanding overbidding: using the neural circuitry of reward to design economic auctions. *Science* 321, 1849–1852.
- Dember, W.N. (1956). Response by the rat to environmental change. *J. Comp. Physiol. Psychol.* 49, 93.
- Dember, W.N., and Earl, R.W. (1957). Analysis of exploratory, manipulatory, and curiosity behaviors. *Psychol. Rev.* 64, 91–96.
- Denison, S., Bonawitz, E., Gopnik, A., and Griffiths, T.L. (2013). Rational variability in children's causal inferences: the Sampling Hypothesis. *Cognition* 126, 285–300.
- Dyer, F.C. (1991). Bees acquire route-based memories but not cognitive maps in a familiar landscape. *Animal Behav.* 41, 239–246.
- Engel, S. (2011). Children's need to know: curiosity in schools. *Harv. Educ. Rev.* 81, 625–645.
- Engel, S. (2015). *The Hungry Mind: The Origins of Curiosity in Childhood* (Harvard University Press).
- Fant, R.L. (1964). Visual experience in infants: decreased attention to familiar patterns relative to novel ones. *Science* 146, 668–670.
- Farroni, T., Johnson, M.H., Menon, E., Zülke, L., Faraguna, D., and Csibra, G. (2005). Newborns' preference for face-relevant stimuli: effects of contrast polarity. *Proc. Natl. Acad. Sci. USA* 102, 17245–17250.
- Fehr, E., and Camerer, C.F. (2007). Social neuroeconomics: the neural circuitry of social preferences. *Trends Cogn. Sci.* 11, 419–427.
- Gipson, C.D., Alessandri, J.J., Miller, H.C., and Zentall, T.R. (2009). Preference for 50% reinforcement over 75% reinforcement by pigeons. *Learn. Behav.* 37, 289–298.
- Gittins, J., and Jones, D. (1974). *Progress in Statistics* (North-Holland).
- Gopnik, A. and Schulz, L., eds. (2007). *Causal Learning: Psychology, Philosophy, and Computation* (Oxford University Press).
- Gopnik, A., Meltzoff, A.N., and Kuhl, P.K. (1999). *The Scientist in the Crib: Minds, Brains, & How Children Learn* (William Morrow & Co.).
- Gopnik, A., Sobel, D.M., Schulz, L.E., and Glymour, C. (2001). Causal learning mechanisms in very young children: two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation. *Dev. Psychol.* 37, 620–629.
- Gottlieb, J. (2012). Attention, learning, and the value of information. *Neuron* 76, 281–295.
- Gottlieb, J., Oudeyer, P.-Y., Lopes, M., and Baranes, A. (2013). Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends Cogn. Sci.* 17, 585–593.
- Gottlieb, J., Hayhoe, M., Hikosaka, O., and Rangel, A. (2014). Attention, reward, and information seeking. *J. Neurosci.* 34, 15497–15504.
- Gould, J.L. (1986). The locale map of honey bees: do insects have cognitive maps? *Science* 232, 861–863.
- Gray, P. (2013). *Free to Learn: Why Unleashing the Instinct to Play Will Make Our Children Happier, More Self-Reliant, and Better Students for Life* (Basic Books).
- Gruber, M.J., Gelman, B.D., and Ranganath, C. (2014). States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron* 84, 486–496.
- Gureckis, T.M., and Markant, D.B. (2012). Self-Directed Learning: A Cognitive and Computational Perspective. *Perspect. Psychol. Sci.* 7, 464–481.
- Gweon, H., Pelton, H., Konopka, J.A., and Schulz, L.E. (2014). Sins of omission: children selectively explore when teachers are under-informative. *Cognition* 132, 335–341.
- Haith, M.M. (1980). *Rules That Babies Look By* (L. Erlbaum Associates).
- Hall, G.S., and Smith, T.L. (1903). Curiosity and interest. *Pedagogical Seminary* 10, 315–358.
- Harlow, H.F. (1950). Learning and satiation of response in intrinsically motivated complex puzzle performance by monkeys. *J. Comp. Physiol. Psychol.* 43, 289–294.
- Harlow, H.F., and McClearn, G.E. (1954). Object discrimination learned by monkeys on the basis of manipulation motives. *J. Comp. Physiol. Psychol.* 47, 73.
- Harlow, H.F., Blazek, N.C., and McClearn, G.E. (1956). Manipulatory motivation in the infant rhesus monkey. *J. Comp. Physiol. Psychol.* 49, 444.
- Harlow, H.F., Harlow, M.K., and Meyer, D.R. (1950). Learning motivated by a manipulation drive. *J. Exp. Psychol.* 40, 228–234.
- Harty, H., and Beall, D. (1984). Toward the development of a children's science curiosity measure. *J. Res. Sci. Teach.* 21, 425–436.
- Hayden, B.Y., Nair, A.C., McCoy, A.N., and Platt, M.L. (2008). Posterior cingulate cortex mediates outcome-contingent allocation of behavior. *Neuron* 60, 19–25.
- Hayden, B.Y., Pearson, J.M., and Platt, M.L. (2009). Fictive reward signals in the anterior cingulate cortex. *Science* 324, 948–950.
- Heilbrunner, S.R., and Hayden, B.Y. (2013). Contextual factors explain risk-seeking preferences in rhesus monkeys. *Front. Neurosci.* 7, 7.
- Heilbrunner, S.R., and Platt, M.L. (2013). Causal evidence of performance monitoring by neurons in posterior cingulate cortex during learning. *Neuron* 80, 1384–1391.
- Heilbrunner, S.R., Hayden, B.Y., and Platt, M.L. (2011). Decision salience signals in posterior cingulate cortex. *Front. Neurosci.* 5, 55.
- Horvitz, J.C. (2000). Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neuroscience* 96, 651–656.
- Hughes, R.N. (1968). Behaviour of male and female rats with free choice of two environments differing in novelty. *Anim. Behav.* 16, 92–96.
- Hunter, M.A., and Ames, E.W. (1988). A multifactor model of infant preferences for novel and familiar stimuli. *Adv. Infancy Res.* 5, 69–95.
- James, W. (1899). *Talks to Teachers on Psychology: And to Students on Some of Life's Ideals* (Henry Holt & Company).
- James, W. (1913). *The Principles of Psychology, Volume II* (Henry Holt & Company).
- Jepma, M., Verdonschot, R.G., van Steenbergen, H., Rombouts, S.A.R.B., and Nieuwenhuis, S. (2012). Neural mechanisms underlying the induction and relief of perceptual curiosity. *Front. Behav. Neurosci.* 6, 5.
- Jirout, J., and Klahr, D. (2012). Children's scientific curiosity: in search of an operational definition of an elusive concept. *Dev. Rev.* 32, 125–160.
- Johnson, M.H., Dziurawiec, S., Ellis, H., and Morton, J. (1991). Newborns' preferential tracking of face-like stimuli and its subsequent decline. *Cognition* 40, 1–19.
- Kang, M.J., Hsu, M., Krajbich, I.M., Loewenstein, G., McClure, S.M., Wang, J.T.Y., and Camerer, C.F. (2009). The wick in the candle of learning: epistemic curiosity activates reward circuitry and enhances memory. *Psychol. Sci.* 20, 963–973.
- Kemler Nelson, D.G., Jusczyk, P.W., Mandel, D.R., Myers, J., Turk, A., and Gerken, L. (1995). The head-turn preference procedure for testing auditory perception. *Infant Behav. Dev.* 18, 111–116.
- Kidd, C., Piantadosi, S.T., and 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, e36399.
- Kidd, C., Piantadosi, S.T., and Aslin, R.N. (2014). The Goldilocks effect in infant auditory attention. *Child Dev.* 85, 1795–1804.

- King-Casas, B., Tomlin, D., Anen, C., Camerer, C.F., Quartz, S.R., and Montague, P.R. (2005). Getting to know you: reputation and trust in a two-person economic exchange. *Science* 308, 78–83.
- Kinney, D.K., and Kagan, J. (1976). Infant attention to auditory discrepancy. *Child Dev.* 47, 155–164.
- Kivy, P.N., Earl, R.W., and Walker, E.L. (1956). Stimulus context and satiation. *J. Comp. Physiol. Psychol.* 49, 90.
- Knutson, B., Adams, C.M., Fong, G.W., and Hommer, D. (2001). Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *J. Neurosci.* 21, RC159.
- Kreps, D.M., and Porteus, E.L. (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46, 185–200.
- Lieberman, D.A., Cathro, J.S., Nichol, K., and Watson, E. (1997). The role of S- in human observing behavior: bad news is sometimes better than no news. *Learn. Motiv.* 28, 20–42.
- Loewenstein, G. (1994). The psychology of curiosity: a review and reinterpretation. *Psychol. Bull.* 116, 75–98.
- Luhmann, C.C., Chun, M.M., Yi, D.J., Lee, D., and Wang, X.J. (2008). Neural dissociation of delay and uncertainty in intertemporal choice. *J. Neurosci.* 28, 14459–14466.
- Matsumoto, M., and Hikosaka, O. (2009). Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature* 459, 837–841.
- Monosov, I.E., and Hikosaka, O. (2012). Regionally distinct processing of rewards and punishments by the primate ventromedial prefrontal cortex. *J. Neurosci.* 32, 10318–10330.
- Najemnik, J., and Geisler, W.S. (2005). Optimal eye movement strategies in visual search. *Nature* 434, 387–391.
- O’Keefe, J., and Nadel, L. (1978). *The Hippocampus as a Cognitive Map* (Oxford, UK: Clarendon Press).
- O’Neill, M., and Schultz, W. (2010). Coding of reward risk by orbitofrontal neurons is mostly distinct from coding of reward value. *Neuron* 68, 789–800.
- Ongür, D., and Price, J.L. (2000). The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cereb. Cortex* 10, 206–219.
- Oudeyer, P.-Y., and Kaplan, F. (2007). What is intrinsic motivation? A typology of computational approaches. *Front. Neurobot.* 1, 6.
- Pavlov, I.P. (1927). *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex* (Oxford University Press).
- Pearson, J.M., Hayden, B.Y., Raghavachari, S., and Platt, M.L. (2009). Neurons in posterior cingulate cortex signal exploratory decisions in a dynamic multioption choice task. *Curr. Biol.* 19, 1532–1537.
- Pearson, J.M., Heilbronner, S.R., Barack, D.L., Hayden, B.Y., and Platt, M.L. (2011). Posterior cingulate cortex: adapting behavior to a changing world. *Trends Cogn. Sci.* 15, 143–151.
- Pelz, M., Yung, A., and Kidd, C. (2015). Quantifying curiosity and exploratory play on touchscreen tablets. G. Gordon, J. Jirout, S. Engel, and A. Chang, eds. *Proceedings of the IDC 2015 Workshop on Digital Assessment and Promotion of Children’s Curiosity*, <http://www.bcs.rochester.edu/people/mpelz/PelzYungKiddIDC2015.pdf>.
- Penney, R.K., and McCann, B. (1964). The children’s reactive curiosity scale. *Psychol. Rep.* 15, 323–334.
- Phillips, K.A., Subiaul, F., and Sherwood, C.C. (2012). Curious monkeys have increased gray matter density in the precuneus. *Neurosci. Lett.* 518, 172–175.
- Piaget, J. (1945). *Play, Dreams, and Imitation in Childhood* (Heinemann).
- Piantadosi, S.T., Kidd, C., and Aslin, R. (2014). Rich analysis and rational models: inferring individual behavior from infant looking data. *Dev. Sci.* 17, 321–337.
- Prokasy, W.F., Jr. (1956). The acquisition of observing responses in the absence of differential external reinforcement. *J. Comp. Physiol. Psychol.* 49, 131–134.
- Redgrave, P., and Gurney, K. (2006). The short-latency dopamine signal: a role in discovering novel actions? *Nat. Rev. Neurosci.* 7, 967–975.
- Rilling, J., Gutman, D., Zeh, T., Pagnoni, G., Berns, G., and Kilts, C. (2002). A neural basis for social cooperation. *Neuron* 35, 395–405.
- Roder, B.J., Bushnell, E.W., and Sasseville, A.M. (2000). Infants’ preferences for familiarity and novelty during the course of visual processing. *Infancy* 1, 491–507.
- Rose, S.A., Gottfried, A.W., Melloy-Carminar, P., and Bridger, W.H. (1982). Familiarity and novelty preferences in recognition memory: implications for information processing. *Dev. Psychol.* 18, 704–713.
- Rushworth, M.F.S., Noonan, M.P., Boorman, E.D., Walton, M.E., and Behrens, T.E. (2011). Frontal cortex and reward-guided learning and decision-making. *Neuron* 70, 1054–1069.
- Salapatek, P., and Kessen, W. (1966). Visual scanning of triangles by the human newborn. *J. Exp. Child Psychol.* 3, 155–167.
- Schoenbaum, G., Takahashi, Y., Liu, T.L., and McDannald, M.A. (2011). Does the orbitofrontal cortex signal value? *Ann. N Y Acad. Sci.* 1239, 87–99.
- Schultz, W., and Dickinson, A. (2000). Neuronal coding of prediction errors. *Annu. Rev. Neurosci.* 23, 473–500.
- Schultz, W., Dayan, P., and Montague, P.R. (1997). A neural substrate of prediction and reward. *Science* 275, 1593–1599.
- Schulz, L., and Bonawitz, E.B. (2007). Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Dev. Psychol.* 43, 1045–1050.
- Schulz, L.E., Gopnik, A., and Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Dev. Sci.* 10, 322–332.
- Siqueland, E.R., and DeLucia, C.A. (1969). Visual reinforcement of nonnutritive sucking in human infants. *Science* 165, 1144–1146.
- Skinner, B.F. (1938). *The Behavior of Organisms: An Experimental Analysis* (MIT Press).
- So, N., and Stuphorn, V. (2012). Supplementary eye field encodes reward prediction error. *J. Neurosci.* 32, 2950–2963.
- Sokolov, E. (1963). *Perception and the Conditioned Reflex* (Pergamon).
- Stagner, J.P., and Zentall, T.R. (2010). Suboptimal choice behavior by pigeons. *Psychon. Bull. Rev.* 17, 412–416.
- Strait, C.E., Blanchard, T.C., and Hayden, B.Y. (2014). Reward value comparison via mutual inhibition in ventromedial prefrontal cortex. *Neuron* 82, 1357–1366.
- Strait, C.E., Sleezer, B.J., and Hayden, B.Y. (2015). Signatures of value comparison in ventral striatum neurons. *PLoS Biol.* 13, e1002173.
- Sutton, R.S., and Barto, A.G. (1998). *Introduction to Reinforcement Learning* (MIT Press).
- Takahashi, Y.K., Roesch, M.R., Wilson, R.C., Toreson, K., O’Donnell, P., Niv, Y., and Schoenbaum, G. (2011). Expectancy-related changes in firing of dopamine neurons depend on orbitofrontal cortex. *Nat. Neurosci.* 14, 1590–1597.
- Tinbergen, N. (1963). On aims and methods of ethology. *Z. Tierpsychol.* 20, 410–433.
- Tolman, E.C. (1948). Cognitive maps in rats and men. *Psychol. Rev.* 55, 189–208.
- Tolman, E.C., and Gleitman, H. (1949). Studies in learning and motivation; equal reinforcements in both end-boxes; followed by shock in one end-box. *J. Exp. Psychol.* 39, 810–819.
- van Schijndel, T.J.P., Visser, I., van Bers, B.M.C.W., and Raijmakers, M.E.J. (2015). Preschoolers perform more informative experiments after observing theory-violating evidence. *J. Exp. Child Psychol.* 131, 104–119.

- Vergassola, M., Villermaux, E., and Shraiman, B.I. (2007). 'Infotaxis' as a strategy for searching without gradients. *Nature* 445, 406–409.
- Wagner, S.H., and Sakovits, L.J. (1986). A process analysis of infant visual and crossmodal recognition memory: implications for an amodal code. *Advances in Infancy Research* 4, 195–217.
- Wehner, R., Gallizzi, K., Frei, C., and Vesely, M. (2002). Calibration processes in desert ant navigation: vector courses and systematic search. *J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol.* 188, 683–693.
- Whittle, P. (1988). Restless bandits: Activity allocation in a changing world. *J. Appl. Probab.* 25, 287–298.
- Wilson, R.C., Geana, A., White, J.M., Ludvig, E.A., and Cohen, J.D. (2014). Humans use directed and random exploration to solve the explore-exploit dilemma. *J. Exp. Psychol. Gen.* 143, 2074–2081.
- Wittmann, B.C., Daw, N.D., Seymour, B., and Dolan, R.J. (2008). Striatal activity underlies novelty-based choice in humans. *Neuron* 58, 967–973.
- Wyckoff, L.B., Jr. (1952). The role of observing responses in discrimination learning. *Psychol. Rev.* 59, 431–442.
- Yarbus, A.L. (1956). The motion of the eye in the process of changing points of fixation. *Biofizika* 1, 76–78.
- Zeil, J. (1998). Homing in fiddler crabs (*Uca lactea annulipes* and *Uca vomeris*: Ocypodidae). *J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol.* 183, 367–377.
- Zentall, T.R., and Stagner, J. (2011). Maladaptive choice behaviour by pigeons: an animal analogue and possible mechanism for gambling (sub-optimal human decision-making behaviour). *Proc. Biol. Sci.* 278, 1203–1208.