Curiosity from the Perspective of Systems Neuroscience

Roberto Lopez Cervera, Maya Zhe Wang, and Benjamin Y. Hayden

Department of Neuroscience,
Center for Magnetic Resonance Research,
and Center for Neuroengineering
University of Minnesota, Minneapolis MN 55455

Contact Information:

Roberto Lopez Cervera
Department of Neuroscience and Center for Magnetic Resonance Research
University of Minnesota, Minneapolis MN 55455
Email: lopez927@umn.edu

ABSTRACT

Curiosity refers to a demand for information that has no instrumental benefit. Because of its critical role in development and in the regulation of learning, curiosity has long fascinated psychologists. However, it has been difficult to study curiosity from the perspective of the single neuron, the circuit, and systems neuroscience. Recent advances, however, have made doing so more feasible. These include theoretical advances in defining curiosity in animal models, the development of tasks that manipulate curiosity, and the preliminary identification of circuits responsible for curiosity-motivated learning. Taken together, resulting scholarship demonstrates the key roles of executive control, reward, and learning circuits in driving curiosity; and has helped us to understand how curiosity relates to information-seeking more broadly. This work has implications for mechanisms of reward-based decisions in general. Here we summarize these results and highlight important remaining questions for the future of curiosity studies.

MAIN TEXT

Introduction

Pitfall is a classic Atari game in which a player must navigate an avatar around a virtual screen to explore and score points [1]. Although a typical seven-year old can do quite well at the game with minimal practice, well-trained deep learning agents that can master many Atari games were unable to score even a single point until 2019 [2]. The reason is that Pitfall is a “hard-exploration” problem - it has several factors that punish type of learning style of typical deep learning agents [3]. Importantly, survival for animals in the real world has many of the same features, meaning that life, in general, is likely also a hard-exploration problem. AI agents can now beat the best humans at Pitfall and the key innovation was endowing them with features that mimic curiosity: an internal reward for simply gaining information [2]. Curiosity, it seems, may be an indispensable element in our cognitive repertoires [4,5,6,7].

That is not to say we now know all we need to about curiosity. There are many ways in which AI curiosity is inferior to the real kind (including problems like detachment and derailment, reviewed in [2]). Moreover, we have yet to learn how to harness curiosity drives to improve our educational systems, with corresponding benefits for human welfare [8,9,10]. And we have yet to harness the study of curiosity for diseases associated with aberrant patterns of curiosity [11]. Nor do we know much about how different taxa differ in how their curiosity is deployed. Clearly, a greater understanding of curiosity is called for.

For these reasons, understanding the neural basis of curiosity has become an important issue in cognitive neuroscience. Indeed, in recent years, curiosity has moved from a peripheral to a central position in psychology and cognitive neuroscience [5,7,12,13,14,15,16]. The study of curiosity has lagged, however, in branches of neuroscience concerned with how neurons perform task-relevant computations (we call this field systems neuroscience). However, because neural activity is the ultimate driver of behavior, such understanding is critical for a full accounting of the mechanisms of curiosity in the field. Here we summarize recent advances.

Defining curiosity

One major barrier to understanding curiosity has long been the difficulty of defining it. Most definitions have tended to be heuristic. We have argued that this is a good thing - that definitions need to wait for sufficient empirical advances because we need to know what it is before we can precisely define it [7]. Nonetheless, there is still value in makeshift definitions. They allow us to do the types of experiments that will allow further work. This is especially important in animal models, which currently provide the best tools for systems neuroscience, but cannot provide reports about internal states and motivations (i.e. how curious they feel). We therefore have developed a rough-and-ready definition [17,18] that is inspired by an integrative reading of the recent literature and especially by the foundational work of Loewenstein [5,13,14,19].

Our proposal is that a research subject demonstrates curiosity if it (1) is willing to pay a real price solely for information, (2) that information is demonstrably non-strategic,

and (3) the subject's demand scales with the amount of information (over some range). This definition has limitations; in particular, it is overly conservative (i.e. it rejects many likely cases of curiosity), and it is based on the assumption that subjects have acquired and are using the correct model of the task/environment in directing their information-seeking behaviors, such that they understood that the information was non-strategic. However, we believe this definition is valuable for the time being.

The observing task

Among tasks that generally satisfy criteria for curiosity, the *observing task* has proven to be especially useful [20,21,22,23,24,25,26]. In this task, the decision-maker is faced with risky options, where resolution (that is, the information about the gamble's outcome) takes place after some delay relative to the choice. The decision-maker then has the opportunity to choose an option that provides earlier resolution but does not affect the reward's probability, size, or timing. Preference for this observing option can be taken as an indicator of curiosity. This task neatly excludes strategic information seeking because the reward itself is dispensed at the same time regardless of choice.

It is important to note that this task is no magic bullet. Indeed, there is residual debate about whether information-seeking behavior in the task reflects curiosity [27,28,29]. Specifically, the apparent preference for advance information could be accounted for by greater levels of attentional engagement on informative trials, leading to increased associative learning for informative cues, and thus promoting their choice even in the absence of curiosity [27]. The general argument is plausible, and differential attention may explain some of the effect; however, it appears insufficient to explain several results [18,30,31]. To give one example, lateral habenula neurons reliably signal information prediction errors for informative as well as uninformative cues, indicating that animals do not forget their reward predictions due to task disengagement even in uninformative trials [32].

In a different RL-inspired critique, Iigaya et al. [28] propose that viewing information provides a boost to the value of anticipating primary reward, which can be thought of as a time-dependent savoring effect that increases the overall value of a primary reward, and does not require curiosity. This proposal is supported by both behavior [28] and neuroimaging data [29]. This model, while likely valid, seems unlikely to fully explain observing behavior; for example it cannot explain willingness to pay for counterfactual information, a hallmark of human curiosity that is also demonstrated in monkeys [18].

Functional neuroanatomy of curiosity

Functional neuroanatomical studies have generally demonstrated the critical importance of regions involved in (1) reward, (2) learning/memory, and (3) control - a finding that is reassuring if not surprising given that curiosity can roughly be defined as the control of self-motivated learning.

Reward areas

Neuroimaging literature supports the involvement of the three "usual suspects" reward regions - dopaminergic midbrain, striatum, and the ventral orbital surface - in

curiosity. In perhaps the first neuroscientific study of human curiosity, Kang and colleagues reported activation in the striatum when subjects were anticipating the answer to a trivia question that they were curious about [33]. Since then, several studies have reported activation in the midbrain and striatum in anticipation of information through a variety of tasks including trivia, lottery, and observing tasks [34,30,35]. These results, along with others, support the idea that curiosity works by commandeering the brain's usual reward regions to drive information-seeking. In other words, curiosity works just like any other motivated process [36,37,38, 39,40].

Electrophysiological results support the involvement of reward areas in curiosity and provide insight into computational processes. Bromberg-Martin and colleagues used the observing task to understand the role of midbrain dopamine (DA) neurons in curiosity [26]. In that study, rhesus macaques had to choose between two risky options that were equivalent except that one option resolved outcome uncertainty early (2.25 seconds). The critical finding of the study was that dopaminergic midbrain neurons that responded to the expectation of primary reward also responded to the expectation of receiving information. This suggested that reward-seeking and information-seeking share a "common currency" or neural code - or, perhaps, that curiosity and non-curiosity motivation converge at the level before the DA system.

In a follow-up study, Bromberg-Martin et al. [32] used a similar task, but included an option that only revealed outcome information 50% of the time. This allowed them to calculate information prediction errors (IPEs) when monkeys chose the semi-informative option and were unexpectedly delivered or denied information (this is analogous to how the brain calculates reward prediction errors, RPEs, from the delivery or denial of primary reward). The critical finding of this study was that a subpopulation of neurons in the lateral habenula (LHb), a subcortical reward structure, signaled IPEs in addition to conventional RPEs. This result further supported the "common currency" hypothesis. Moreover, because of projections from LHb to the midbrain DA system, these results presented the start of a possible hierarchical anatomical pathway along which curiosity representations are transformed into choices.

In a related study performed in our lab, we examined responses in the orbitofrontal cortex (OFC, Area 13) while monkeys performed a version of the task with titrated rewards [41]. This is a brain region associated with valuation processes, but also, to some extent, with executive ones [42,43,44]. This task allowed us to assess the specific value of information for each subject. We found that subjects were willing to pay up to 20-30% of their regular reward intake for advance information. We found that instantaneous firing rates of neurons in this region encode both the value of information (on a continuous scale) and the value of the primary (liquid) rewards, but use orthogonal ensemble coding formats to do so. This fact indicates that OFC tracks multiple influences on choice, but does not integrate them into a coherent abstract value variable, and suggests that such integration into an abstract value variable, if it occurs, takes place in a downstream structure. (Note that a good deal of evidence supports the idea that the midbrain DA system is downstream of OFC [45,46,47]. More broadly, these results fit into theories about hierarchies in reward processing [48,49,50,51].

Our data also bear on the debate about the causes of information-seeking. On one hand, subjects may value information because it allows them to physically or mentally prepare for reward delivery [19,28,29,52]. If this were the case, informative offers would increase primary reward-related signals in OFC. Alternately, the brain could give reward a distinct value assignment - if so, then OFC primary reward signals should have no net enhancement by information. Our data supported the second inherent value tagging view. Specifically, signals coding water reward were not affected by the promise of information about those rewards. Indeed, OFC neurons appeared to use distinct (orthogonal) codes for information about upcoming rewards and unsignaled rewards (which also provided information), suggesting that they encode the value of the information and not the information itself [53].

Though not mutually exclusive, White et al. [31] proposed a cortico-basal ganglia circuit to also explain information-seeking using an observing task that was adapted from Bromberg-Martin and Hikosaka, 2011 [32], but their model showed that ACC and striatum both track uncertainty, and that the delivery of information alone (even before reward delivery) quenches this tracking activity. In addition, striatum was found to causally direct gaze towards informative targets while anterior pallidum was found to inhibit gaze to uninformative targets. These results further delineate the circuitry of curiosity.

Involvement of learning areas in curiosity

Information-seeking implies learning, which implies memory. Indeed, curiosity enhances memory formation for acquired information. It is not surprising, then, that it naturally involves learning [8,33,54,34,55,56]. Most importantly, curiosity activates the hippocampus and the associated parahippocampal gyrus [33, 54, 34]. It was also found that variation in hippocampal involvement predicted subsequent information recall accuracy. Moreover, increased activity while waiting for the answer to a high-curiosity question and after incorrectly answering a high-curiosity question lead to more accurate recall.

There are currently no electrophysiological studies on the role of memory regions in curiosity-motivated learning. Clearly, this is an important job for future research. In particular, it remains critical to record in hippocampal regions and in regions that are likely to anatomically mediate between the hippocampus and the reward/punishment systems, such as OFC and posterior cingulate cortex [57], vmPFC [29], and amygdala [58]

Involvement of control areas in curiosity

Perhaps less obvious than motivation and learning, curiosity involves executive control. That is, it involves carefully managing the tradeoff between competing interests, including the ability to de-emphasize the demand for immediate reward in favor of the most indirect benefits of information. This fact may explain the involvement of classic control areas, the most important of which is the dorsal anterior cingulate cortex (dACC). Several theories suggest that dACC serves to monitor the demand for control or the need to adjust foraging strategy, and send this information to downstream structures that regulate behavior, such as guiding gaze to salient stimuli [59,60,61,62,63,64,65,66,67]. Curiosity about lottery outcomes was shown to disinhibit the ACC in humans via the increased amplitude of a “feedback-related negativity” EEG signal that is believed to be related to

dopaminergic projections to the ACC [68]. ACC was also shown to be responsive to anticipation of information during an observing task in monkeys [31].

In addition to regions that monitor the demand for control, downstream effector structures are also necessary to carry out the command. Gottlieb and colleagues have shown that neurons in the monkey lateral intraparietal cortex (LIP) produce stronger responses to informative task cues (after they appear in the monkey's peripheral vision, but before a saccade is executed) compared to noninformative cues in active sampling tasks [12,69,70,71,72]. A recent study provides human support for these findings, where the authors found that the inferior parietal lobule (IPL, which is believed to be the human homolog of monkey LIP) encoded uncertainty in a lottery task. This finding was based on an increased BOLD response during the presentation of lottery probabilities, which was highest when outcome uncertainty was also highest [73]. These results are important because they demonstrate that gaze can be controlled by a drive to resolve uncertainty independent of reward value and suggest that different circuits may govern different types of executive control [72].

Conclusion

Our brains did not evolve to perform laboratory tasks; they evolved to solve complex foraging problems embedded in a rich natural world [74,75,76,77,78]. In the natural world, information is almost always missing; natural foragers therefore have a constant pressing hunger for information [79]. A clever forager ought to devote a good deal of its resources simply to learning about the world in an effort to reduce uncertainty because even a small amount of information gain can provide a strong competitive advantage. The value of information (or resolving uncertainty) has likely endowed our cognitive repertoires with an intrinsic information-seeking drive, which we call curiosity [7]. This drive is likely reflected in the organization of specific brain circuits and the underlying computations they implement.

So far research implicates three types of systems in curiosity - reward, learning, and control. This work is, naturally, prey to standard problems of reverse inference. Therefore, future studies will be needed to work out the specific processes associated with curiosity. In particular, we need to fully delineate the circuits. Second, we have to understand how curiosity both resembles and differs from other basic drives. Third, we will have to understand the role of valence - how curiosity relates to the price people are willing to pay to avoid information [80,55,30]. Fourth, we need to integrate the study of curiosity into the systems neuroscience of learning and motivation more broadly, especially as how it can teach us about intrinsic motivational variables. This information will ultimately help us harness curiosity to improve education, theories of learning, psychiatry, and our understanding of ourselves and our animal friends.

FIGURES

A

Figure 1. Observing task used in rhesus macaques [41]. A. Schematic of the computerized version of the observing task. Macaque subjects choose on each trial between 50/50 gambles with differing and randomly chosen stakes (indicated by size of inscribed white bar). Informative offers (cyan color) promise gamble resolution immediately after choice; uninformative ones (magenta color) promise task-irrelevant decoy information. Stochastic reward occurs after 2.25 seconds regardless of choice. B. Behavior of two example subjects on the task. Y-axis indicates indifference point (equivalence point) for informative and uninformative options. Both subjects placed value on advance information (as did a third subject, data not shown) and the value of that information increased roughly linearly with the stakes of the gamble.

ACKNOWLEDGMENTS

This work was supported by the National Institutes of Health grant T32 NS105604 (R.L.C.) and by the National Institutes of Health Medical Scientist Training Program grant T32 GM008244 (R.L.C.).

REFERENCES

1. Montfort, N., & Bogost, I. (2009). *Racing the beam: the Atari video computer system*. MIT Press.
2. Ecoffet, A., Huizinga, J., Lehman, J., Stanley, K. O., & Clune, J. (2019). Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*.
3. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533. <https://doi.org/10.1038/nature14236>
4. Urgan Schmidhuber, J. J. (1991). *A Possibility for Implementing Curiosity and Boredom in Model-Building Neural Controllers*. MIT Press/Bradford Books.
5. Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, 116(1), 75–98. <https://doi.org/10.1037/0033-2909.116.1.75>
6. Oudeyer, P. Y., & Kaplan, F. (2009). What is intrinsic motivation? A typology of computational approaches. *Frontiers in neurorobotics*, 1, 6.
7. Kidd, C., & Hayden, B. Y. (2015). The Psychology and Neuroscience of Curiosity. *Neuron*, 88(3), 449–460. <https://doi.org/10.1016/j.neuron.2015.09.010>
8. Wade, S., & Kidd, C. (2019). The role of prior knowledge and curiosity in learning. *Psychonomic Bulletin and Review*, 1377–1387. <https://doi.org/10.3758/s13423-019-01598-6>
9. Engel, S. (2011). Children's need to know: Curiosity in schools. *Harvard Educational Review*, 81(4), 625-645.
10. Gordon, G., Breazeal, C., & Engel, S. (2015, March). Can children catch curiosity from a social robot?. In 2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI) (pp. 91-98). IEEE.
11. Jovanovic, V., & Brdaric, D. (2012). Did curiosity kill the cat? Evidence from subjective well-being in adolescents. *Personality and Individual Differences*, 52(3), 380–384. <https://doi.org/10.1016/j.paid.2011.10.043>
12. Gottlieb, J., Oudeyer, P. Y., Lopes, M., & Baranes, A. (2013). Information-seeking, curiosity, and attention: Computational and neural mechanisms. *Trends in Cognitive Sciences*, 17(11), 585–593. <https://doi.org/10.1016/j.tics.2013.09.001>
13. Golman, R., & Loewenstein, G. (2015). Curiosity, information gaps, and the utility of knowledge. *Information Gaps, and the Utility of Knowledge* (April 16, 2015).
14. Golman, R., & Loewenstein, G. (2018). Information gaps: A theory of preferences regarding the presence and absence of information. *Decision*, 5(3), 143.

15. Gottlieb, J., & Oudeyer, P. Y. (2018). Towards a neuroscience of active sampling and curiosity. *Nature Reviews Neuroscience*, 19(12), 758–770.
<https://doi.org/10.1038/s41583-018-0078-0>
16. Gruber, M. J., & Ranganath, C. (2019). How Curiosity Enhances Hippocampus-Dependent Memory: The Prediction, Appraisal, Curiosity, and Exploration (PACE) Framework. *Trends in Cognitive Sciences*, 1–12.
<https://doi.org/10.1016/j.tics.2019.10.003>
17. Wang, M. Z., Sweis, B. M., & Hayden, B. Y. (2018). A Testable Definition of Curiosity. *IEEE CDS Newsletter*.
18. Wang, M. Z., & Hayden, B. Y. (2019). Monkeys are curious about counterfactual outcomes. *Cognition*, 189(September 2018), 1–10.
<https://doi.org/10.1016/j.cognition.2019.03.009>
19. Loewenstein, G. (1987). Anticipation and the Valuation of Delayed Consumption. *The Economic Journal*, 97(387), 666. <https://doi.org/10.2307/2232929>
20. Wyckoff, L. B. (1952). The role of observing responses in discrimination learning. Part I. *Psychological Review*, 59(6), 431–442. <https://doi.org/10.1037/h0053932>
21. Prokasy, W. F. (1956). The acquisition of observing responses in the absence of differential external reinforcement. *Journal of Comparative and Physiological Psychology*, 49(2), 131–134. <https://doi.org/10.1037/h0046740>
22. Ward, E. F. (1971). Acquisition and extinction of the observing response as a function of stimulus predictive validity. *Psychonomic Science*, 24(3), 139–141.
<https://doi.org/10.3758/BF03331790>
23. Dinsmoor, J. A. (1983). Observing and conditioned reinforcement. *Behavioral and Brain Sciences*, 6(4), 693–704. <https://doi.org/10.1017/S0140525X00017969>
24. Spetch, M. L., Belke, T. W., Barnet, R. C., Dunn, R., & Pierce, W. D. (1990). SUBOPTIMAL CHOICE IN A PERCENTAGE-REINFORCEMENT PROCEDURE: EFFECTS OF SIGNAL CONDITION AND TERMINAL-LINK LENGTH. *Journal of the Experimental Analysis of Behavior*, 53(2), 219–234.
<https://doi.org/10.1901/jeab.1990.53-219>
25. Lieberman, D. A., Cathro, J. S., Nichol, K., & Watson, E. (1997). The role of S-in human observing behavior: Bad news is sometimes better than no news. *Learning and Motivation*, 28(1), 20–42. <https://doi.org/10.1006/lmot.1997.0951>
26. Bromberg-Martin, E. S., & Hikosaka, O. (2009). Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron*, 63(1), 119–126. <https://doi.org/10.1016/j.neuron.2009.06.009>
27. Beierholm, U. R., & Dayan, P. (2010). Pavlovian-instrumental interaction in “observing behavior.” *PLoS Computational Biology*, 6(9).
<https://doi.org/10.1371/journal.pcbi.1000903>
28. Iigaya, K., Story, G. W., Kurth-Nelson, Z., Dolan, R. J., & Dayan, P. (2016). The modulation of savouring by prediction error and its effects on choice. *ELife*, 5(APRIL2016), 1–24. <https://doi.org/10.7554/eLife.13747>
29. Iigaya, K., Hauser, T. U., Kurth-Nelson, Z., O’Doherty, J. P., Dayan, P., & Dolan, R. J. (2019). The value of what’s to come: neural mechanisms coupling prediction error and reward anticipation. *BioRxiv*, 588699. <https://doi.org/10.1101/588699>

30. Charpentier, C. J., Bromberg-Martin, E. S., & Sharot, T. (2018). Valuation of knowledge and ignorance in mesolimbic reward circuitry. *Proceedings of the National Academy of Sciences of the United States of America*, 115(31), E7255–E7264. <https://doi.org/10.1073/pnas.1800547115>
31. White, J. K., Bromberg-Martin, E. S., Heilbronner, S. R., Zhang, K., Pai, J., Haber, S. N., & Monosov, I. E. (2019). A neural network for information seeking. *Nature communications*, 10(1), 1-19.
32. Bromberg-Martin, E. S., & Hikosaka, O. (2011). Lateral habenula neurons signal errors in the prediction of reward information. *Nature Neuroscience*, 14(9), 1209–1218. <https://doi.org/10.1038/nn.2902>
33. Kang, M. J., Hsu, M., Krajchich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T., & Camerer, C. F. (2009). The Wick in the Candle of Learning. *Psychological Science*, 20(8), 963–973. <https://doi.org/10.1111/j.1467-9280.2009.02402.x>
34. Gruber, M. J., Gelman, B. D., & Ranganath, C. (2014). States of Curiosity Modulate Hippocampus-Dependent Learning via the Dopaminergic Circuit. *Neuron*, 84(2), 486–496. <https://doi.org/10.1016/j.neuron.2014.08.060>
35. Kobayashi, K., & Hsu, M. (2019). Common neural code for reward and information value. *Proceedings of the National Academy of Sciences of the United States of America*, 116(26), 13061–13066. <https://doi.org/10.1073/pnas.1820145116>
36. Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. E. (2006). Reward-Motivated Learning: Mesolimbic Activation Precedes Memory Formation. *Neuron*, 50(3), 507–517. <https://doi.org/10.1016/j.neuron.2006.03.036>
37. Fehr, E., & Camerer, C. F. (2007). Social neuroeconomics: the neural circuitry of social preferences. *Trends in cognitive sciences*, 11(10), 419-427.
38. D’Ardenne, K., McClure, S. M., Nystrom, L. E., & Cohen, J. D. (2008). BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science*, 319(5867), 1264–1267. <https://doi.org/10.1126/science.1150605>
39. Hikosaka, O., Bromberg-Martin, E., Hong, S., & Matsumoto, M. (2008). New insights on the subcortical representation of reward. *Current Opinion in Neurobiology*, 18(2), 203–208. <https://doi.org/10.1016/j.conb.2008.07.002>
40. Shohamy, D., & Adcock, R. A. (2010). Dopamine and adaptive memory. *Trends in Cognitive Sciences*, 14(10), 464–472. <https://doi.org/10.1016/j.tics.2010.08.002>
41. Blanchard, T. C., Hayden, B. Y., & Bromberg-Martin, E. S. (2015). Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron*, 85(3), 602–614. <https://doi.org/10.1016/j.neuron.2014.12.050>
42. Wallis, J. D. (2007). Orbitofrontal cortex and its contribution to decision-making. *Annu. Rev. Neurosci.*, 30, 31-56.
43. Izquierdo, A. (2017). Functional heterogeneity within rat orbitofrontal cortex in reward learning and decision making. *Journal of Neuroscience*, 37(44), 10529-10540.

44. Slezzer, B. J., Castagno, M. D., & Hayden, B. Y. (2016). Rule encoding in orbitofrontal cortex and striatum guides selection. *Journal of Neuroscience*, 36(44), 11223–11237.
45. Takahashi, Y. K., Stalnaker, T. A., Roesch, M. R., & Schoenbaum, G. (2017). Effects of inference on dopaminergic prediction errors depend on orbitofrontal processing. <https://doi.org/10.1037/bne0000192>
46. Stalnaker, T. A., Liu, T. L., Takahashi, Y. K., & Schoenbaum, G. (2018). Orbitofrontal neurons signal reward predictions, not reward prediction errors. *Neurobiology of Learning and Memory*, 153(July 2017), 137–143. <https://doi.org/10.1016/j.nlm.2018.01.013>
47. Namboodiri, V. M. K., Otis, J. M., van Heeswijk, K., Voets, E. S., Alghorazi, R. A., Rodriguez-Romaguera, J., ... Stuber, G. D. (2019). Single-cell activity tracking reveals that orbitofrontal neurons acquire and maintain a long-term memory to guide behavioral adaptation. *Nature Neuroscience*, 22(7), 1110–1121. <https://doi.org/10.1038/s41593-019-0408-1>
48. Hunt, L. T., Malalasekera, W. N., de Berker, A. O., Miranda, B., Farmer, S. F., Behrens, T. E., & Kennerley, S. W. (2018). Triple dissociation of attention and decision computations across prefrontal cortex. *Nature neuroscience*, 21(10), 1471.
49. Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., ... & Wang, X. J. (2014). A hierarchy of intrinsic timescales across primate cortex. *Nature neuroscience*, 17(12), 1661.
50. Yoo, S. B. M., & Hayden, B. Y. (2018). Economic Choice as an Untangling of Options into Actions. *Neuron*, 99(3), 434–447.
51. Pezzulo, G., & Cisek, P. (2016). Navigating the Affordance Landscape: Feedback Control as a Process Model of Behavior and Cognition. *Trends in Cognitive Sciences*, 20(6), 414–424. <https://doi.org/10.1016/j.tics.2016.03.013>
52. Perkins Jr, C. C. (1955). The stimulus conditions which follow learned responses. *Psychological Review*, 62(5), 341.
53. Wang, M. Z., & Hayden, B. Y. (2017). Reactivation of associative structure specific outcome responses during prospective evaluation in reward-based choices. *Nature communications*, 8, 15821.
54. Jepma, M., Verdonchot, R. G., van Steenbergen, H., Rombouts, S. A. R. B., & Nieuwenhuis, S. (2012). Neural mechanisms underlying the induction and relief of perceptual curiosity. *Frontiers in Behavioral Neuroscience*, 6. <https://doi.org/10.3389/fnbeh.2012.00005>
55. Marvin, C. B., & Shohamy, D. (2016). Curiosity and reward: Valence predicts choice and information prediction errors enhance learning. *Journal of Experimental Psychology: General*, 145(3), 266–272. <https://doi.org/10.1037/xge0000140>
56. Stare, C. J., Gruber, M. J., Nadel, L., Ranganath, C., & Gómez, R. L. (2018). Curiosity-driven memory enhancement persists over time but does not benefit from post-learning sleep. *Cognitive Neuroscience*, 9(3–4), 100–115. <https://doi.org/10.1080/17588928.2018.1513399>

57. Wikenheiser, A. M., & Schoenbaum, G. (2016). Over the river, through the woods: Cognitive maps in the hippocampus and orbitofrontal cortex. *Nature Reviews Neuroscience*, 17(8), 513–523. <https://doi.org/10.1038/nrn.2016.56>
58. Alvarez, R. P., Biggs, A., Chen, G., Pine, D. S., & Grillon, C. (2008). Contextual fear conditioning in humans: Cortical-hippocampal and amygdala contributions. *Journal of Neuroscience*, 28(24), 6211–6219. <https://doi.org/10.1523/JNEUROSCI.1246-08.2008>
59. Botvinick, M., Nystrom, L. E., Fissell, K., Carter, C. S., & Cohen, J. D. (1999). Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature*, 402(6758), 179.
60. Botvinick, M. M., & Cohen, J. D. (2014). The Computational and Neural Basis of Cognitive Control: Charted Territory and New Frontiers. *Cognitive Science*, 38(6), 1249–1285. <https://doi.org/10.1111/cogs.12126>
61. Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013, July 24). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron*. <https://doi.org/10.1016/j.neuron.2013.07.007>
62. Kolling, N., Behrens, T. E. J., Mars, R. B., & Rushworth, M. F. S. (2012). Neural mechanisms of foraging. *Science*, 335(6077), 95–98. <https://doi.org/10.1126/science.1216930>
63. Kolling, N., Behrens, T. E. J., Wittmann, M. K., & Rushworth, M. F. S. (2016). Multiple signals in anterior cingulate cortex. *Current opinion in neurobiology*, 37, 36-43.
64. Heilbronner, S. R., & Hayden, B. Y. (2016). Dorsal Anterior Cingulate Cortex: A Bottom-Up View. *Annual Review of Neuroscience*, 39(1), 149–170. <https://doi.org/10.1146/annurev-neuro-070815-013952>
65. Smith, E. H., Horga, G., Yates, M. J., Mikell, C. B., Banks, G. P., Pathak, Y. J., Schevon, C. A., McKann, G. A., Hayden, B. Y., Botvinick, M. M., & Sheth, S. A. (2019). Widespread temporal coding of cognitive control in the human prefrontal cortex. *Nature Neuroscience*, 1-9.
66. Eisenreich, B. R., Akaishi, R., & Hayden, B. Y. (2017). Control without Controllers: Toward a Distributed Neuroscience of Executive Control. *Journal of Cognitive Neuroscience*, 29(10), 1684–1698. https://doi.org/10.1162/jocn_a_01139
67. Rushworth, M. F. S., Noonan, M. A. P., Boorman, E. D., Walton, M. E., & Behrens, T. E. (2011, June 23). Frontal Cortex and Reward-Guided Learning and Decision-Making. *Neuron*. <https://doi.org/10.1016/j.neuron.2011.05.014>
68. Brydevall, M., Bennett, D., Murawski, C., & Bode, S. (2018). The neural encoding of information prediction errors during non-instrumental information seeking. *Scientific Reports*. <https://doi.org/10.1038/s41598-018-24566-x>
69. Gottlieb, J., Hayhoe, M., Hikosaka, O., & Rangel, A. (2014). Attention, reward, and information seeking. *Journal of Neuroscience*, 34(46), 15497–15504. <https://doi.org/10.1523/JNEUROSCI.3270-14.2014>
70. Foley, N. C., Kelly, S. P., Mhatre, H., Lopes, M., & Gottlieb, J. (2017). Parietal neurons encode expected gains in instrumental information. *Proceedings of the National Academy of Sciences*, 114(16), E3315-E3323.

- 483 71. Daddaoua, N., Lopes, M., & Gottlieb, J. (2016). Intrinsically motivated oculomotor
484 exploration guided by uncertainty reduction and conditioned reinforcement in non-
485 human primates. *Scientific Reports*, 6. <https://doi.org/10.1038/srep20202>
486 72. Horan, M., Daddaoua, N., & Gottlieb, J. (2019). Parietal neurons encode
487 information sampling based on decision uncertainty. *Nature Neuroscience*.
488 <https://doi.org/10.1038/s41593-019-0440-1>
489 73. Van Lieshout, L. L. F., Vandenbroucke, A. R. E., Müller, N. C. J., Cools, R., & de
490 Lange, F. P. (2018). Induction and relief of curiosity elicit parietal and frontal
491 activity. *Journal of Neuroscience*, 38(10), 2579–2588.
492 <https://doi.org/10.1523/JNEUROSCI.2816-17.2018>
493 74. Pearson, J. M., Watson, K. K., & Platt, M. L. (2014). Decision making: the
494 neuroethological turn. *Neuron*, 82(5), 950-965.
495 75. Calhoun, A. J., & Hayden, B. Y. (2015). The foraging brain. *Current Opinion in*
496 *Behavioral Sciences*, 5, 24-31.
497 76. Krakauer, J. W., Ghazanfar, A. A., Gomez-Marín, A., MacIver, M. A., & Poeppel,
498 D. (2017). Neuroscience needs behavior: correcting a reductionist bias. *Neuron*,
499 93(3), 480-490.
500 77. Brown, A. E., & De Bivort, B. (2018). Ethology as a physical science. *Nature*
501 *Physics*, 14(7), 653-657.
502 78. Hayden, B. Y. (2018). Economic choice: the foraging perspective. *Current Opinion*
503 *in Behavioral Sciences*, 24, 1-6.
504 79. Ebitz, R. B., Sleezer, B. J., Jedema, H. P., Bradberry, C. W., & Hayden, B. Y.
505 (2019). Tonic exploration governs both flexibility and lapses. *PLoS computational*
506 *biology*, 15(11).
507 80. Berns, G. S., Chappelow, J., Cekic, M., Zink, C. F., Pagnoni, G., & Martin-Skurski,
508 M. E. (2006). Neurobiological substrates of dread. *Science*, 312(5774), 754-758.