PNAS RESEARCH ARTICLE | NEUROSCIENCE
PSYCHOLOGICAL AND COGNITIVE SCIENCES

Learning reinforces curiosity for related information

Yaniv Abir[a,b,1], Jane Mok[c], Christopher A. Baldassano[a], Caroline B. Marvin[d,2], and Daphna Shohamy[a,e,2]

Edited by Dani S. Bassett, University of Pennsylvania, Philadelphia, PA; received October 14, 2025; accepted March 18, 2026, by Editorial Board Member Michael S. Gazzaniga

Human curiosity is dynamic, however the principles governing its fluctuations remain debated. Here, we test two competing hypotheses about how past learning shapes subsequent curiosity and memory. The first, based on the "optimal arousal" theory, proposes that satisfying curiosity reduces subsequent curiosity. The second, grounded in reinforcement learning, suggests that satisfying curiosity strengthens it. To distinguish between these accounts, we analyzed information-seeking decisions from 5,831 participants, who chose whether to wait for answers to a range of questions. We examined how engagement with questions and answers, as well as information prediction errors, influenced subsequent curiosity. Reading satisfying answers increased curiosity compared to reading dissatisfying ones. Critically, this depended on semantic similarity: prior learning enhances subsequent curiosity only when new information is related to previously learned content. These results suggest that curiosity operates as an information-seeking policy learned through reinforcement. Humans may therefore seek information not only to improve future instrumental decisions, but also to learn what to be curious about.

curiosity | reinforcement learning | motivation

Significance

In everyday life, we constantly consume information—but not indiscriminately. Our curiosity directs our search, influenced by what we have just learned. Yet most empirical studies of human curiosity use tasks that isolate it from prior learning. In this study, we examined how curiosity and information-seeking unfold dynamically in response to recent learning. We found that learning satisfying information reinforces curiosity toward related content, challenging theories that frame curiosity as an easily satisfied intrinsic drive to seek information regardless of its utility. These findings offer a foundation for interventions aimed at sparking curiosity in education, public communication, and the media.

Curiosity has long been thought of as inherently dynamic. The 18th century philosopher Edmund Burke wrote that curiosity "changes its object perpetually; it has an appetite which is very sharp, but very easily satisfied; [...] Curiosity from its nature is a very active principle" (1). Despite this active nature, the nascent science of curiosity has often emphasized the static features of curiosity, most notably the personality traits associated with or constituting curiosity (see ref. 2, for a review), and how the informational content of different questions determines curiosity (e.g., refs. 3–5). Relatively little attention has been paid to understanding how curiosity changes from moment to moment. Instead, the dynamics of information seeking were often treated as noise and tasks were typically designed to discourage knowledge attained on one trial from guiding choices on subsequent trials (6–8). Therefore, our understanding of how curiosity fluctuates from one experience to the next, or what triggers sudden interest in previously unappealing questions, is only beginning to develop.

The aim of this study is to examine how each learning experience influences subsequent curiosity, thereby gaining an understanding of the temporal dynamics of curiosity. We consider a learning experience as a question, either posed by the environment or internally generated by the individual (9), and an answer that is revealed if the individual seeks it (4). Accordingly, participants' epistemic behavior concerning question-and-answer pairs was assessed, measuring how the response to learning an answer impacts curiosity for subsequently encountered questions.

We tested two contrasting hypotheses regarding the impact of past learning experiences on subsequent curiosity, rooted in different theoretical perspectives: one grounded in reinforcement learning and the other in optimal arousal theory. The first hypothesis suggests a reinforcing effect, where discovering a satisfying answer enhances curiosity for similar information, while a dissatisfying answer reduces subsequent curiosity. Satisfaction here is used in a general sense, encompassing usefulness (10, 11), the amount of information (12), and valence (13). This hypothesis is inspired by reinforcement learning theory (14, 15), which models curiosity as a motivation shaped by trial-and-error learning to maximize the utility of information. Under this framework, curiosity is adjusted based on the perceived value of past learning, guiding future information-seeking toward useful content. Consequently, the reinforcing effect should be specific: a learning experience should enhance curiosity primarily for similar information but not for dissimilar topics (10, 16).

Alternatively, prior learning experiences may trigger a compensatory effect on subsequent curiosity, as predicted by optimal arousal theory (17, 18). This theory suggests that individuals seek to maintain a stable level of arousal by adjusting their

Author affiliations: [a]Department of Psychology, Columbia University, New York, NY 10027; [b]Max Planck Centre for Computational Psychiatry and Ageing Research, Queen Square Institute of Neurology, University College London, London WC1B 5EH, United Kingdom; [c]Department of Cognitive Science, Barnard College, New York, NY 10027; [d]School of General Studies, Columbia University, New York, NY 10027; and [e]Zuckerman Mind Brain Behavior Institute and Kavli Institute for Brain Science, Columbia University, New York, NY 10027

Author contributions: Y.A., C.B.M., and D.S. designed research; Y.A. and J.M. performed research; C.A.B. contributed new reagents/analytic tools; Y.A. analyzed data; C.A.B. analysis supervision; C.B.M. and D.S. supervision; and Y.A. and D.S. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. D.S.B. is a guest editor invited by the Editorial Board.

[1] To whom correspondence may be addressed. Email: y.abir@ucl.ac.uk.

[2] C.B.M. and D.S. contributed equally to this work.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2529176123/-/DCSupplemental.

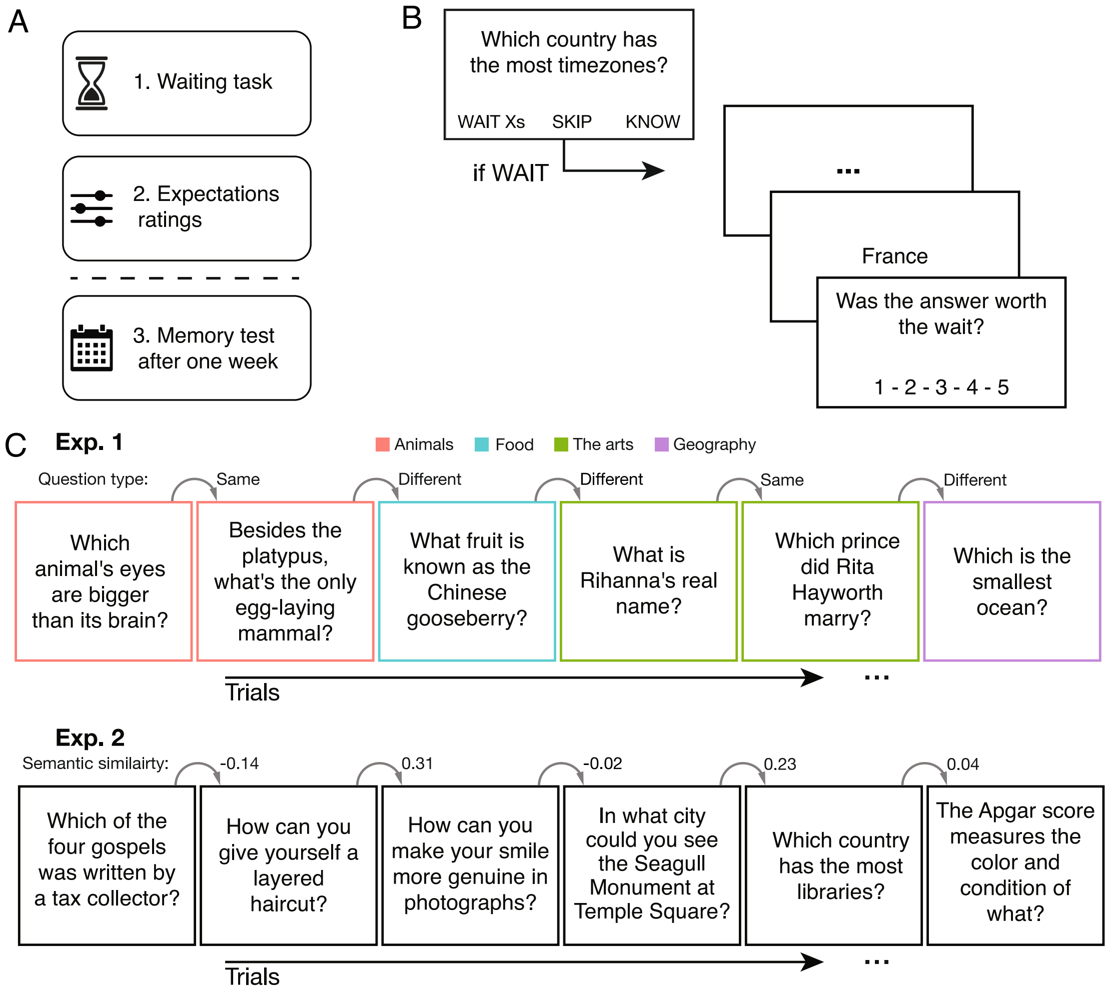

Fig. 1. Measuring the dynamics of information seeking across multiple experiences. (*A*) Information seeking was measured using the waiting task, in which participants decided whether to seek answers to questions of various knowledge domains. Participants also rated their curiosity about a randomly selected subset of questions. One week later, they were asked to recall the answers they had read. (*B*) In each trial of the waiting task, participants were shown a question. If they already knew the answer, they could indicate this. Otherwise, they could choose whether to wait a specified duration (3 to 9 s) to view the answer, after which they rated their satisfaction with the answer. (*C*) In Experiment 1, questions belonged to four distinct categories, presented in random order. In Experiment 2, no distinct categories were used; instead, pairwise semantic similarity served as a continuous measure of semantic space. Example answers: Ostrich, Echidna, Kiwi, Robyn Fenty, Aly Khan, Arctic Ocean, Matthew, make a ponytail on the front and cut, squint, Salt Lake City, probably Russia, newborns.

information-seeking behavior, therefore predicting that curiosity will decrease after a particularly satisfying answer and increase after a dissatisfying one (2, 19). Unlike reinforcement learning, which emphasizes the value of specific information, optimal arousal theory conceptualizes curiosity as a drive for a steady flow of information, regardless of its content (2, 17, 20, c.f., ref. 21). As a result, the compensatory effect should generalize broadly: a satisfying answer should reduce curiosity for any subsequent question, regardless of topic, while a dissatisfying answer should increase curiosity across the board.

To adjudicate between these alternative hypotheses, we measured answer-seeking choices across knowledge domains and levels of complexity (Fig. 1). Experiment 1 investigated this question in a controlled setting with questions from a limited range of categories. Participants' choices to seek answers were measured for four distinct question types presented randomly in the same experimental block: questions about animals, the arts, food, and geography. Curiosity ratings, reflecting expectations elicited by each question, and satisfaction ratings, gauging the response to each answer, were collected. Because curiosity and satisfaction are known to be correlated (10, 13), we decomposed them into two orthogonal components parameterizing each learning experience: the sum of engagement elicited by the question and answer, and the prediction errors elicited by the answer. We examined the effects of these components on subsequent answer-seeking within and across question types, including their impact on memory recall one week later.

In Experiment 2 we tested the same hypotheses in a more complex and naturalistic context. Here, question stimuli varied continuously in content rather than being grouped into distinct categories. This design allowed us to ask how the effect of previous learning experiences on subsequent curiosity generalizes in a continuous fashion across the semantic space of natural language.

Results

Curiosity-driven behavior was assessed in 5,831 participants using the waiting task (10, 13, 22). On each trial, participants read a question and decided whether to wait a variable number of seconds to see the answer. Their choices to wait rather than skip, given different durations, serve as our measure of information

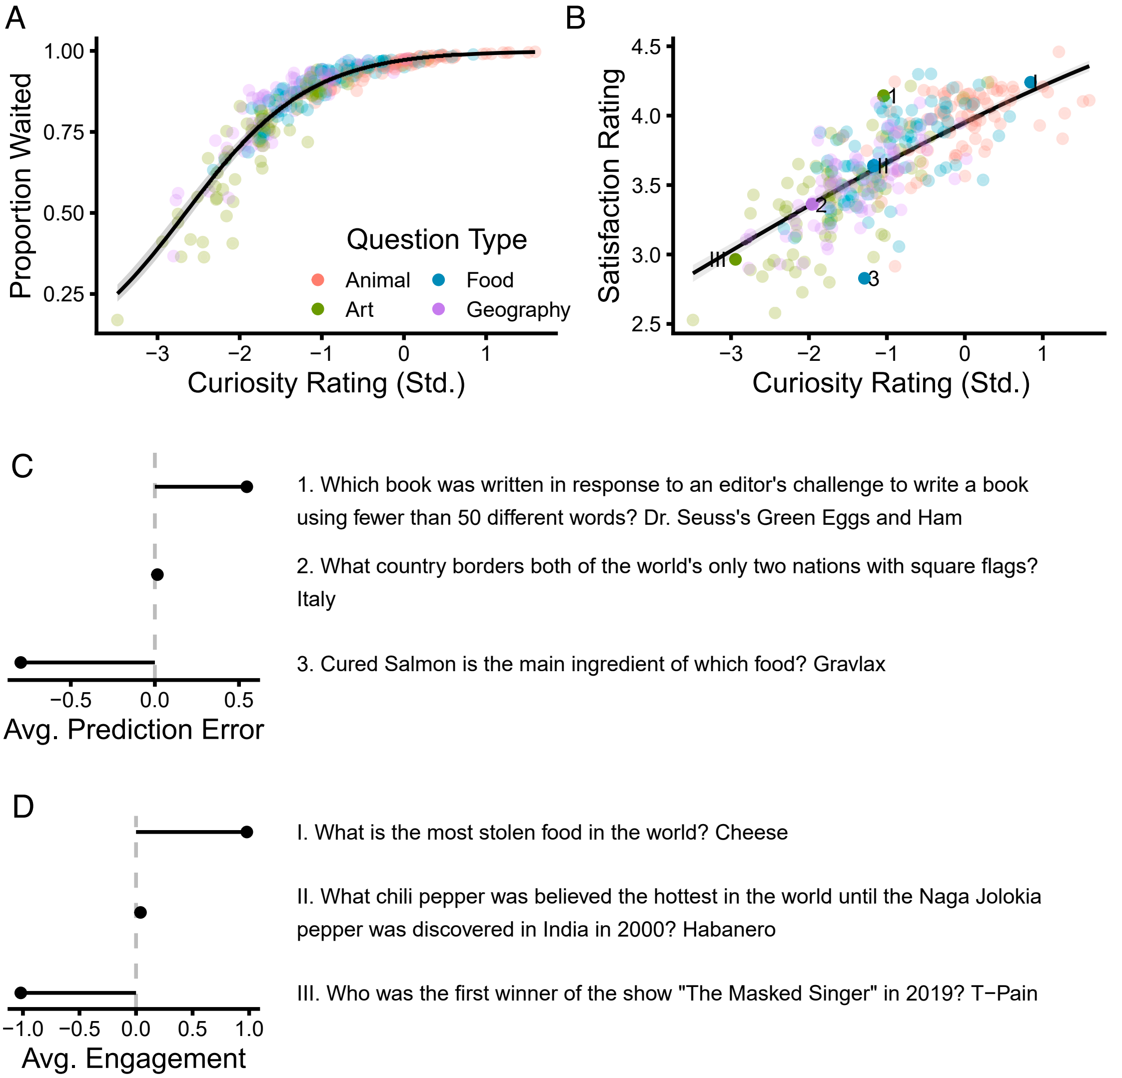

1. Which book was written in response to an editor's challenge to write a book using fewer than 50 different words? Dr. Seuss's Green Eggs and Ham

2. What country borders both of the world's only two nations with square flags? Italy

3. Cured Salmon is the main ingredient of which food? Gravlax

I. What is the most stolen food in the world? Cheese

II. What chili pepper was believed the hottest in the world until the Naga Jolokia pepper was discovered in India in 2000? Habanero

III. Who was the first winner of the show "The Masked Singer" in 2019? T−Pain

Fig. 2. Prediction errors and engagement summarize learning on each trial. (*A*) Curiosity ratings robustly predict participants' decisions to wait for answers. (*B*) Curiosity ratings predict participants' satisfaction with answers they had read. The prediction error is defined as the difference between satisfaction and curiosity, while engagement is defined as the sum of satisfaction and curiosity. (*C*) A selection of question-and-answer stimuli and the prediction error they elicited on average. (*D*) A selection of question-and-answer stimuli and the engagement they elicited on average. Questions from (*C*) and (*D*) are marked as numerals in (*B*). Lines show median predictions from multilevel regression models, ribbons denote 50% PI. Colorful points are per-question adjusted mean from the regression model.

seeking. Participants also rated their satisfaction with the answers they read and were tested on their recall one week afterward. At the end of the experiment, expectations elicited by a randomly held-out set of questions were measured, either as curiosity (Experiment 1) or expected usefulness ratings (Experiment 2). Critically, we varied the semantic content of the question-and-answer stimuli: across four discrete categories in Experiment 1 and along a continuum of semantic space in Experiment 2. This trial-by-trial variation allowed us to examine how expectations and response to one question influence curiosity toward the following similar or dissimilar question.

Overall, the key findings from Experiment 1 were successfully replicated in Experiment 2. We report Experiment 2 results in the main text only where its design meaningfully differed from that of Experiment 1. Results representing direct replications are provided in *SI Appendix*.

Information-Seeking Is Aligned with Participants' Expectations.

Before testing the two competing hypotheses about how learning influences curiosity, we validated our measures of information-seeking, expectations, and satisfaction. As expected, participants' information-seeking was closely aligned to their expectations. In Experiment 1, participants' curiosity ratings were coherent as a group, with high interjudge agreement (ICC2k = 0.99, 95% CI = [0.99,0.99]). Participants' decisions to wait for answers rather than skip to the next question were tightly linked to average curiosity ratings, as revealed by multilevel logistic regression (b = 1.32, 95% posterior interval (PI) = [1.19,1.46]; Fig. 2*A*). Curiosity ratings were well calibrated, as participants tended to be more satisfied with answers to questions they were more curious about, as revealed by multilevel ordinal regression (b = 0.58, 95% PI = [0.48,0.68]; Fig. 2*B*). See *SI Appendix* for a replication of these patterns using usefulness ratings in Experiment 2.

Previous Learning Reinforces Subsequent Information-Seeking.

To test how learning affects curiosity, we calculated engagement and prediction error scores for each answer in the waiting task. Prediction error was the difference between a participant's satisfaction and the estimated expectation for that question.

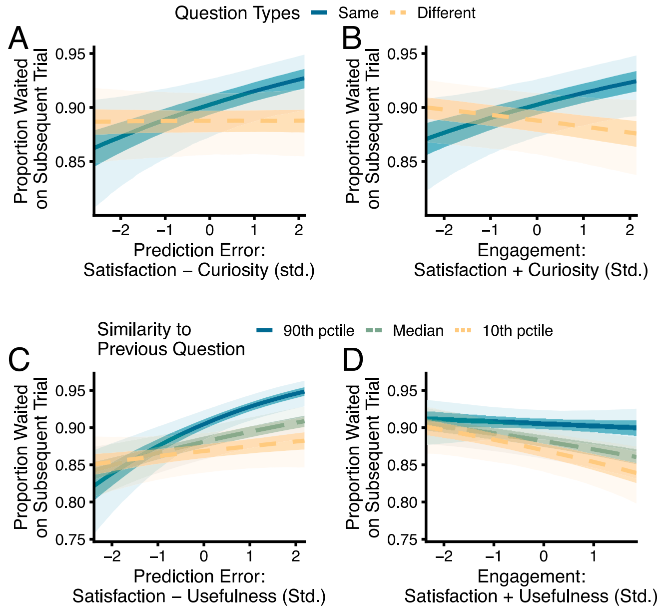

Fig. 3. Previous learning reinforces subsequent waiting for answers. (*A*) In Experiment 1, prediction errors elicited by answers are positively associated with waiting for a subsequent answer of the same type, with no significant effect for different types. (*B*) Engagement from a question and its answer is also positively associated with waiting for a same-type answer, with no significant effect for different types. (*C*) Experiment 2 replicates the reinforcing effect of prediction error on subsequent waiting, moderated by a continuous measure of semantic similarity, with a stronger effect for highly similar questions. (*D*) A similar pattern is observed for engagement, where the interaction with similarity is significant only for the low variability block. Lines are median predictions from multilevel regression models, dark ribbons denote 50% PIs, and light ribbons 95% PIs.

Engagement was the sum of satisfaction and expectation. Expectations were estimated from other participants' curiosity or usefulness ratings, avoiding bias from having participants both rate and choose the same question.

We tested for a reinforcing vs. compensatory effect of learning on subsequent curiosity by examining how these two learning-related scores predict subsequent waiting choices. Both the prediction error and the level of engagement elicited by a previous question-and-answer pair were robust predictors of waiting choices on the next trial. We found a significant reinforcing effect of prediction errors on subsequent waiting in Experiment 1 (b = 0.08, 95% PI = [0.01,0.14]), such that surprisingly satisfying answers increased participants' willingness to wait on the next trial, and surprisingly dissatisfying answers diminished it. Critically, this effect was modulated by the similarity between the previous and subsequent questions (b = 0.07, 95% PI = [0.01,0.14]). While prediction errors exerted a significant influence on subsequent waiting for questions of the same type (b = 0.15, 95% PI = [0.04,0.26]), no significant effect was found for questions of different types (b = 0.005, 95% PI = [−0.06,0.07]; Fig. 3*A*).

The level of engagement elicited by a previous question-and-answer pair was similarly linked to waiting on the next trial. We found a significant interaction between engagement and the match in question types (b = 0.09, 95% PI = [0.03,0.15]). Previous trial engagement was positively associated with choosing to wait for an answer of the same type (b = 0.13, 95% PI = [0.03,0.24]). When the subsequent trial had a different question type, there was a trend toward a negative association between previous trial engagement and waiting (b = −0.05, 95% PI = [−0.12,0.007], Fig. 3*B*). Thus, question-and-answer engagement also had a reinforcing effect on subsequent information-seeking choices.

The effects of prediction errors and engagement were consistent across all four question types. A multilevel logistic regression model with question type interaction terms revealed no significant interaction of question type with any of the effects reported above (*SI Appendix*, Table S2). Comparing models with and without question type interactions using 10-fold cross-validation strongly favored the simpler model, with the difference in prediction score far exceeding the SE (ELPD difference = −37.46, SE = 13.94). This indicates that the reinforcing effect of learning holds consistently across question types.

Learning Reinforces Curiosity Across Continuous Semantic Space. Experiment 2 allowed us to examine the generalization of previous learning across continuous semantic space, rather than discrete question categories. The experiment consisted of two blocks, one of relatively low semantic variability (with questions all concerning different aspects of the COVID-19 pandemic), and one of higher variability (questions concerning all manner of trivial and useful knowledge; *SI Appendix*, Fig. S1). To compare stimuli with varying semantic content, we computed similarity scores between each pair of questions using the Sentence-BERT model (23). Similarity scores were computed as the cosine between model embeddings for a pair of stimuli, ranging from 0 (dissimilar) to 1 (identical). See *SI Appendix* and *SI Appendix*, Fig. S2 for validation of this measure.

The same patterns observed in Experiment 1 emerged in Experiment 2. Prediction errors and the engagement elicited by a previous question-and-answer pair were robust predictors of waiting choices on the next trial, and these effects were moderated by semantic similarity. Prediction errors were positively associated with waiting on subsequent trials (b = 0.16, 95% PI = [0.09,0.23]), and especially so when questions on the current trial were more semantically similar to those on the previous trial (interaction b = 0.09, 95% PI = [0.03,0.16]; Fig. 3*C*).

Additionally, there was a three-way interaction between previous trial engagement, question similarity, and block (b = 0.08, 95% PI = [0.02,0.15]). In the low-variability block, there was a significant interaction between previous engagement and question similarity. High engagement in this block was associated with decreased waiting for dissimilar subsequent questions (b = 0.12, 95% PI = [0.03,0.21]). This relationship was not significant in the high variability block (b = −0.05, 95% PI = [−0.15,0.05]). On average, there was a negative effect of engagement on subsequent waiting (b = −0.09, 95% PI = [−0.15,−0.03]; Fig. 3*D*). Together, these results suggest that learning not only drives curiosity for similar content but that this generalization is shaped by semantic structure.

Individual Differences in Interests and Learning Effects. Accounting for individual differences in interest was essential for detecting the reinforcing effect of learning. When prediction errors and engagement scores were computed using group-averaged satisfaction ratings rather than each participant's own ratings, neither showed significant effects in Experiment 1, nor significant interactions with semantic similarity in Experiment 2 (*SI Appendix*, Fig. S3). This suggests that the observed effects reflect individual learning processes, rather than common structure in the questions and answers. Furthermore, across both experiments, we found robust differences in overall curiosity and in the reinforcing effects of prediction errors and engagement,

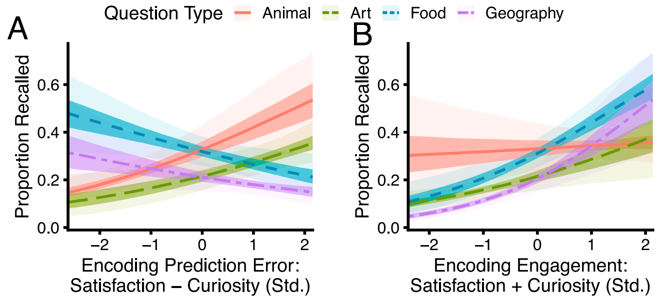

Fig. 4. Memory one week later is dependent on question-and-answer content and engagement. (*A*) In Experiment 1, prediction error elicited by an answer is not a consistent predictor of memory, with different associations for each question type. (*B*) However, reading high engagement questions and answers is consistently associated with better memory for answers. Lines are median predictions from multilevel regression models, dark ribbons denote 50% PIs and light ribbons 95% PIs.

in line with prior work documenting individual differences in the determinants of curiosity (24–26). However, we found no evidence for distinct clusters of curiosity styles (*SI Appendix*), though detecting such structure may require paradigms with substantially more trials per participant.

Memory Recall Driven by Content, Not Prior Learning. We next asked whether the same learning signals—prediction error and engagement—also influence memory for the information learned. Prior studies have shown that prediction errors at encoding can enhance memory (13, 27–29). Using data from Experiment 1, we fit a model predicting recall from both prediction error and engagement associated with each answer. Results were mixed. Engagement had a positive overall effect on memory ($b = 0.42$, 95% PI = [0.27, 0.57]), but this effect varied by question type and was significant only for food and geography questions (Fig. 4*B* and *SI Appendix*, Table S4). Prediction error had no significant main effect ($b = 0.06$, 95% PI = [−0.06, 0.19]), but interacted with question type: only answers about animals and art showed improved memory with higher prediction error (Fig. 4*A* and *SI Appendix*, Table S4).

To directly test the effect of prediction error, we fit a reduced model excluding engagement. Here, prediction error significantly predicted memory ($b = 0.38$, 95% PI = [0.30, 0.46]; see *SI Appendix*, Table S5 for effects by question type). Experiment 2 showed similar patterns (*SI Appendix*, Fig. S6). Finally, we asked whether learning signals from one trial affected memory for the next. We did not find support for this idea: neither prediction error nor engagement, nor their interactions with semantic similarity, reliably predicted next-trial recall (*SI Appendix*, Table S3). Thus, memory depended on the content of each question, regardless of whether curiosity arose from the question itself or was shaped by prior learning.

Discussion

Together, these findings show that information-seeking behavior is shaped by the context and history of learning. Learning satisfying information increases curiosity for related content, while reducing curiosity for unrelated content. This pattern, predicted by reinforcement learning accounts of motivation, is consistent with viewing curiosity as an information-seeking policy learned from experience.

This reinforcing effect of learning on curiosity involves two aspects: an increase in curiosity for similar information and a decrease for dissimilar information. However, one effect does not necessarily imply the other. Curiosity can increase for similar information without a corresponding decrease for dissimilar information. This independence is a feature of the waiting task, which like free operant paradigms in animal models (30), allows participants to separately control both the amount and type of information they seek. Furthermore, these two effects may arise from distinct learning mechanisms.

Enhanced curiosity for related information is readily explained by standard value-based reinforcement learning models, such as Q-learning (31), in which prediction errors selectively update the value of the chosen information domain. However, explaining the concurrent decrease in curiosity for unrelated information requires a mechanism that allows for reciprocal updating: as the value of one domain increases, the relative value of others must decrease. This can be achieved either by extending Q-learning to include reciprocal value updates, or by adopting policy-gradient models (32, 33), in which learning updates a vector of action propensities, increasing curiosity for similar information while suppressing curiosity for alternatives (34).

A conceptually related account is provided by foraging theories (35, 36), in which agents compare the value of a current option against an evolving average of available alternatives. When curiosity rises for one category, the cross-category average increases, raising the opportunity cost of seeking information from other domains and thereby reducing curiosity for them. These frameworks make overlapping predictions but they differ in their underlying mechanisms. Further work will be needed to disentangle whether curiosity declines for dissimilar information due to reciprocal value updating or average-based comparison processes.

Interpreting our findings through the lens of reward-based learning theories requires that we recognize the important distinctions in the underlying constructs between the two domains. In particular, the prediction error derived from the waiting task may not directly correspond to a reward prediction error. Its components, reported satisfaction and expectations of curiosity or usefulness, are not direct analogues of reward and expected value, in part because they draw on semantic memory as well as valuation processes (see *SI Appendix*, Fig. S6 for a preliminary analysis). Further work is needed to understand how semantic memory contributes to reinforcement learning in this task by supporting separable option representations and the computation of expectations and satisfaction.

In both experiments (and prior work, 10, 13), satisfaction was derived from participants' judgment of whether an answer was "worth the wait." Because participants trade off waiting durations against curiosity in their choices (10, 22), satisfaction reports may already be normalized relative to expectations (c.f., ref. 37). Decomposing satisfaction and expectations into orthogonal components helps disambiguate these findings, but refinement of these measures is necessary to fully capture the psychological constructs underlying prediction errors in the waiting task. Another limitation is the use of other participants' ratings of curiosity or usefulness to approximate each participant's expectations. While this approach reduces potential response bias, it disregards individual interests by relying on a group average. Nevertheless, the high intersubject agreement in curiosity and usefulness ratings observed here and in prior work (10, 13, 28), partially mitigates this concern, though modifications to the waiting task may be needed in future studies to fully address it.

Recent theories have emphasized the significance of information prediction errors for curiosity (38–40). Examining the role of information prediction errors in behavior has proven a successful approach to discovering neural correlates of curiosity (7, 8, 39, 41, 42), and understanding how curiosity affects subsequent memory (13). The present study provides a direct test of the idea that curiosity is learned from experience, with prediction error serving as the learning signal.

This study directly tests the distinction, articulated by White (21), between learning-based theories of motivation and optimal arousal theory, a member of the drive-based family. In machine learning and robotics, a related distinction was drawn by Oudeyer and Kaplan (43) between homeostatic motivations, which regulate behavior around a set point, and heterostatic motivations, which drive agents away from their habitual state. In that literature, heterostatic, reinforcement-based approaches have proven highly effective: curiosity is modeled as an intrinsic reward that drives agents to seek experiences maximizing learning progress (43–48). Recent work in humans using finite-information tasks, carefully designed exploration games where information is quantifiable, has revealed analogous effects, showing how people track and seek out learnable experiences (49–51). Our study extends these insights to open-ended curiosity, where the space of potential information to be learned is effectively unbounded, enabling predictions about how curiosity evolves across learning episodes.

Our findings join a small group of recent studies applying new tools to examine the dynamics of curiosity. Using network-theoretic measures, Lydon-Staley et al. (25) have described curiosity as a random walk through semantic space. Using a task in which information unfolds over time, Hsiung and colleagues (52) described how curiosity evolves in tandem with emotion within a single learning experience. Related approaches have also been used to study how the structure and sequencing of questions shape explanation-seeking and generation over time (53–55). Our study extends this work by identifying the key factors that determine how curiosity changes from one experience to the next, enabling us to predict how the locus of curiosity shifts learning episodes.

The idea that curiosity is learned from experience is consistent with a broader reevaluation of motivations traditionally viewed as intrinsic and distinct from extrinsic motivation (56, c.f., refs. 21 and 57). The findings presented here support the evolving understanding that intrinsically motivated behaviors, such as curiosity, idea generation and question asking, self-efficacy, and play, adhere to the same systematic learning rules as reward-based behaviors, and do not necessarily need to be considered as a separate natural kind (41, 57–63).

Contemporary scientific discourse often characterizes human curiosity as irrationally excessive (2, 20), noting that the direct instrumental value of information people seek is often lower than the costs incurred to obtain it (2, 24, 64). This characterization rests on the assumption that information is valuable only insofar as it guides reward-maximizing decisions. Recognizing that information also teaches individuals what to be curious about challenges this view. Formulating a good information-seeking policy is itself not trivial, in part because estimating the value of information before it is acquired is computationally difficult (14, 65). As with other flexible behaviors, information-seeking must therefore be learned from experience, a process that requires information as its input.

Materials and Methods

Data Collection and Participants. Data were collected from two independent samples. In Experiment 1, 500 participants were recruited via Amazon Mechanical Turk ($n = 200$) and CloudResearch Connect ($n = 300$) (66). Following Amazon's suspension of the first author's account, data collection resumed on Connect, where participants were also invited to complete a second session one week later (79.33% returned). In Experiment 2, we reanalyzed data from 5,376 participants collected on Amazon Mechanical Turk during the early days of the COVID-19 pandemic (March to May 2020; 10, 67).

The same exclusion criteria were applied in both experiments, based on previous studies using this task (10, 13). In Experiment 1, data from 18 participants (3.65%) reporting less than perfect English language fluency and 27 participants (5.48%) who interacted with other applications more than five times were excluded from the analysis. Data from participants whose mean response time was >2 SD longer than the group average ($n = 3$, 0.61%) were further excluded, leaving 445 participants for analysis (median age = 38, range = 18 to 78; 192 female, 247 male, 6 other). We separately excluded data from the second session if participants had more than five application interactions during the recall task ($n = 13$, 2.92%). In Experiment 2, the final analyzed sample consisted of 5,376 participants (median age = 36, range = 18 to 89; 2,818 female). Further details on exclusions for Experiment 2 are available in ref. 10.

All participants provided informed consent, and all protocols were approved by the Columbia University Institutional Review Board.

Stimuli and Task. Stimuli consisted of trivia question-and-answer pairs. In Experiment 1, 300 questions were evenly divided into four categories (animals, the arts, geography, and food), while in Experiment 2, 104 questions were presented in two blocks: one with only COVID-19-related questions and one with a variety of general questions. Questions in Experiment 2 were not categorized otherwise. Questions were sourced from previous studies (10, 13) and online trivia databases. Both experiments measured four aspects of epistemic behavior: information-seeking, expectations, satisfaction, and memory. *Information-seeking* was assessed using a validated waiting task (10, 13, 22, 28). On each trial of the waiting task, participants were presented with a question, and three choice buttons. If they knew the answer to the question, they were instructed to press "know." Otherwise, they could choose to wait a specified duration for the answer by pressing "wait Xs" (3 to 9 s in Experiment 1, 4 to 16 s in Experiment 2, randomly assigned), or else press the "skip" button, which terminated the trial. The proportion of "wait" vs. skip responses at variable waiting durations serves as our main index of information-seeking. Participants were instructed that task duration was fixed, and were therefore encouraged to use their own interest to decide whether to wait. After receiving an answer, they rated their *satisfaction* ("Was the answer worth the wait?") on a 1 to 5 Likert scale (Fig. 1*B*).

Expectations were measured differently across experiments. In Experiment 1, curiosity ratings were collected for a separate set of 30 questions. In Experiment 2, participants instead rated a set of 10 questions on the predicted usefulness of each answer, a measure that similarly predicts waiting behavior (10) and aligns with computational models of curiosity (3, 11, 60).

Memory was tested after 7 to 8 d, using a recall task on previously waited-for questions. Participants indicated whether they remembered the answer and provided a response, which was coded as correct, partially correct, or incorrect by a research assistant blind to the conditions and hypotheses.

Finally, both experiments included motivation and affect questionnaires, though these data were not analyzed here.

Data Analysis. The primary goal of analysis was to test whether learning from previous trials influenced subsequent information seeking, either through a reinforcing effect, a compensatory effect, or no effect at all. To this end, we developed measures summarizing learning on the previous trial and examined whether they predicted subsequent waiting choices. We further investigated how these effects were modulated by the similarity between consecutive stimuli and whether they generalized across different question types and experimental contexts.

Since subjective ratings of curiosity, satisfaction, and usefulness are known to be correlated (10, 13), we decomposed them into two orthogonal components to avoid multicollinearity and enhance interpretability. The first component, prediction error, was defined as the difference between satisfaction and expectations (curiosity ratings in Experiment 1; usefulness ratings in Experiment 2). This measure captures participants' responses to answers that defied expectations, a key concept in reinforcement learning (14, 15). The second component, engagement, was defined as the sum of satisfaction and expectations, capturing the extent to which participants' responses aligned with their prior expectations.

Expectation estimates in both experiments were obtained using a multilevel item-response model for rating data (*SI Appendix*). Prior to computing the two component scores, we standardized expectation estimates and the satisfaction ratings (the latter separately for each participant) to account for differences in the use of the rating scales.

To assess generalization across semantic space, we quantified the similarity between each question and the one preceding it using Sentence-BERT embeddings (23). The cosine similarity between embeddings provided a continuous measure ranging from 0 (totally dissimilar) to 1 (identical). See *SI Appendix* and *SI Appendix*, Fig. S2 for validation details.

Data, Materials, and Software Availability. Data and experimental code for Experiment 1 are available on OSF (https://osf.io/h9fqy) (66). Data and experimental code for Experiment 2 are available on OSF (https://osf.io/gjcu9) (67). The analysis code and computational environment are also available at https://osf.io/h9fqy (66).

ACKNOWLEDGMENTS. Portions of this work were developed from the doctoral thesis of Y.A. We thank the Shohamy lab for their insightful discussion of the project. We are grateful for funding support from the NSF (Award #1822619 to D.S.), NIMH/NIH (#MH121093 to D.S.) and the Templeton Foundation (#60844 to D.S.).

1. E. Burke, *The Works of the Right Honourable Edmund Burke, Vol. 01 (of 12)*. (Henry G. Bohn, London, 1854).
2. G. Loewenstein, The psychology of curiosity: A review and reinterpretation. *Psychol. Bull.* **116**, 75–98 (1994).
3. R. Dubey, T. L. Griffiths, Reconciling novelty and complexity through a rational analysis of curiosity. *Psychol. Rev.* **127**, 455–476 (2019).
4. R. Golman, G. Loewenstein, Information gaps: A theory of preferences regarding the presence and absence of information. *Decision* **5**, 143–164 (2018).
5. T. Sharot, C. R. Sunstein, How people decide what they want to know. *Nat. Hum. Behav.* **4**, 14–19 (2020).
6. L. FitzGibbon, A. Komiya, K. Murayama, The lure of counterfactual curiosity: People incur a cost to experience regret. *Psychol. Sci.* **32**, 241–255 (2021).
7. L. L. F. van Lieshout, A. R. E. Vandenbroucke, N. C. J. Müller, R. Cools, F. P. de Lange, Induction and relief of curiosity elicit parietal and frontal activity. *J. Neurosci.* **38**, 2579–2588 (2018).
8. C. J. Charpentier, E. S. Bromberg-Martin, T. Sharot, Valuation of knowledge and ignorance in mesolimbic reward circuitry. *Proc. Natl. Acad. Sci. U.S.A.* **115**, E7255–E7264 (2018).
9. T. Raz, Y. Kenett, Questions in information seeking. A primer. PsyArXiv [Preprint] (2024). https://doi.org/10.31234/osf.io/qzsgf_v1 (Accessed 6 April 2026).
10. Y. Abir *et al*., An energizing role for motivation in information-seeking during the early phase of the COVID-19 pandemic. *Nat. Commun.* **13**, 2310 (2022).
11. R. Dubey, T. L. Griffiths, T. Lombrozo, If it's important, then I'm curious: Increasing perceived usefulness stimulates curiosity. *Cognition* **226**, 105193 (2022).
12. N. Chater, G. Loewenstein, The under-appreciated drive for sense-making. *J. Econ. Behav. Organ.* **126**, 137–154 (2016).
13. C. B. Marvin, D. Shohamy, Curiosity and reward: Valence predicts choice and information prediction errors enhance learning. *J. Exp. Psychol. Gen.* **145**, 266–272 (2016).
14. R. S. Sutton, A. G. Barto, *Reinforcement Learning: An Introduction* (The MIT Press, ed. 2, 2018), p. xxii, 526.
15. Y. Niv, Reinforcement learning in the brain. *J. Math. Psychol.* **53**, 139–154 (2009).
16. Y. Niv, D. Joel, P. Dayan, A normative perspective on motivation. *Trends Cogn. Sci.* **10**, 375–381 (2006).
17. D. E. Berlyne, Curiosity and exploration. *Science* **153**, 25–33 (1966).
18. D. O. Hebb, Drives and the CNS (conceptual nervous system). *Psychol. Rev.* **62**, 243 (1955).
19. R. Ligneul, M. Mermillod, T. Morisseau, From relief to surprise: Dual control of epistemic curiosity in the human brain. *NeuroImage* **181**, 490–500 (2018).
20. C. Kidd, B. Y. Hayden, The psychology and neuroscience of curiosity. *Neuron* **88**, 449–460 (2015).
21. R. W. White, Motivation reconsidered: The concept of competence. *Psychol. Rev.* **66**, 297 (1959).
22. M. J. Kang *et al*., The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychol. Sci.* **20**, 963–973 (2009).
23. N. Reimers, I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks" in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing* (Association for Computational Linguistics, 2019).
24. K. Kobayashi, S. Ravaioli, A. Baranès, M. Woodford, J. Gottlieb, Diverse motives for human curiosity. *Nat. Hum. Behav.* **3**, 587–595 (2019).
25. D. M. Lydon-Staley, D. Zhou, A. S. Blevins, P. Zurn, D. S. Bassett, Hunters, busybodies and the knowledge network building associated with deprivation curiosity. *Nat. Hum. Behav.* **5**, 327–336 (2021).
26. D. Zhou *et al*., Architectural styles of curiosity in global Wikipedia mobile app readership. *Sci. Adv.* **10**, eadn3268 (2024).
27. G. M. Fastrich, T. Kerr, A. D. Castel, K. Murayama, The role of interest in memory for trivia questions: An investigation with a large-scale database. *Motiv. Sci.* **4**, 227–250 (2018).
28. E. A. Lang, C. Geen, E. Tedeschi, C. B. Marvin, D. Shohamy, Learned temporal statistics guide information seeking and shape memory. *J. Exp. Psychol. Gen.* **151**, 986 (2022).
29. Y. Fandakova, M. J. Gruber, States of curiosity and interest enhance memory differently in adolescents and in children. *Dev. Sci.* **24**, e13005 (2021).
30. Y. Niv, N. Daw, P. Dayan, How fast to work: Response vigor, motivation and tonic dopamine. *Adv. Neural Inf. Process. Syst.* **18**, 1019–1026 (2005).
31. C. J. Watkins, P. Dayan, Q-learning. *Mach. Learn.* **8**, 279–292 (1992).
32. R. J. Williams, Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.* **8**, 229–256 (1992).
33. J. Li, N. D. Daw, Signals in human striatum are appropriate for policy upyear rather than value prediction. *J. Neurosci.* **31**, 5504–5511 (2011).
34. N. Biderman, S. J. Gershman, D. Shohamy, The role of memory in counterfactual valuation. *J. Exp. Psychol. Gen.* **152**, 1754 (2023).
35. R. Ackerman, The diminishing criterion model for metacognitive regulation of time investment. *J. Exp. Psychol.: Gen.* **143**, 1349 (2014).
36. T. T. Hills, P. M. Todd, D. Lazer, A. D. Redish, I. D. Couzin, Exploration versus exploitation in space, mind, and society. *Trends Cogn. Sci.* **19**, 46–54 (2015).
37. E. Liquin, More time and effort, same curiosity: Expected effort does not impact curiosity. *Open Mind* **10**, 391–414 (2026).
38. M. J. Gruber, C. Ranganath, How curiosity enhances hippocampus-dependent memory: The prediction, appraisal, curiosity, and exploration (PACE) framework. *Trends Cogn. Sci.* **23**, 1014–1025 (2019).
39. E. S. Bromberg-Martin, O. Hikosaka, Lateral habenula neurons signal errors in the prediction of reward information. *Nat. Neurosci.* **14**, 1209–1216 (2011).
40. F. Poli, J. X. O'Reilly, R. B. Mars, S. Hunnius, Curiosity and the dynamics of optimal exploration. *Trends Cogn. Sci.* **28**, 441–453 (2024).
41. E. S. Bromberg-Martin, O. Hikosaka, Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron* **63**, 119–126 (2009).
42. I. E. Monosov, Curiosity: Primate neural circuits for novelty and information seeking. *Nat. Rev. Neurosci.* **25**, 195–208 (2024).
43. P. Y. Oudeyer, F. Kaplan, V. V. Hafner, Intrinsic motivation systems for autonomous mental development. *IEEE Trans. Evol. Comput.* **11**, 265–286 (2007).
44. D. Pathak, P. Agrawal, A. A. Efros, T. Darrell, "Curiosity-driven exploration by self-supervised prediction" in *International Conference on Machine Learning* (PMLR, 2017), pp. 2778–2787.
45. Y. Burda *et al*., Large-scale study of curiosity-driven learning. arXiv [Preprint] (2018). https://doi.org/10.48550/arXiv.1808.04355 (Accessed 6 April 2026).
46. N. Haber, D. Mrowca, S. Wang, L. F. Fei-Fei, D. L. Yamins, Learning to play with intrinsically-motivated, self-aware agents. *Adv. Neural Inf. Process. Syst.* **31**, 8388–8399 (2018).
47. R. Sekar *et al*., "Planning to explore via self-supervised world models" in *International Conference on Machine Learning* (PMLR, 2020), pp. 8583–8592.
48. N. Rhinehart *et al*., Information is power: Intrinsic control via information capture. *Adv. Neural Inf. Process. Syst.* **34**, 10745–10758 (2021).
49. A. Ten, P. Kaushik, P. Y. Oudeyer, J. Gottlieb, Humans monitor learning progress in curiosity-driven exploration. *Nat. Commun.* **12**, 5972 (2021).
50. Y. Abir, M. Shadlen, D. Shohamy, Human exploration strategically balances approaching and avoiding uncertainty. *eLife* **13**, RP94231 (2024).
51. J. Gottlieb, P. Y. Oudeyer, M. Lopes, A. Baranes, Information-seeking, curiosity, and attention: Computational and neural mechanisms. *Trends Cogn. Sci.* **17**, 585–593 (2013).
52. A. Hsiung, J. H. Poh, S. A. Huettel, R. A. Adcock, Curiosity evolves as information unfolds. *Proc. Natl. Acad. Sci. U.S.A.* **120**, e2301974120 (2023).
53. K. Kedrick, R. Golman, Explanation generation: Questions direct exploration in semantic space. PsyArXiv [Preprint] (2025). https://doi.org/10.31234/osf.io/shnkf_v1 (Accessed 6 April 2026).
54. K. Kedrick, W. Yang, T. Gebhart, Y. Wang, R. J. Funk, Opening knowledge gaps drives scientific progress. arXiv [Preprint] (2025). https://doi.org/10.48550/arXiv.2509.21899 (Accessed 6 April 2026).
55. E. G. Liquin, T. Lombrozo, Motivated to learn: An account of explanatory satisfaction. *Cogn. Psychol.* **132**, 101453 (2022).
56. E. L. Deci, R. M. Ryan, *Intrinsic Motivation and Self-Determination in Human Behavior* (Springer, New York, NY, 2010).
57. A. W. Kruglanski *et al*., A structural model of intrinsic motivation: On the psychology of means-ends fusion. *Psychol. Rev.* **125**, 165 (2018).
58. T. Doan, A. Castro, E. Bonawitz, S. Denison, "Wow, I did it!": Unexpected success increases preschoolers' exploratory play on a later task *Cogn. Dev.* **55**, 100925 (2020).
59. J. K. L. Lau, H. Ozono, K. Kuratomi, A. Komiya, K. Murayama, Shared striatal activity in decisions to satisfy curiosity and hunger at the risk of electric shocks. *Nat. Hum. Behav.* **4**, 531–543 (2020).
60. P. Schwartenbeck *et al*., Computational mechanisms of curiosity and goal-directed exploration. *eLife* **8**, 1–45 (2019).
61. Y. N. Kenett, S. Humphries, A. Chatterjee, A thirst for knowledge: Grounding curiosity, creativity, and aesthetics in memory and reward neural systems. *Creat. Res. J.* **35**, 412–426 (2023).
62. S. Moreno-Rodriguez, B. Béranger, E. Volle, A. Lopez-Persem, The human reward system encodes the subjective value of ideas during creative thinking. *Commun. Biol.* **8**, 37 (2025).
63. K. Ding, R. He, X. Wang, Q. Chen, Y. N. Kenett, Recognizing ideas generated in a creative task: The roles of the hippocampus and medial prefrontal cortex in facilitating self-generated learning. *Cereb. Cortex* **34**, bhae219 (2024).
64. T. C. Blanchard, B. Y. Hayden, E. S. Bromberg-Martin, Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron* **85**, 602–614 (2015).
65. E. Schulz, S. J. Gershman, The algorithmic architecture of exploration in the human brain. *Curr. Opin. Neurobiol.* **55**, 7–14 (2019).
66. Y. Abir, J. Mok, C. Baldassano, C. Marvin, D. Shohamy, Learning reinforces curiosity. Open Science Framework. https://doi.org/10.17605/OSF.IO/H9FQY. Deposited 6 April 2026.
67. Y. Abir, C. Marvin, C. van Geen, M. Leshkowitz, R. Hassin, D. Shohamy, An energizing role for motivation in information-seeking during the early phase of the COVID-19 pandemic. Open Science Framework. https://doi.org/10.17605/OSF.IO/GJCU9. Deposited 6 April 2026.