bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Title

From curiosity relief to epistemic surprise: complementary roles of the prefrontal cortex and the ventral striatum in the neural valuation of knowledge

Authors

Romain Ligneul[1,2]*, Martial Mermillod[3], Tiffany Morisseau[1,4]

Affiliations

[1]Institut des Sciences Cognitives, CNRS, Lyon, France

[2]Donders Institute for Brain, Cognition and Behavior, Nijmegen, The Netherlands

[3]Université Pierre Mendès France, LPNC, Grenoble, France

[4]Central European University, Budapest, Hungary

*To whom correspondence should be addressed (r.ligneul@donders.ru.nl)

Number of figures: 6 (+ 3 supplementary)

Main text: 4759 words

Online methods: 2949 words

Figure captions: 1133 words

Abstract: 187 words

References (including Methods): 53

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Abstract

Epistemic curiosity (EC) is a cornerstone of human cognition that contributes to the actualization of our cognitive potential by stimulating a myriad of information-seeking behaviors. Understanding the neural control of EC requires interdisciplinary crosstalks at the theoretical and methodological levels. Using a trivia quiz performed under fMRI in which answer uncertainty was manipulated, we provide behavioral and neural evidence for an integrative model of EC inspired by predictive coding. Behavioral analyses supported a hypothesis derived from this theoretical framework according to which previously experienced surprise should reduce subsequent EC levels. While suppression of neural activity in the rostrolateral prefrontal cortex implemented this key regulatory mechanism, the ventromedial prefrontal cortex coordinated with an array of other brain regions to integrate several dimensions of knowledge valuation, including surprise itself. Following the logics of temporal-difference learning, the ventral striatum encoded curiosity relief only when answer delivery was stochastic. Finally, curiosity, prior knowledge and surprise concurred to predict subsequent memory recall, with surprise mediating curiosity-driven memory benefits. By reconciling different views on the neurocognitive underpinnings of knowledge valuation, these findings may provide a fertile ground for the burgeoning neuroscience of curiosity.

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Introduction

Termed epistemic curiosity (EC), the "motivation to know" predicts educational success[1], orients our attention[2,3] and underlies many decisions of our everyday life, such as opening a book, browsing the internet, watching movies or engaging in trivia quizzes. To date, the complexity of curiosity-related behaviors remains however a challenging issue for the neurosciences of cognitive control, reinforcement-learning and memory. Deeply involved in individuals' success for survival and reproduction, EC seems to outreach the information required to fulfill these essential needs. Indeed, it extends to issues with unclear or indirect biological value such as philosophy, cosmology or art. Following Aristotle's thesis that "all men by nature desire to know"[4], it has long been suggested that knowledge might act as an intrinsic reward and curiosity as an innate drive in humans[5–7]. Yet, such accounts implicitly turn knowledge itself into an evolutionary goal and leave aside the developmental and situational determinants of curiosity.

The predictive coding framework provides promising hypotheses to go beyond this conception of knowledge as a reward. Indeed, the fundamental principle of predictive coding is that a primary function of our cognitive systems is to actively reduce uncertainty relative to the upcoming states of the world, which constitutes a straightforward rationale for information-seeking behaviors. However, predictive coding accounts of curiosity must confront the "dark room" problem stating that, once specific sources of uncertainty have been addressed, avoiding stimulation and refraining from acting appear as the most efficient way to escape new sources of uncertainty[8,9]. Therefore, another principle is required to reconcile this framework with the manifold exploratory behaviors — including those energized by epistemic curiosity — that are not restricted to the short-term minimization of uncertainty (e.g googling the correct spelling of a word). Indeed, exploration often transiently increases uncertainty (e.g pressing the "random" button of a Wikipedia page) and occasionally leads to sustained doubtful states (e.g reading Descartes). In order to accommodate this objection, proponents of predictive coding have suggested that individuals would be born with (and would continuously update) second-order expectations regarding the average amount of surprise experienced when interacting with their environment, and that they would actively try to fulfill these expectations by engaging or disengaging behaviors susceptible of eliciting surprise[9–11].

Here, we embed these two predictive coding principles (reducing uncertainty and adjusting the level of surprise experienced during exploratory behaviors) into a single, integrative model of epistemic curiosity (Fig. 1a). Inspired by the former concepts of "specific" and "diversive" EC early proposed by

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Berlyne[5,12], the model distinguishes two components of curiosity whose behavioral weights may vary from one person to another and from one context to another[6,13]. Specific EC is defined as the transient motivation to seek a solution to well-defined epistemic problems (e.g. reading a tutorial to make something work). By contrast, nonspecific (or diversive) EC is defined as the motivation to seek epistemic stimulation in general (e.g. reading a newspaper or going to cinema). Therefore, only nonspecific EC can stimulate the quest of genuinely new epistemic problems, as seen for example in people who appear to pursue knowledge "for its own sake"[14]. Nonspecific EC is also the only component adjusted to equate experienced and expected surprise, which turns into a clear-cut behavioral prediction: nonspecific EC level should increase when the environment is, on average, less surprising than expected, and vice versa (Fig. 1b).

Consequently, our model dissociates two cognitive processes elicited by the obtention of new information: curiosity relief and surprise (see also Table S1 summarizing the different concepts related to Fig. 1). Curiosity relief reflects the motivational process engaged whenever one becomes aware of a "gap" in one's own knowledge[7]. In line with the first principle of predictive coding described above, these information gaps are hypothesized to elicit uncertainty states of negative valence[13,15,16], so that their resolution might participate in the reinforcement of specific EC likewise pain relief participates in the reinforcement of escape behaviors[17]. By contrast, epistemic surprise goes beyond the reduction of knowledge gaps and relates to the incongruence of preexisting representations and incoming information. Indeed, while the obtention of new information can relieve specific states of ignorance, it can also challenge the coherence and the completeness of one's own prior knowledge about a given topic. Finally, these transient surprise signals may be used update an estimate of the average surprise experienced by the organism leading to the adjustment nonspecific EC levels.

Supporting the assumptions of predictive coding, numerous findings in humans and animals indicate that attentional allocation and exploratory behaviors increase for stimuli or contexts associated with intermediate amounts of uncertainty, while they decrease when uncertainty becomes too high or too low[2,3,7,12]. For example, in trivia quizzes, participants report higher curiosity when confidence in their guesses is intermediate[18]. Yet, no study has explained how fluctuations in uncertainty dynamically control EC levels, nor has addressed the neural implementation of this homeostatic regulation, which is timely for curiosity research[19,20]. In the field of reinforcement-learning, the combination of elaborate decision-making tasks, computational modeling and neuroimaging recently showed the rostrolateral prefrontal cortex (rlPFC) monitors uncertainty and mediates its impact on exploration[21–23]. Since exploratory decisions are generally associated with heightened rlPFC responses[24] and that exogenous inhibition of rlPFC activity causally reduces the frequency of such decision[25], it would therefore be

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

logical for the suppression of EC by average surprise (as hypothesized in our model) to depend on a suppression of rlPFC activity in response unexpected pieces of knowledge.

Unfortunately, while surprise can be readily manipulated and framed as an information-theoretic or Bayesian quantity in decision-making tasks involving monetary reinforcers, surprise in the context of epistemic curiosity research depends on high-level representations about the world, which are themselves dependent upon an open-ended prior knowledge shaped by language. Although the hippocampus seem to play an important role by acting as a mismatch detector[26,27] when clear-cut expectations can be computed (e.g based the training phase of associative learning tasks), the number of possible alternatives is often too wide to be represented *a priori* in more ecological situations such as trivia quizzes. Therefore, contrary to the well-known mechanisms underlying reward prediction errors and despite its putative importance for curiosity-driven memory encoding[28,29], the neurobiological origins of epistemic surprise remain virtually unknown when computed *a posteriori*. Nevertheless, because of its role in schema-based memory[30–32] and its intense connectivity with the hippocampus[33], the ventromedial prefrontal cortex (vmPFC) may appear as a good candidate for deriving such signals based on an *ex post* comparison of new information with the content of episodic or semantic memory.

Regarding curiosity relief, electrophysiological studies demonstrated that the satisfaction of rudimentary curiosity states recruits dopamine neurons in the brainstem, hence explaining why non-human primates are willing to endure costs in order to obtain advance information about upcoming rewards[19,34]. Given that striatal BOLD responses are themselves modulated by dopamine[35], it was tempting to speculate that the relief of EC would recruit this key reward-related area in humans. Consequently, two fMRI studies based on trivia quizzes suggested that the striatum encodes curiosity states at the question stage, in its dorsal[18] or ventral parts[36]. Yet, these studies reported no modulation of the ventral striatum or other reward-related areas at the time of answers — which were systematically delivered — whereas such activations were observed in a perceptual curiosity paradigm[37] — in which curiosity was relieved in 50% of the trials. Since this dependency to uncertainty is highly consistent with the principle of temporal-difference (TD) learning[17], we hypothesized that the striatal encoding of curiosity relief might be observed when trivia answers are delivered in a stochastic rather than systematic manner (Fig. 1c).

In order to provide behavioral evidence for this integrative model of EC and to test our hypotheses regarding its neural implementation, we used a two-step trivia quiz (Fig. 2) designed to induce curiosity and manipulate answer uncertainty in 22 participants undergoing fMRI. Our quiz

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

focused on cinema because of the widespread interest in this domain across sexes, cultures, and education levels. This choice also facilitated the standardization of answers (which were always movie titles) and the evaluation of prior knowledge related to them (titles watched or not by the participant). During the first part of the quiz (run 1), participants rated their curiosity for 60 cinema-related questions. After each rating, the answer to the question was either revealed (50%) or replaced by hashtags (50%), hence generating uncertainty regarding the relief of curiosity. In the second part (run 2), the same 60 questions were presented again and participants were asked to indicate whether they remembered the answer or not. At this point, questions that had not been answered in run 1 could still elicit curiosity and their associated (new) answers could still elicit surprise, whereas remembered items served as controls, matched with the former in terms of visual stimulation and epistemic content thanks to counterbalancing. After the main task, a localizer involving individualized sets of new movie titles was used to reveal the brain regions responding to prior knowledge in a task-independent fashion (run 3). Unannounced post-test questionnaires were finally administered outside the scanner, including a recall test as well as surprise and interest ratings for each trivia item.

Results

Interplay of prior knowledge, curiosity and surprise for memory encoding

Behavioral analyses demonstrated that possessing some prior knowledge related to the answers increased both curiosity ratings ($t(21)=2.46$, $p=0.023$) and surprise ratings ($t(21)=4.86$, $p<0.001$)(Fig. 3a). Moreover, curiosity and surprise were positively correlated to each other ($r=0.22$, $p<0.001$, ratings z-scored for each participant individually; Fig. 3b) and a repeated-measures ANOVA confirmed that curiosity predicted surprise ratings ($F(2,42)=27.4$, $p<0.001$). Moreover, curiosity, surprise and prior knowledge were all positively associated with recall performances (whether the response was correct or not in the post-scan memory test) according to median-split analyses separating items as a function of high and low curiosity ($z=3.88$, $p<0.001$), surprise ($z=3.32$, $p<0.001$) and prior knowledge ($z=3.81$, $p<0.001$; Fig. 3c).

In order to exclude the possibility that prior knowledge, curiosity and surprise would simply reflect a common latent variable (e.g. attention), we used a Generalized Estimating Equations approach (see Methods, Behavioral analyses). The three factors were individually significant in this analysis (curiosity: beta=1.27, Wald $\chi2 = 15.9$, $p<0.001$; surprise: beta=0.43, Wald $\chi2 = 10.2$, $p=0.001$; prior

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

knowledge: beta = 1.58, Wald $\chi2$ = 74.4, p<0.001), hence confirming their additive contribution to memory encoding.

While the relationship between curiosity and surprise as well as the positive impact of curiosity, prior knowledge and surprise on recall were expected based on studies by Berlyne[28] and others[20,29,36,38], a logistic mediation analysis including prior knowledge and condition (whether the answer had been seen once or twice during the experiment) further showed that surprise partly mediated the beneficial effects of curiosity on recall performances (indirect path: z=3.47, p<0.001; direct path: z=3.64, p<0.001; Fig. 3d; Supplementary Fig. 1a). Finally, items rated as more interesting were associated with higher surprise ratings (t(21)=2.48, p=0.02). However, although interest was also strongly related to memory performance (direct path: z=3.72, p<0.001) as expected from previous findings[38], it did not mediate curiosity-driven memory benefits (indirect path: z=1.31, p=0.21).

Surprise-dependent control of curiosity

Our main hypothesis concerned the variation of nonspecific curiosity levels over time, as a function of the average amount of surprise recently experienced. In the computational approach used to tackle this issue (see Methods), we assumed that the subjective level of curiosity reported in participants' ratings (run 1) resulted from two distinct influences: (i) the motivation to relieve an acute ignorance state induced by the specific content of the question presented in any trial *t* (specific EC); (ii) the motivation to be exposed to any new information (nonspecific EC) conceived as an item-independent variable fluctuating slowly throughout the quiz. Importantly, the Rescorla-Wagner algorithm used to monitor the average amount of surprise was totally blind to the content of the questions and to the outcome of a trial *t*: consequently, it could only explain the variance associated with the nonspecific component of EC, based on previous items (*t-1*, *t-2*, etc).

Supporting our hypothesis, a delta-rule that updated the average amount of surprise experienced in the quiz (Q{sur}) outperformed a model that updated only the probability of knowledge delivery (Q{0-1}) and models that included time as a regressor, either alone or in combination with any of the two delta-rules (Fig. 4a). Bayesian group comparisons treating model attribution as a random effect indicated that this conclusion held both when curiosity ratings were considered as a normally distributed variable or when they were binarized into high/low categories and predicted by means of logistic regression (see Methods and Supplementary Fig. 1b-d for details). Crucially, the overall effect of expected surprise on curiosity ratings was negative in both the continuous (t(21)=-2.95, p=0.008) and

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

binomial cases (t(21)=-3, p=0.007), which confirmed that this variable exerted significant suppressive effects on EC ratings.

Monitoring of average surprise in the rostrolateral PFC

Next, we studied how the brain tracked changes in average surprise from one trial to another. To do so, we investigated the parametric encoding of surprise prediction errors (PE{sur}) at the answer stage of run 1, during which variations of curiosity levels were assessed. Formally, this trial-wise variable corresponds to the surprise experienced in each trial minus the average surprise recently experienced. Restricted to a prefrontal mask spanning all voxels anterior to the head of the caudate (MNI: Y>22, see Methods), our analyses revealed a significant encoding of PE{sur} within the right rostrolateral prefrontal cortex (Fig. 4b). The GLM excluded potential confounding effects of displaying answer *versus* hashtag (modeled as separate events), curiosity relief and prior knowledge (both included before PE{sur} in the serial orthogonalization procedure implemented by SPM, see Methods). Since parameter estimates were negative, we used a Finite Impulse Response model distinguishing outcomes (answer and hashtag) as a function of PE{sur}. It confirmed that stronger surprise prediction errors triggered proportional deactivations of the rlPFC (Fig. 4c). Additional ROI analyses demonstrated that this rlPFC area also encoded positively Q{sur}, with higher average surprise values associated with stronger rlPFC responses at the question stage (t(21)=3.03, p=0.006). Finally, replacing PE{sur} by surprise ratings in the GLM (keeping every other aspect equal) demonstrated that the rlPFC not only encoded Q{sur} and PE{sur} but also surprise itself (t(21)=-2.85, p=0.009).

Given that the cluster reported in Fig. 4b appeared more anterior and medial than expected based on the literature on uncertainty-driven exploration, we also assessed the effect of Q{sur} and PE{sur} at previously reported locations[21–24]. Interestingly, a distinct activation pattern was observed at 3 of these locations where Q{sur} was negatively encoded (Fig. 4e; peak from Daw et al: t(21)=-4.07, p<0.001; peak from Boorman et al: t(21)=-3.52, p=0.002; peak from Donoso et al: t(21)=-2.81,p=0.010) and where no significant modulations by PE{sur} emerged at the group level. However, still at these 3 locations, the neural encoding of PE{sur} predicted the dynamical influence of Q{sur} on curiosity ratings from one subject to another (Fig. 4f; peak from Daw et al: ρ=0.61, p=0.003; Boorman et al: ρ=0.69, p<0.001; Donoso et al., ρ=0.53, p=0.014), while it was not the case at the peak reported in Badre et al. (ρ=0.24, p=0.29) or at the peak reported Fig. 4b (ρ=-0.09, p=0.68). In other words, a more posterior and lateral portion of the rlPFC seemed to implement the control of nonspecific EC levels based on computations performed in the more medial and anterior site highlighted by our initial analysis.

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Genesis of epistemic surprise in the medial PFC

The vmPFC was among the various brain areas that discriminated strongly and reliably new answers from hashtag in the first run, as revealed by whole-brain analyses (Fig. 5a). The comparison of new answers and old answers in the second run produced a very similar pattern of activations which again included the hippocampus, the superior temporal sulcus (STS), the dlPFC and the precuneus (Fig. S2b; Table S3). Since vmPFC activations regularly co-occurred with activations in these structures and others, we systematically investigated the selectivity profile of the vmPFC together with 7 other regions of interest (Fig. 5b; see Methods for details), including the dorsomedial PFC (dmPFC) and the inferior parietal lobe (IPL) found to encode curiosity at the question stage (Fig. S2a), as well as the ventral striatum involved in curiosity relief (see Fig. 6a).

This multiple ROI analysis indicated that the vmPFC had an ideal selectivity profile to compute epistemic surprise signals. First, together with the precuneus, it discriminated genuinely new answers from old but forgotten ones (vmPFC: $t(21)=3.74$, $p=0.001$; precuneus: $t(21)=3.94$, $p<0.001$, Fig. 5c). Second, the vmPFC and to a lesser extent the precuneus were the only regions sensitive to the prior knowledge variable (watched or not watched), during the trivia quiz (vmPFC: ($t(21)=2.82$, $p=0.010$; Fig. 5d) and the localizer task (vmPFC: $t(21)=3.14$, $p=0.005$; Fig. 5e; Table S4). Third, along with the dmPFC and to a lesser extent the STS, the vmPFC encoded epistemic surprise itself ($t(21)=3.47$, $p=0.002$; dmPFC: $z=2.88$, $p=0.005$; STS: $t(21)=2.52$, $p=0.019$; Fig. 5f).

In line with the influence of epistemic surprise on memory encoding, higher vmPFC activity also predicted subsequent recall of trivia answers in the post-scan memory test ($t(21)=2.21$, $t=0.03$; Fig. S2c). Although this effect did not survive the false-discovery rate criterion used to correct for multiple comparisons across ROIs, this finding remains remarkable since the regressor indexing "subsequent recall" was orthogonalized on curiosity, prior knowledge and PE{sur}. Finally, we observed that the encoding of epistemic surprise in the vmPFC correlated with how the model-based variable Q{sur} impacted nonspecific EC levels across participants ($r=0.45$, $p=0.039$) — a result consistent with the idea that the vmPFC could forward the results of its computations to the rlPFC.

Curiosity relief and tip-of-the-tongue events in the ventral striatum

Epistemic surprise and its temporal integration into an average surprise signals are crucial for curiosity-driven memory benefits and for the regulation of nonspecific EC levels, respectively. However, our model also stresses that information-seeking behaviors are proximally motivated by the

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

relief of ignorance regarding specific information gaps. Thus, we first investigated curiosity relief in the first run of our task, which delivered answers in 50% of the trials. A whole-brain analysis showed that the ventral striatum was the only region that responded parametrically to curiosity at the answer stage (Fig. 6a). Importantly, this conclusion held when restricted to an anatomical mask of the nucleus accumbens (NAcc; t(21)=2.54, p=0.02; Fig. 6b), where no significant modulation by curiosity occurred when hashtags were displayed (t=0.03, p=0.97), hence excluding the possibility that curiosity-dependent modulation originated from the question or waiting stages (as reported by Gruber and colleagues[36]). Finally, in line with previous null findings[18,36], no modulation of ventral striatal or NAcc activities were detected in the second run, in which curiosity was relieved in 100% of the trials (striatal ROI from Fig. 6a: t(21)=0.51, p=0.6; NAcc: t(21)=-1.34, p=0.19; Fig. 6b).

Subsequent analyses revealed that activations of the ventral striatum were not only elicited by stochastic curiosity relief: together with the hippocampus (see Fig. S2e-g), it was more activated by questions whose answers were known before the experiment (or remembered from run 1), as compared to question whose answers were unknown (run 1: z=3.25, p=0.001, three participants ignorant of all answers excluded; run 2: z=2.77, p=0.006; Table S5). Intriguingly, the ventral striatum was also more activated by questions whose answers were presumably on the “tip of the tongue” (answered in run 1 but reported as forgotten in run 2) than for questions with unknown answers (t(21)=3.13, p=0.005: Fig. 6c), which indicates that this activation did not depend on successful retrieval contrary what was observed in the hippocampus (Fig. S2). Finally, among the eight ROIs discussed above (see Fig. 5b), the ventral striatum and the hippocampus were the only structures correlating negatively with Q{sur}, the model-based variable tracking average surprise (ventral striatum: t(21)=-3.04, p=0.006, hippocampus: t(21)=-2.94, p=0.008; Fig. S2h).

Discussion

Taken together, our behavioral and neuroimaging results provide evidence for an integrative model of epistemic curiosity (EC) which dissociates several cognitive processes simultaneously triggered by the reception of new information (Fig. 1): (i) the relief of curiosity itself, encoded in the ventral striatum when knowledge delivery is stochastic rather than systematic; (ii) the comparison of new facts with prior knowledge, resulting in an epistemic surprise signal possibly generated by the vmPFC and facilitating memory encoding; (iii) the update of an average surprise variable computed by the rlPFC and involved in the dynamical control of nonspecific EC levels. These findings demonstrate

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

the complementary roles played by the prefrontal cortex and the ventral striatum in the neural valuation of knowledge and the regulation of EC. Moreover, they validate a non-trivial assumption of the predictive coding framework regarding the relationship between experienced surprise and subsequent curiosity levels.

First, our data indicates that curiosity-related activities elicited in the ventral striatum occurs only when knowledge delivery is stochastic, which supports our hypothesis that EC relief engages computations related to temporal-difference (TD) learning. Indeed, one prediction of TD learning is that the affective value of any awaited outcomes is represented in the reward circuitry as soon as it can be anticipated[39]. Thus, in contexts where questions are systematically answered, the motivational or affective signaling associated with EC should mostly occur at the question stage, as suggested by previous studies[18,36]. On the contrary, when answers are uncertain, higher EC levels can either translate into more interesting or more frustrating outcomes — depending on whether awaited information is delivered or not — so that outcome value cannot be anticipated. This latter context corresponds exactly to the first run of our trivia task delivering answers in only 50% of the trials and where ventral striatal signals proportional to EC relief were actually observed. This finding provides an important cog in the motivational machinery of information-seeking behaviors as it may help to reconcile two opposing views regarding the affective valence of acute curiosity states. Indeed, EC is sometimes envisioned as an appetitive state during which "epistemic rewards" are anticipated[18,36] or — instead — as an aversive state of ignorance to be relieved through information-seeking and exploratory behaviors[7,40]. From a TD-learning perspective, these two views may actually apply to different contexts, with appetitive processes being predominant when answer uncertainty is null or low (Fig. 1c, right panel) and aversive processes being predominant when it is high (Fig. 1c, left panel). In addition, our data suggests that the ventral striatum may encode the affective or motivational value of situations in which one has the feeling to know and/or anticipates the confirmation of an expected answer, as it was clearly more activated — in the second run — by question which had previously been answered as compared to unanswered ones. Yet, an alternative interpretation may be that ventral striatal activities at the question stage reflect the energization of goal-directed memory retrieval. Consistent with its hypothesized role in the motivational and cognitive control of mnemonic processes[41], this view would explain why the striatum was also activated in front of questions for which answers remained the tip of the tongue, contrary to the hippocampus.

Second, our model emphasizes the pivotal importance of surprise signals, which go beyond the motivational processes engaged by curiosity relief. Indeed, if answers usually reduces the uncertainty associated with their corresponding questions, they can also be incongruent with prior knowledge, hence

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

triggering surprise. Surprise was positively correlated with curiosity ratings. This may be due to the fact that curiosity tends to increase attention[2], which in turns modulate the amplitude surprise-related neural activities such as mismatch responses and prediction errors[42,43]. Moreover, consistent with the hypothesis that surprise gates the update of internal representations (i.e. encoding) like prediction errors do in reinforcement-learning or predictive coding[44], it predicted subsequent memory recall and partly mediated curiosity-driven memory benefits. At the neural level, more surprising answers elicited stronger vmPFC responses, likewise answers associated with more prior knowledge and answers subsequently recalled in the postscan memory test. Therefore, the vmPFC seems ideally suited to compute epistemic surprise and to control memory encoding, as observed in our behavioral data. Yet, no definitive conclusions can be made at this stage because the precedence of surprise signaling in vmPFC compared to the dmPFC — which also encoded surprise — cannot be firmly assessed (given the limited temporal resolution of fMRI). Alternatively, the vmPFC may integrate inputs originating from different brain regions representing different variables (surprise, novelty, prior knowledge, etc.) into a global value signal (i.e. interest). A last possibility would be that different population of vmPFC neurons compute these surprise and value signals concurrently. Indeed, recent fMRI experiments and electrophysiological recordings showed that vmPFC activities can represent (i) knowledge value and confidence about knowledge[45] or (ii) reward valence and advanced information about rewards[46] in a multiplexed manner.

Third, using a simple delta-rule to track the average amount of surprise experienced during the trivia quiz, we provide evidence for a key assumption of the predictive coding framework regarding the regulation of EC. Namely, we demonstrated that this average surprise variable robustly suppressed nonspecific curiosity levels, which fluctuated over time independently of questions' contents. This observation suggests that individuals might indeed regulate their level of curiosity and associated exploratory behaviors so that experienced surprise aligns with their expectations. The rlPFC appeared to implement this key process. Indeed, it encoded the amount of surprise recently experienced when participants processed new questions and surprise prediction errors when participants were presented with new answers. Importantly, stronger surprise prediction errors were associated with stronger deactivations of the rlPFC structure in response to more surprising answers. Although the functional meaning of negative BOLD responses constitutes a debated topic in the neuroimaging community[47,48] due notably to the paucity of papers reporting peristimulus time-course histograms, the activation pattern observed here is compatible with previous studies. For example, clear-cut deactivations of this prefrontal area have been documented in the context of semantic judgments[49] and in reinforcement-learning[24] tasks, where they have been associated with exploitative rather than explorative decisions. Yet, the representation of average surprise in the rlPFC proved to be highly complex. First, its encoding

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

was positive or negative depending on the exact spatial coordinates under scrutiny. Second, the analysis of interindividual differences revealed a positive correlation between the suppression of curiosity levels by surprise and the neural encoding of surprise prediction errors, but only in the posterior and lateral portion of the rlPFC cortex. This suggest a possible dissociation — within the rlPFC — of the neural populations monitoring average surprise and those implementing the influence of this variable over curiosity-related behaviors. Accordingly, the study of interindividual differences in a reinforcement-learning task recent suggested that the rlPFC area monitoring uncertainty may not be the exactly same as the rlPFC area implementing the influence of uncertainty on exploratory decision[21].

Taken together, our findings support a neurocognitive model reconciling several processes involved in the upstream causes and the downstream consequences of epistemic curiosity. This model incorporates elements from the reinforcement-learning and the predictive coding literatures in order to articulate the "affective" and "cognitive" dimensions of EC, its relief by information and its dynamic regulation over time, as well as its tight relationship with memory encoding. Yet, further research is warranted to overcome some limitations of our study. First, it would be important to confirm the suppressive effect of average surprise in other tasks probing the willingness-to-pay or the willingness-to-wait for answers[18,50] and to cross-validate surprise ratings using related measures such as pupil dilation[51] or eye movements[20]. Second, our assessment of prior knowledge was confounded by the fact that people usually watch movies following willful, value-based choices. Although people naturally tend to value more cultural or scientific domains for which they are more knowledgeable (and vice-versa), this confound could be excluded by inducing — rather than merely assessing — prior knowledge in future experiments. Achieving more stringent control over prior knowledge will be key to confirm that introspective ratings of epistemic surprise indeed reflect a mismatch between the content of episodic or semantic memory and new information, as suggested here. Third, assessing the presence and the accuracy of expectations when processing trivia questions may help to disentangle the neural circuits signaling (i) surprise as a violation of active *a priori* expectations, and (ii) surprise as an incongruence with episodic or semantic representations retrieved *a posteriori*.

To conclude, understanding the regulation of EC and its neural implementation in the human brain will require intense research efforts in domains as diverse as memory, attention, linguistics and decision-making. This endeavour should build on the complementary insights provided by the reinforcement-learning and predictive coding frameworks, as well as other information-seeking principles such as learning progress maximization[2]. Taking into account the dynamical relationship between surprise and curiosity may help integrate these diverse literatures and open the path to more

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

autonomous systems in artificial intelligence and to new strategies of knowledge transfer in the classroom.

Online Methods

Participants

Twenty-two right-handed students (11 females, 11 males; mean age: 22.9; range: 19-28) were recruited through advertisements in an art cinema and via university mailing lists. This sample size matched the range of existing neuroimaging studies on epistemic curiosity[18,36]. No participant was excluded from data analyses. All were paid at the fixed rate of 60€ for their participation in the study. A few days before the experimental session, participants signed informed consent after exhaustive explanations were provided. They were also given a list of 215 movie titles and asked to indicate for each of them to what extent they knew the movie (from 1= *never heard of it* to 4 = *seen it several times*). Target movie titles were covertly included in this list, which enabled us to quantify prior knowledge about trivia items (i.e. watched/unwatched status). For the behavioral experiment performed to select and validate the trivia used in the fMRI study, 64 participants of all ages were invited to complete a computerized evaluation of candidate trivia items (Supplementary Fig. 3a) after filing a consent form. The behavioral experiment took place in an art cinema (Comoedia, Lyon). After completing the task (about 20 minutes), they were offered to pick a book among a large selection of novels and essays. The entire protocol was approved by the local ethics committee of Sud-Est II, France (authorization number: 2011-056-2).

Stimuli

Sixty question-answer pairs were selected amongst the 120 pre-screened trivia items (Supplementary Fig. 3a and Table S1). The trivia questions included in the fMRI experiment were chosen to maximize reported surprise, interest and knowledge about the target movies (Fig S3b). In addition, items were selected and designed to minimize the chances that participants would know or guess the answers. In the fMRI experiment, the frequency of known answers was therefore very low (5.5±6.1%: range: 0-18%) and known items were always modeled separately and excluded from all analyses (except for Fig. 5b). Moreover, we ensured that items associated with the two main conditions (i.e. items answered or not in the first run, see Fig. 1) were highly matched for characters count (for

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

both questions and answers), curiosity, surprise and interest (all p>0.75). Finally, we counterbalanced across participants the subsets of items associated with each condition.

For the prior knowledge localizer task, we used the 215-items questionnaire to create two personalized sets of movie titles, different from those encountered in the main trivia task: 30 watched movies (if possible, watched less than two years before the experimental session) and 30 unwatched movies (if possible, with titles known). For two participants who had not seen enough movies in the list, we included movies seen more than two years before within the pool of watched movies (10 and 26 items, respectively).

Time course of the fMRI experiment

At their arrival to the MRI lab, participants were reminded that they would be exposed to cinema-related trivia questions and warned that those questions had been selected for being interesting but rarely known, even to cinema lovers. Once in the scanner, they completed sixteen training items to improve self-calibration in curiosity ratings (those items were not redundant with those of the main task) after receiving the following instructions (hereafter translated from French):

You are about to begin an experiment about intellectual curiosity and cinema in the specific conditions of the MRI scanner.

[Slide 1] During the calibration of the scanner and the acquisition of the anatomical image of our brain, you will practice the task that you will be doing while we will record your cerebral activities.

This training must in particular enable you to manipulate properly the response gauge with which you will indicate to what extent you are curious to know the answers to the questions we are going to present you.

[Slide 2, showing a fixation cross] Each trial will begin with a small symbol signaling that a question is about to appear.

[Slide 3, showing a dummy trivia question] After a few seconds, the question will appear. Take the time to read it properly. Once you have read it, press the left button (index).

[Slide 4, showing the gauge and the question] Once you press the button, the curiosity gauge will appear. By keeping the left button pressed, you can increase the gauge up to how much you are curious.

[Slide 5, showing the gauge and the question] If you are sure to know the answer, press the right button (major finger) when the answer comes to your mind. NB: if you think you know the answer but don't remember it (answer on the tip of the tongue), don't answer with the right finger but indicate your curiosity level.

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

[Slide 6] Once you have raised the gauge up to the level corresponding to your curiosity, a fixation cross will appear on the screen.

[Slide 7, showing the answer to the dummy question] Finally, the answer to the question will be displayed during a few seconds. However, during the first part of the experiment, we will only delivered 50% of the answers. NB: there is NO relationship between your curiosity rating and the likelihood of of receiving or not the answer.

In the first functional run, each trial started with a jittered fixation cross (exponential distribution; mean: 4.2s: range: 3-7.5 seconds). Then, participants had to read one of the 60 pre-screened trivia questions and to signal end of reading with a button press (right index finger; average reading time: 5.1±1.49s). After a fixed interval of 750ms, a continuous gauge appeared. Participants had then to use their index finger to rate their curiosity by keeping the left button pressed until the gauge reach the desired point (maximum curiosity: 2.5s). In case they would know the answer already, they were instructed to answer with the right finger and then had to wait for 2s. Another jittered fixation cross (exponential distribution; mean: 4.2s: range: 3-7.5s) preceded the delivery of either an answer (50% of the trials) or hash tags "#" (3s, fixed duration). The temporal order of items was randomized for each participant independently.

In the second functional run, participants were verbally instructed that they would be presented again with all the questions, and that this time they would simply have to indicate whether the correct answer came spontaneously to their mind or not (average response time: 3.7±0.83s). To do so, they had to select either a "light bulb" or a "cloud" associated with each situation, respectively (black and white drawings of similar size displayed on the left and right of the question; side counterbalanced across trials). All questions were again preceded and followed by a fixation cross (exponential distribution; mean: 4.2s: range: 3-7.5s). In this second run, answers were delivered in all trials (3s, fixed duration). The temporal order of items was re-randomized for each participant independently.

In the third functional run, participants were presented with 30 watched movie titles, 30 unwatched movie titles, and 30 hashtags "#". Each trial began with a fixation cross (mean: 2.5s: range: 2-6s). Then a target was appeared on the screen for a fixed duration (3s) together with two dots, associated with the mentions "seen" and "unseen" (on the left and right of the movie title) or "skip" (on both side, in case of hash tags). The side of "seen" and "unseen" mentions was counterbalanced across trials and the temporal order of items was randomized for each participant independently.

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Once outside the scanner, participants were first presented with an unexpected memory test in which they had to write down the answer of the 60 trivia questions encountered in the task. At this stage, they also reported which answers they were expecting (13.2±8.8%) or knew already for sure (4.4±5.1%) before the task. Then, all questions and answers were shown together, and participants were asked to rate their surprise levels (from 1 "not at all" to 5 "yes, a lot") and to report the thirty items they found the most interesting. To conclude, they filled an epistemic curiosity questionnaire [6] designed to capture specific (i.e deprivation) and diversive (i.e interest) EC. All behavioral tasks were programmed using Presentation (www.neurobs.com).

fMRI acquisition

Imaging was conducted on a Siemens Sonata scanner (1.5T), using an eight-channel head coil. Twenty six interleaved slices tilted relative to the anterior commissure – posterior commissure line (20-30°) were acquired per volume. We acquired an average of 837 echo-planar T2*-weighted functional volumes per subject (TR = 2.5; TE = 60 ms; FOV = 220 mm; matrix = 64 x 64; voxel size = 3.4 x 3.4 x 4mm). Following the fMRI session, a high-resolution T1-weighted anatomical scan was acquired. Before the functional acquisition, a gradient-field map was acquired using a gradient echo sequence and was applied for distortion-correction of the acquired functional images in order to improve local field homogeneity and minimize susceptibility artifacts, for example in the ventral parts of the prefrontal cortex.

fMRI preprocessing

All preprocessing steps were performed using SPM8. The first four volumes of each run were removed to allow for T1 equilibrium effects. For each participant, functional images were time-corrected, realigned, unwarped using the magnitude and phase images, and coregistered to the anatomical scan. The six movement parameters were derived from the iterative realignment procedure carried out by SPM8 (three for translation, three for rotation). The anatomical scan was then normalized to the MNI space using the ICBM152 template brain and the resulting non-linear transformation matrix was applied to the functional images. Finally, the normalized functional images were spatially smoothed with an 8 mm Gaussian kernel.

fMRI analyses

Statistical analyses of fMRI signals were performed using a conventional two-levels random-effects approach with SPM8. All general linear models (GLM) described below included the 6

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

unconvolved motion parameters from the realignment step, in order to covary out potential movement-related artifacts in the BOLD signal. All regressors of interest were convolved with the canonical hemodynamic response function (HRF). All GLM models included a high-pass filter to remove low-frequency artifacts from the data (cut-off = 128s) as well as a run-specific intercept. Temporal autocorrelation was modeled using an AR(1) process. All motor responses recorded were modeled using a zero-duration Dirac function. Voxel-wise thresholds used to generate SPM maps were either $p<0.005^{UNC}$ (parametric contrasts) or $p<0.001^{UNC}$ (categorical contrasts), unless notified otherwise. All statistical inferences based on whole-brain analyses satisfied the standard multiple comparison threshold ($p<0.05^{FWE}$) at the cluster level.

In the first run (GLM1), the question, rating and outcome stages were modeled separately using boxcar functions set to the duration of each individual event. This decision to use boxcars was justified by an analysis of the residuals produced by the GLMs at the first level, compared with those from the homologous model using Dirac functions (difference in log-likelihood (LL) against homologous Dirac model: 271.7). Questions for which the participant did not know the answer were parametrically modulated by four regressors, orthogonalized in the following order:

1°) Q{sur}: value of the surprise accumulator (see "Behavioral analyses" section, below).

2°) Prior knowledge: 1 if target movie title had been watched by the participant, 0 otherwise.

3°) Curiosity: value from 0 (excluded) to 1 (maximum curiosity).

4°) Subsequent recall: 1 if item subsequently recalled, 0 otherwise.

At the outcome stage, answers and hash tags were also parametrically modulated using four regressors, orthogonalized in the following order:

1°) Curiosity.

2°) Prior knowledge.

3°) Surprise prediction error (PE{sur}) or Surprise (see below).

4°) Subsequent recall

Questions and answers for which participants knew the answer before starting the experiment were modeled separately and not included in any contrast, except for the contrast reported Fig. S2f. In order to uncover the neural correlates of surprise in the first run (i.e. only for Fig. 5f), surprise ratings were simply substituted to surprise prediction errors, keeping all other aspects of the analysis identical.

In the second run, questions and answers were both modeled using Dirac functions. Again, this decision was principled by the analysis of first-level residuals (difference in LL against homologous boxcar model: 23.0). We splitted questions and answers regressors as a function of their status in the

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

first run (i.e. items answered or not in run 1) and participants' ability to recall spontaneously the answer or not. This resulted in two "HIT" regressors (items previously answered and remembered, at the question and answer stages) and two "correct rejection" (CR) regressors (unanswered and correctly classified as such, also at both stages). Questions (HIT and CR) were parametrically modulated using 4 regressors, orthogonalized in the following order:

1°) Curiosity
2°) Prior Knowledge
3°) Surprise
4°) Subsequent recall

Answers (HIT and CR) were also modulated using 4 regressors, orthogonalized in the following order:

1°) Curiosity
2°) Prior Knowledge
3°) Surprise
4°) Subsequent recall

Items which had been answered in the first run but could not be spontaneously recalled by the participants were modeled separately (MISS regressors). Items which were already known before starting the experiment were also modeled separately and not included in any analysis.

In the third run, we modelled the onset of hashtags, watched movies and unwatched movies separately using zero-duration Dirac functions. Given the short duration of each trial, we lowered the cut-off of the high-pass filter (64s instead of 128s).

Concerning ROI analyses, the mask used to extract effects from the peaks previously reported in the literature study the contribution of the rlPFC to uncertainty-driven exploration were 3mm-radius spheres centered around the MNI coordinates reported in the original papers (explicitly displayed on Fig. 4b). For the multiple ROIs analyses reported Fig. 5b-f, Fig. 6c and Fig. S2b-h, we used the following method: (i) clusters surviving a voxel-wise threshold of $p<0.05$FWE were extracted from the [new answer>hashtag] contrast (run 1; dlPFC, vmPFC, HPC, STS, Precuneus), (ii) clusters surviving a cluster-wise threshold of $p<0.05$FWE (voxel-wise threshold: $p<0.005$unc) were extracted from the parametric curiosity contrasts at the question (dmPFC, IPL) and answer (ventral striatum) stages of run 1. For each of the 8 regions, the mirror (x-flipped) ROI was added to the mask itself, so that every ROIs were strictly symmetric and identical across the two hemispheres. Finally, the nucleus accumbens mask (Fig. 6b) was based on an anatomical probabilistic atlas of the basal ganglia[52].

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Peristimulus time-course histograms (PSTH, sampled at 1Hz) were computed using the toolbox *rfxplot* for Matlab [53]. These time-decomposed effects were thus re-estimated using the first eigenvariate extracted from the regions of interest, after adjustment for run intercept and movement-related variance.

Behavioral analyses

The modeling of nonspecific EC levels as a function of epistemic surprise used the following delta-rule:

$$Q_{t+1} = Q_t + \alpha(R - Q_t) \quad \text{(Equation 1)}$$

where Q is initialized at 0 and updated on each trial by the prediction error term R-Q, times a learning rate α. In the most simple model termed Q{0-1}, the delivery of an answer was coded as R=1 while the absence of answer was coded as R=0, so that the variable Q represents the amount of knowledge recently delivered to the participant, which enabled us to explore whether knowledge tended to reinforce or saturate curiosity over time. In the best-fitted model termed Q{sur},the delivery of an answer was coded as R=S while the absence of answer was coded as R=0, with S corresponding to the surprise rating given by the participant for that particular item.

In order to ascertain that this approach was useful to explain variance in curiosity ratings, we compared a range of alternative models using a Bayesian group comparison approach (Fig. 4a, Fig. S1b-e), as implemented in the toolbox VBA[24] for Matlab (http://mbb-team.github.io/VBA-toolbox/). Alternative models were: Q{0-1}, Q{sur}, time, Q{0-1} & time, Q{sur} & time. The "time" model was used to ascertain that our delta-rule was not merely capturing a linear (increasing or decreasing) trend in curiosity ratings but rather an information-dependent process. An intercept was included in all models. Learning rates were treated as a fixed-effect in order to limit model complexity and facilitate the interpretation of individual differences and correlates of surprise accumulation. Subject-level estimations were performed using the *fitglm* algorithm provided in Matlab. When the fit was performed on continuous curiosity ratings, we assumed that those were normally distributed. However, because this assumption was violated in 4 participants, we also check that the same results could be observed using binarized curiosity levels (ie. superior or inferior to 50%, corresponding to the half-maximum of the curiosity gauge).

To confirm the complementary contributions of prior knowledge, surprise and curiosity in facilitating recall performance, we performed a Generalized Estimating Equations analysis (GEE) analysis, as implemented by SPSS 21 (https://www.ibm.com/analytics/us/en/technology/spss/). Successful recall was coded as 1 and unsuccessful recall as 0 and predicted by mean of a logistic

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

regression. The analysis included a participant-specific intercept, trivia id as a within-participant effects, and curiosity, prior knowledge and surprise as random effects.

The multi-level mediation analysis (Fig. 2e) was performed using the Mediation toolbox[26] for Matlab (https://github.com/canlab/MediationToolbox). Curiosity and surprise levels were z-scored for each participant separately, after removing items which were known to the participants before the experiment (according to post-scan task and responses given in run 1). Mediation path coefficients were estimated for each participant independently. Statistical inferences were drawn at the group level for each coefficient using a bias-corrected bootstrap significance test relaxing the normality assumption (10 000 permutations). Averaged paths coefficient and standard deviations are reported directly on Fig. 2e. Individual path coefficients are reported on Fig S1a. The algorithm could not converge for 3 participants, which were excluded, due to their high percentage of correct responses (above 90%). The condition (answers repeated or not during the trivia task) and the presence of prior knowledge (watched or unwatched status of the target movie) were included as covariates of non-interest. The mediation analysis performed on data from the prescreen experiment (Supplementary Fig. 3c) was not logistic (as it applied to continuous interest ratings) and it included no covariates.

ACKNOWLEDGEMENTS

We are grateful to the staff of CERMEP–Imagerie du Vivant for helpful assistance with data collection and to the movie theater Comoedia (Lyon, France) for offering us the opportunity to perform the behavioral experiment in their premises. We must also thank Roshan Cools and three anonymous reviewers for their helpful comments on an earlier version of this manuscript.

AUTHOR CONTRIBUTIONS

R.L and T.M designed and performed the experiments. R.L and T.M analyzed the behavioral data. R.L analyzed the fMRI data. R.L and T.M wrote the paper. M.M funded the experiments and participated in writing the paper.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

References

1 Von Stumm S, Hell B, Chamorro-Premuzic T. The hungry mind intellectual curiosity is the third pillar of academic performance. *Perspect Psychol Sci* 2011;**6**:574–88.

2 Gottlieb J, Oudeyer P-Y, Lopes M, Baranes A. Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends Cogn Sci* 2013;**17**:585–93.

3 Kidd C, Hayden BY. The Psychology and Neuroscience of Curiosity. *Neuron* 2015;**88**:449–60.

4 Aristotle. *Metaphysics*. Ann Arbor, MI: University of Michigan Press; 1952.

5 Berlyne DE. A theory of human curiosity. *Br J Psychol* 1954;**45**:180–91.

6 Litman JA, Spielberger CD. Measuring epistemic curiosity and its diversive and specific components. *J Pers Assess* 2003;**80**:75–86.

7 Loewenstein G. The psychology of curiosity: A review and reinterpretation. *Psychol Bull* 1994;**116**:75.

8 Friston K, Thornton C, Clark A. Free-energy minimization and the dark-room problem. *Front Psychol* 2012;**3**:130.

9 Clark A. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav Brain Sci* 2013;**36**:181–204.

10 Friston K, Rigoli F, Ognibene D, Mathys C, Fitzgerald T, Pezzulo G. Active inference and epistemic value. *Cogn Neurosci* 2015;**6**:187–214.

11 Schwartenbeck P, FitzGerald T, Dolan R, Friston K. Exploration, novelty, surprise, and free energy minimization. *Front Psychol* 2013;**4**:710.

12 Berlyne DE. Curiosity and exploration. *Science* 1966;**153**:25–33.

13 Litman JA. Relationships between measures of I- and D-type curiosity, ambiguity tolerance, and need for closure: An initial test of the wanting-liking model of information-seeking. *Pers Individ Dif* 2010;**48**:397–402.

14 Bateman TS, Hess AM. Different personal propensities among scientists relate to deeper vs. broader knowledge contributions. *Proc Natl Acad Sci U S A* 2015;**112**:3653–8.

15 Litman JA. Interest and deprivation factors of epistemic curiosity. *Pers Individ Dif* 2008/5;**44**:1585–95.

16 Hirsh JB, Mar RA, Peterson JB. Psychological entropy: a framework for understanding uncertainty-related anxiety. *Psychol Rev* 2012;**119**:304–20.

17 Seymour B, O'Doherty JP, Koltzenburg M, Wiech K, Frackowiak R, Friston K, *et al.* Opponent appetitive-aversive neural processes underlie predictive learning of pain relief. *Nat Neurosci* 2005;**8**:1234–40.

18 Kang MJ, Hsu M, Krajbich IM, Loewenstein G, McClure SM, Wang JT-Y, *et al.* The wick in the candle of learning: epistemic curiosity activates reward circuitry and enhances memory. *Psychol Sci* 2009;**20**:963–73.

19 Daddaoua N, Lopes M, Gottlieb J. Intrinsically motivated oculomotor exploration guided by uncertainty reduction and conditioned reinforcement in non-human primates. *Sci Rep* 2016;**6**:20202.

20 Baranes A, Oudeyer P-Y, Gottlieb J. Eye movements reveal epistemic curiosity in human observers. *Vision Res* 2015;**117**:81–90.

21 Badre D, Doll BB, Long NM, Frank MJ. Rostrolateral prefrontal cortex and individual differences in uncertainty-driven exploration. *Neuron* 2012;**73**:595–607.

22 Donoso M, Collins AGE, Koechlin E. Human cognition. Foundations of human reasoning in the prefrontal cortex. *Science* 2014;**344**:1481–6.

23 Boorman ED, Behrens TEJ, Woolrich MW, Rushworth MFS. How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. *Neuron* 2009;**62**:733–43.

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

24 Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ. Cortical substrates for exploratory decisions in humans. *Nature* 2006;**441**:876–9.
25 Raja Beharelle A, Polanía R, Hare TA, Ruff CC. Transcranial Stimulation over Frontopolar Cortex Elucidates the Choice Attributes and Neural Mechanisms Used to Resolve Exploration–Exploitation Trade-Offs. *J Neurosci* 2015;**35**:14544–56.
26 Kumaran D, Maguire EA. Novelty signals: a window into hippocampal information processing. *Trends Cogn Sci* 2009;**13**:47–54.
27 Duncan K, Ketz N, Inati SJ, Davachi L. Evidence for area CA1 as a match/mismatch detector: a high-resolution fMRI study of the human hippocampus. *Hippocampus* 2012;**22**:389–98.
28 Berlyne DE. Uncertainty and epistemic curiosity. *Br J Psychol* 1962;**53**:27–34.
29 Stahl AE, Feigenson L. Observing the unexpected enhances infants' learning and exploration. *Science* 2015.
30 Ghosh VE, Moscovitch M, Melo Colella B, Gilboa A. Schema representation in patients with ventromedial PFC lesions. *J Neurosci* 2014;**34**:12057–70.
31 van Kesteren MTR, Beul SF, Takashima A, Henson RN, Ruiter DJ, Fernández G. Differential roles for medial prefrontal and medial temporal cortices in schema-dependent encoding: from congruent to incongruent. *Neuropsychologia* 2013;**51**:2352–9.
32 Garrido MI, Barnes GR, Kumaran D, Maguire EA, Dolan RJ. Ventromedial prefrontal cortex drives hippocampal theta oscillations induced by mismatch computations. *Neuroimage* 2015;**120**:362–70.
33 van Kesteren MTR, Fernández G, Norris DG, Hermans EJ. Persistent schema-dependent hippocampal-neocortical connectivity during memory encoding and postencoding rest in humans. *Proceedings of the National Academy of Sciences* 2010;**107**:7550–5.
34 Bromberg-Martin ES, Hikosaka O. Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron* 2009;**63**:119–26.
35 Pessiglione M, Seymour B, Flandin G, Dolan RJ, Frith CD. Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature* 2006;**442**:1042–5.
36 Gruber MJ, Gelman BD, Ranganath C. States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron* 2014;**84**:486–96.
37 Jepma M, Verdonschot RG, van Steenbergen H, Rombouts SARB, Nieuwenhuis S. Neural mechanisms underlying the induction and relief of perceptual curiosity. *Front Behav Neurosci* 2012;**6**:5.
38 McGillivray S, Murayama K, Castel AD. Thirst for knowledge: The effects of curiosity and interest on memory in younger and older adults. *Psychol Aging* 2015;**30**:835–41.
39 O'Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ. Temporal difference models and reward-related learning in the human brain. *Neuron* 2003;**38**:329–37.
40 Berlyne DE. Curiosity and learning. *Motiv Emot* 1978;**2**:97–175.
41 Scimeca JM, Badre D. Striatal contributions to declarative memory retrieval. *Neuron* 2012;**75**:380–92.
42 Jiang J, Summerfield C, Egner T. Attention sharpens the distinction between expected and unexpected percepts in the visual brain. *J Neurosci* 2013;**33**:18438–47.
43 Auksztulewicz R, Friston K. Attentional Enhancement of Auditory Mismatch Responses: a DCM/MEG Study. *Cereb Cortex* 2015;**25**:4273–83.
44 Fernández RS, Boccia MM, Pedreira ME. The fate of memory: Reconsolidation and the case of Prediction Error. *Neurosci Biobehav Rev* 2016;**68**:423–41.
45 Lebreton M, Abitbol R, Daunizeau J, Pessiglione M. Automatic integration of confidence in the brain valuation signal. *Nat Neurosci* 2015;**18**:1159–67.
46 Blanchard TC, Hayden BY, Bromberg-Martin ES. Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron* 2015;**85**:602–14.
47 Shmuel A, Augath M, Oeltermann A, Logothetis NK. Negative functional MRI response correlates with decreases in neuronal activity in monkey visual area V1. *Nat Neurosci* 2006;**9**:569–77.

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

48 Weitz AJ, Fang Z, Lee HJ, Fisher RS, Smith WC, Choy M, *et al.* Optogenetic fMRI reveals distinct, frequency-dependent networks recruited by dorsal and intermediate hippocampus stimulations. *Neuroimage* 2015;**107**:229–41.

49 Hayama HR, Rugg MD. Right dorsolateral prefrontal cortex is engaged during post-retrieval processing of both episodic and semantic information. *Neuropsychologia* 2009;**47**:2409–16.

50 Marvin CB, Shohamy D. Curiosity and reward: Valence predicts choice and information prediction errors enhance learning. *J Exp Psychol Gen* 2016;**145**:266–72.

51 Preuschoff K, 't Hart BM, Einhäuser W. Pupil Dilation Signals Surprise: Evidence for Noradrenaline's Role in Decision Making. *Front Neurosci* 2011;**5**:115.

52 Ahsan RL, Allom R, Gousias IS, Habib H, Turkheimer FE, Free S, *et al.* Volumes, spatial extents and a probabilistic atlas of the human basal ganglia and thalamus. *Neuroimage* 2007;**38**:261–70.

53 Gläscher J. Visualization of group inference data in functional neuroimaging. *Neuroinformatics* 2009;**7**:73–82.

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Figures

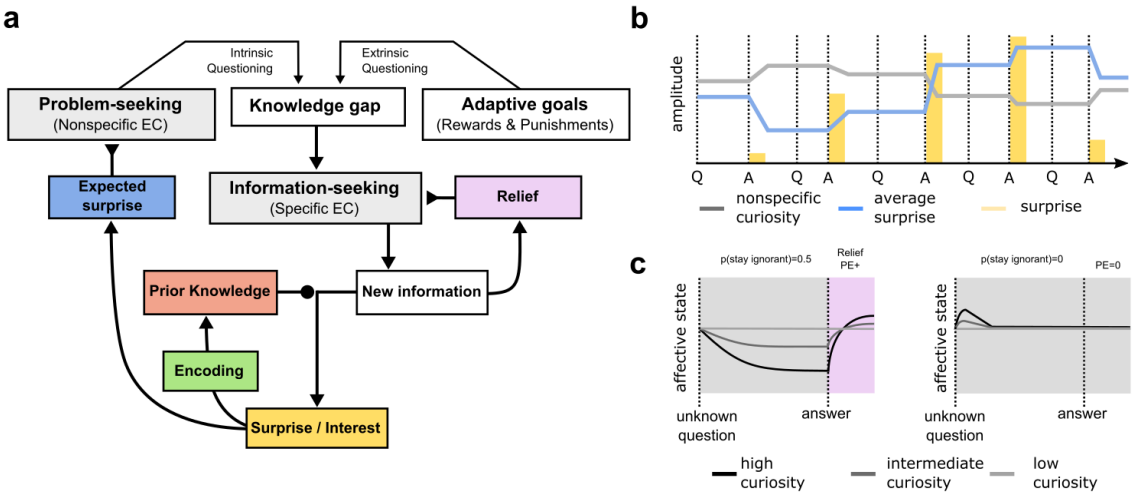

Figure 1. An integrative model of epistemic curiosity (see also Table S1). (a) Graphical representation of the model. The model assumes that all information-seeking behaviors derive from the awareness of a "knowledge gap", resulting from either extrinsic questioning or intrinsic questioning. Corresponding to the specific component of EC, information-seeking consists in exploring the environment in order to gather new information. In turn, the comparison of new information with prior knowledge can elicit epistemic surprise, thereby facilitating memory encoding and updating a dynamic representation of the average surprise experienced in the environment controlling the evolution of nonspecific EC levels. (b) Idealized opponency between nonspecific EC (in grey) and the average amount of surprise experienced (in blue) during a series of questions (Q) and answers (A). The higher this average surprise, the lower the motivational salience of new problems and questions, and vice-versa. (c) Idealized representation of affective states as a function of (specific) EC intensity and the probability of obtaining an answer. When answer delivery is uncertain, aversive states of ignorance may last, so that answer delivery elicits a curiosity-dependent relief-prediction error. Oppositely, when answer delivery is certain, higher curiosity levels predict are only associated with the anticipation of more interesting and relevant information, which may translate into a positive affect.

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

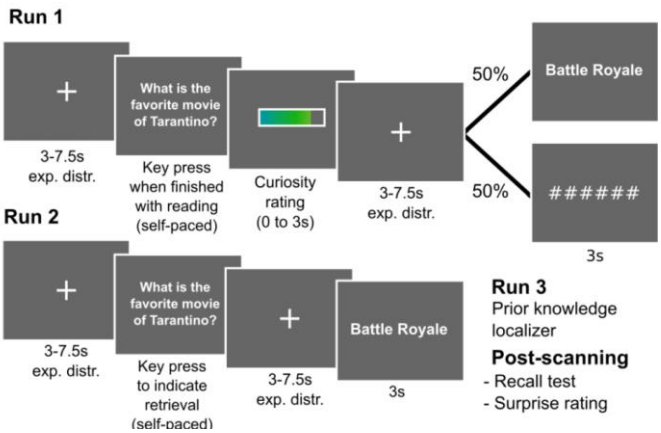

Figure 2. Design of the fMRI study task. In run 1, participants were presented with 60 prescreened trivia items (see also Table S2). After reporting curiosity on a non-numerical continuous gauge, they were presented with either the answer (for half of the trivias), or hashtags (for the other half). In run 2, participants were presented again with the 60 questions and reported whether the answer came spontaneously to their mind (HIT) or not (CR or MISS). Each answer was then revealed, so that half of the answers relieved curiosity whereas the other half merely echoed a previously encountered information. In run 3, participants were presented with an independent set of movie titles they had watched or not (prior knowledge localizer). Once outside the MRI scanner, they were finally asked to report all the answers they could remember and to rate their surprise and interest levels for each trivia answer.

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

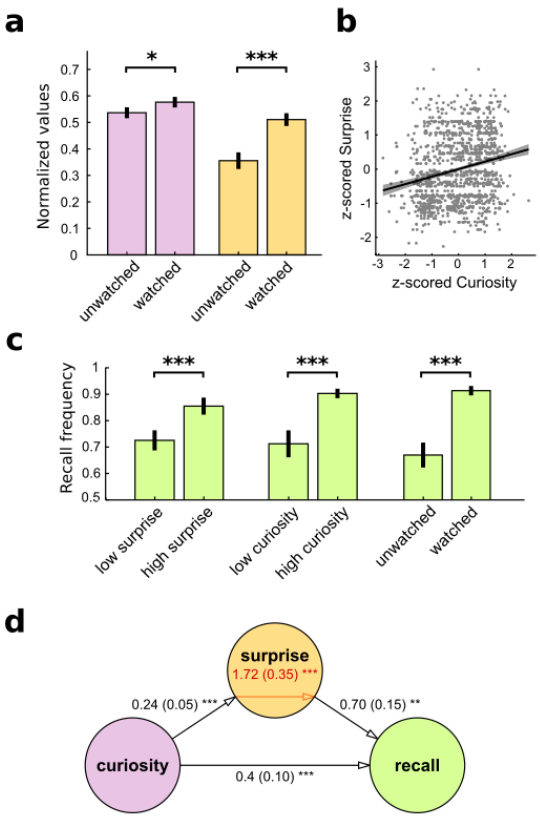

Figure 3. Behavioral results (see also Fig. S1a). (a) The presence of prior knowledge about target answers (whether or not the participant had seen the movie before the experiment) was associated with increased curiosity ratings (rose) and surprise ratings (yellow) in post-scan questionnaires (ratings rescaled between 0 and 1). **(b)** Higher curiosity levels were also associated with increased surprise ratings. **(c)** Recall performances were affected by surprise, curiosity and prior knowledge, as revealed by median-split analyses. **(d)** Participant-wise mediation analyses demonstrated that curiosity induced direct and surprise-mediated benefits for the ability to recall answers in the post-test task (logistic mediation with unsuccessful recall coded as 0 and successful recall coded as 1, with prior knowledge and repetition included as covariates). Error bars represent s.e.m. $p<0.05$*, $p<0.01$**, $p<0.001$**.

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

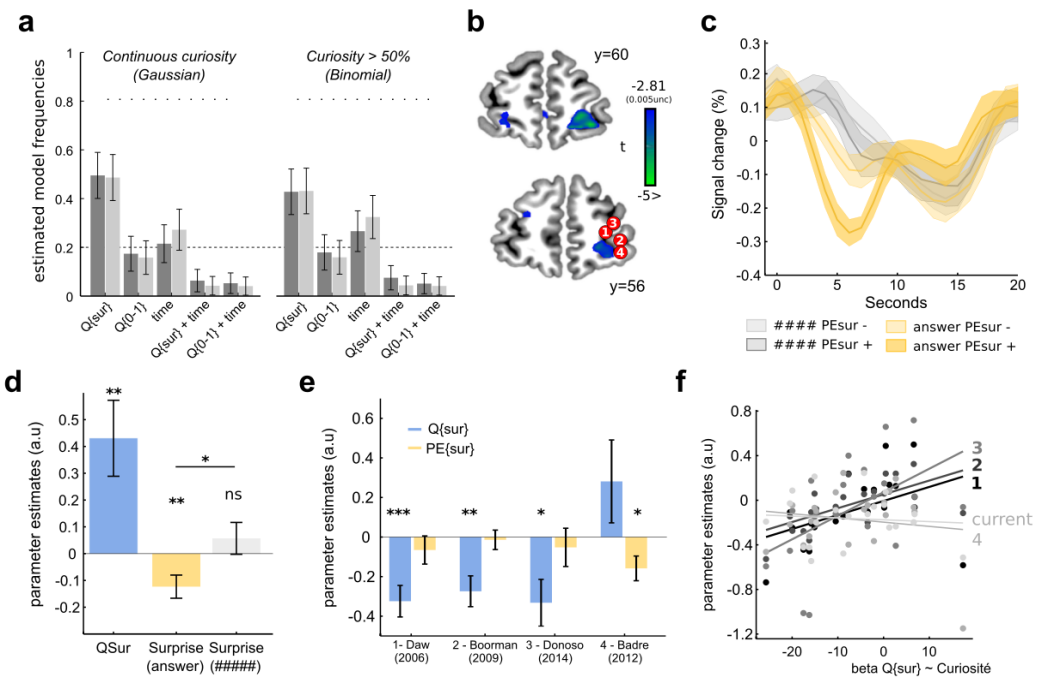

Figure 3. Surprise-dependent control of curiosity and the rlPFC (see also Fig. S1b-e). (a) The model that updated on each trial the average amount of surprise recently experienced outperformed the four alternative models tested to account for the evolution of curiosity ratings represented as Gaussian or binomial variables, for both the Akaike (dark grey) and Bayesian (light grey) Information Criteria. **(b)** Neural correlates of surprise prediction errors at the answer stage (run 1) in the right rostrolateral prefrontal cortex (rlPFC; $p<0.05^{FWE}$, MNI [18 62 -11]; SVC-corrected, prefrontal mask). This functional cluster was close but did not overlap the activation peaks reported in the literature on uncertainty-driven exploration (1: MNI [27 57 6][24], 2: MNI [36, 54, 0][23], 3: MNI [32 56 12][22], 4: MNI [35 56 -8][21]). **(c)** Finite Impulse Response modeling confirmed the presence of genuine rlPFC deactivations in responses to stronger surprise prediction errors, occuring only when answers were actually delivered (yellow). **(d)** The rlPFC also correlated with model-based estimates of average surprise (Q{sur}) at the question stage and with surprise itself when answers but not hashtag were delivered (3mm sphere around peak reported in 4b). **(e)** The profile of activity at previously reported peaks 1-3 was markedly different from the peak reported in 4b. Indeed, in these more lateral and posterior areas of the frontopolar cortex Q{sur} was instead associated with pronounced deactivations at the question stage and no sign of PE{sur} encoding was found at the group level. At peak 4 an intermediate pattern of activity was observed. **(f)** The encoding of PE{sur} at peaks 1-3 was correlated with the the suppression of nonspecific curiosity by Q{sur}, as observed at the behavioral level. p<0.05*, p<0.01**, p<0.001*** (two-tailed). Error bars represent s.e.m. Plotted signals were extracted from 3mm-radius spheres centered around the peaks of interest.

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

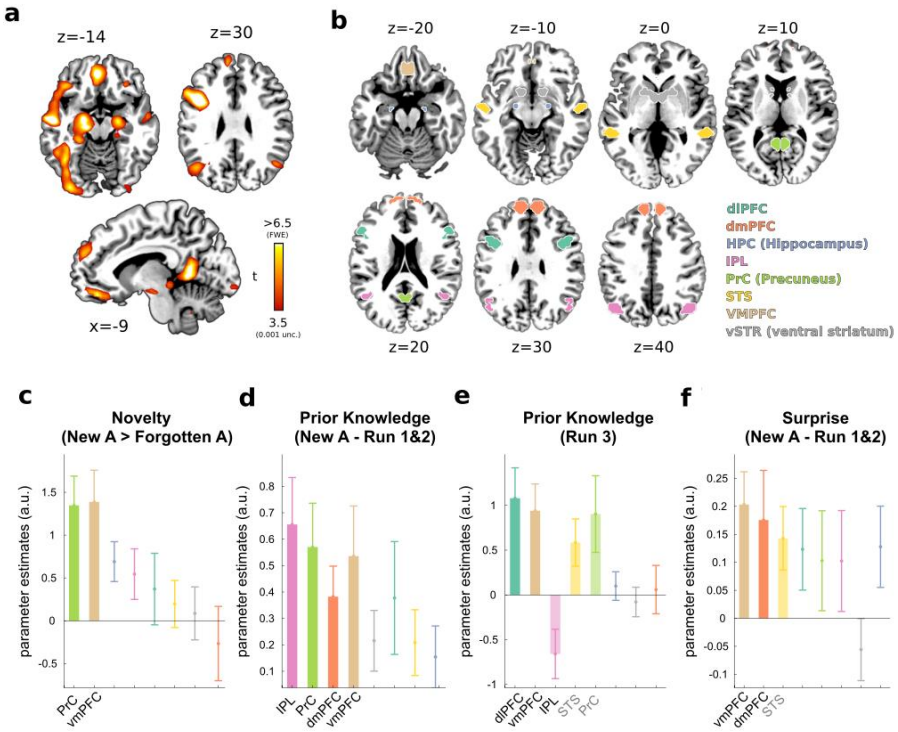

Figure 5. Epistemic surprise signals deriving from the comparison of new information with prior knowledge in the vmPFC (see also Fig. S2). (a) Among other areas, the vmPFC was more activated when processing new answers than hashtags (run 1, $p<0.05^{FWE}$, MNI peak: [-3, 35, -17]) or old answers (run 2, see Table S3). **(b)** Eight bilateral ROIs were systematically investigated to highlight the central role of the vmPFC in the genesis of epistemic surprise (see also Fig. S2a and Fig. 6a). **(c)** Sensitivity to information novelty was revealed by comparing new answers to old but forgotten answers in the second run. **(d)** Encoding of prior knowledge pooled over the two runs of the trivia quiz when processing new answers (watched *versus* unwatched movie titles). **(e)** Encoding of prior knowledge in the localizer task which presented participants with a separate set of watched or unwatched movie titles, independently of any trivia question (run 3). **(f)** Encoding of surprise ratings associated with new answers, pooled over the two runs of the trivia quiz.

In graphs c-f, areas surviving the p(FDR)<0.05 are plotted with plain colors, areas significant only at an uncorrected threshold (p<0.05) are plotted with half-transparent colors and non-significant effects are reported using only error bars. Effect are ordered from left to right as a function of their significance. Error bars represent s.e.m.

bioRxiv preprint doi: https://doi.org/10.1101/157644; this version posted June 30, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

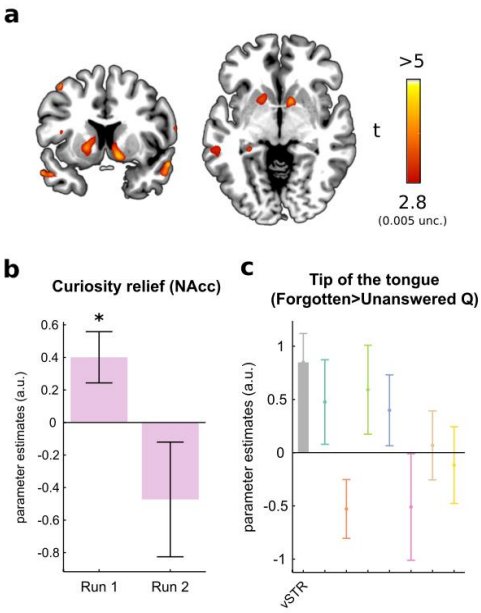

Figure 6. Neural activities related to epistemic curiosity in the ventral striatum. (a) Curiosity levels modulated ventral striatal responses to knowledge delivery during the stochastic trivia quiz (voxel-wise threshold: $p=0.005^{UNC}$; cluster-wise threshold: $p<0.05^{FWE}$; MNI peak: [-6 5 1]). **(b)** In the first run, the effect reported in (A) was also significant in the nucleus accumbens (NAcc) but it disappeared during the second part of the trivia quiz. **(c)** Among the 8 ROIs described in Fig. 5b (same color code), the ventral striatum was the only region to activate significantly more in response to old but forgotten questions in run 2, as compared to never answered questions.

Error bars represent s.e.m. $p<0.05$* (two-tailed).