



Valuation of knowledge and ignorance in mesolimbic reward circuitry

Caroline J. Charpentier^{a,b,1}, Ethan S. Bromberg-Martin^{c,d,e}, and Tali Sharot^{a,1}

^aAffective Brain Lab, Department of Experimental Psychology, University College London, WC1H 0AP London, United Kingdom; ^bDivision of Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125; ^cDepartment of Neuroscience, Columbia University, New York, NY 10032; ^dKavli Institute for Brain Science, Columbia University, New York, NY 10032; and ^eDepartment of Neuroscience, Washington University in St. Louis, St. Louis, MO 63110

Edited by Valerie F. Reyna, Cornell University, Ithaca, NY, and accepted by Editorial Board Member Michael S. Gazzaniga May 30, 2018 (received for review January 10, 2018)

The pursuit of knowledge is a basic feature of human nature. However, in domains ranging from health to finance people sometimes choose to remain ignorant. Here, we show that valence is central to the process by which the human brain evaluates the opportunity to gain information, explaining why knowledge may not always be preferred. We reveal that the mesolimbic reward circuitry selectively treats the opportunity to gain knowledge about future favorable outcomes, but not unfavorable outcomes, as if it has positive utility. This neural coding predicts participants' tendency to choose knowledge about future desirable outcomes more often than undesirable ones, and to choose ignorance about future undesirable outcomes more often than desirable ones. Strikingly, participants are willing to pay both for knowledge and ignorance as a function of the expected valence of knowledge. The orbitofrontal cortex (OFC), however, responds to the opportunity to receive knowledge over ignorance regardless of the valence of the information. Connectivity between the OFC and mesolimbic circuitry could contribute to a general preference for knowledge that is also modulated by valence. Our findings characterize the importance of valence in information seeking and its underlying neural computation. This mechanism could lead to suboptimal behavior, such as when people reject medical screenings or monitor investments more during bull than bear markets.

information seeking | decision making | valence | knowledge | ignorance

People spend a substantial amount of time seeking and consuming information. The quest for knowledge is central both to our modern economy and to our evolutionarily old drive to learn. Despite the importance of information seeking to human behavior we know surprisingly little about what drives the desire for knowledge. Most prevalent theories suggest that humans and other animals are endowed with curiosity because information can help make better decisions that will facilitate obtaining rewards and avoiding harm (1, 2). As information is often useful, curiosity evolved to be a broad feature, which also generalizes to cases when information cannot inform action (3–8). The puzzle, however, is that in domains ranging from health to finance people at times select ignorance, even when knowledge can inform action, such as when people reject medical screenings (9–11).

We pose that valence is central to the process by which the human brain evaluates the opportunity to gain information, explaining why knowledge may not always be preferred. This is because knowledge influences not only people's actions but also their belief state. For example, knowing one has a genetic predisposition for Alzheimer's disease generates a negative belief state, while knowing one is about to receive a promotion generates a positive belief state. This intuition led to the hypothesis that beliefs, just like material goods and services, have utility in and of themselves (12–15). This simple, yet fundamental, notion has significant potential for predicting people's preference for knowledge or ignorance in domains ranging from medicine to politics. If beliefs have utility, people will be motivated to regulate the information they are exposed to (13, 16). In particular,

they will be biased to obtain knowledge that can generate or confirm desirable beliefs (such as seeking positive feedback about their work) (17, 18) and at times to remain ignorant of information that does so for undesirable beliefs (such as avoiding medical tests) (9–11, 19, 20).

Here, we examine whether the human brain represents the value of knowledge as a function of valence and whether this coding is associated with information-seeking behavior. We hypothesized that reward-related brain regions compute errors in predicting the likelihood of obtaining knowledge (21–24) in a valence-dependent manner. Specifically, this suggests that the opportunity to gain knowledge is ascribed greater value when knowledge is expected to generate positive beliefs. This pattern of neural coding would in turn predict a bias to pursue information that can support desirable beliefs over undesirable beliefs. Moreover, if the brain values knowledge as a function of valence, then people should be willing to forgo material rewards to acquire information that can confirm desirable beliefs, even when information cannot inform action, and may at times forgo material rewards to avoid information that can confirm undesirable beliefs.

Our investigation focuses on neural regions most likely to process reward. We used Neurosynth (25), a platform for large-scale, automated synthesis of thousands of fMRI studies, to

Significance

Humans desire to know what the future holds. Yet, at times they decide to remain ignorant (e.g., reject medical screenings). These decisions have important societal implications in domains ranging from health to finance. We show how the opportunity to gain information is valued and explain why knowledge is not always preferred. Specifically, the mesolimbic reward circuitry selectively treats the opportunity to gain knowledge about favorable, but not unfavorable, outcomes as a reward to be approached. This coding predicts biased information seeking: Participants choose knowledge about future desirable outcomes more than about undesirable ones, vice versa for ignorance, and are willing to pay for both. This work demonstrates a role for valence in how the human brain values knowledge.

Author contributions: C.J.C., E.S.B.-M., and T.S. designed research; C.J.C. performed research; C.J.C. and T.S. analyzed data; and C.J.C., E.S.B.-M., and T.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. V.F.R. is a guest editor invited by the Editorial Board.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

Data deposition: Data and code related to this paper are available on GitHub (https://github.com/ccharpen/info_seeking_PNAS).

See Commentary on page 7846.

¹To whom correspondence may be addressed. Email: ccharpen@caltech.edu or t.sharot@ucl.ac.uk.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1800547115/-DCSupplemental.

Published online June 28, 2018.

identify the areas most documented in signaling, processing, and assessing rewards (see also ref. 26 for a similar procedure). This procedure clearly identified the nucleus accumbens (NAc) and midbrain dopaminergic regions ventral tegmental area and substantia nigra (VTA/SN), which have been widely documented in signaling expectations of reward in humans (27–29) and non-human animals (30, 31). To test our hypothesis, we conducted two experiments. In Experiment 1 we tested whether these reward regions represent the opportunity to gain knowledge as a function of valence. In Experiment 2 we tested whether participants pay for knowledge and ignorance and whether these decisions can be explained by the expected valence of knowledge.

Results

In the first experiment we combined a novel behavioral task with functional brain imaging. Participants played a lottery with gain (Fig. 1A) and loss (Fig. 1B) blocks. On gain trials they would either win \$1 or \$0. On loss trials they would either lose \$1 or \$0. At the end of the experiment they received their accumulated earnings. The probability of winning or losing \$1 was displayed on each trial in the form of a pie chart. This allowed explicitly computing the expected value (EV) of the lottery on every trial as the product between outcome probability and magnitude (thus varying from -0.9 to -0.1 for losses and from 0.1 to 0.9 for gains). The participant's task was to indicate whether they would like to reveal the outcome of the lottery. They did so by selecting between two offers, each representing a different probability of having the outcome revealed. Importantly, it was made clear to participants that whether the outcome was revealed had no bearing on their actual earnings. Thereafter either (i) a green knowledge cue appeared, indicating that an informative outcome cue (either win/zero/loss) would follow, or (ii) a red ignorance cue appeared, indicating that a noninformative outcome cue ("XXXX") would follow. Color associations were counterbalanced across participants. After completing this task inside the fMRI scanner participants completed a similar task outside the scanner (Fig. 1C) in which they were shown all lotteries

again, but instead of making choices they indicated how much they would like to know their outcome by moving a cursor on a scale from -300 ("Not at all") to 300 ("Extremely"). All symbols and cues used in the study are described in *SI Appendix, Fig. S1A*.

Preference for Knowledge Is Valence-Dependent. We predicted that participants would prefer information that supports desirable beliefs over undesirable ones (13, 16, 17), and that this bias would sit on top a general curiosity to reveal outcomes (5–7, 22, 23). Indeed, the preference for knowledge was significantly greater when participants expected good outcomes. Participants selected the most informative option (highest information probability offer) on 83.83% of trials [greater than chance: $t(35) = 10.29$, $P < 0.0001$], and they did so more often on gain trials (average = $88.14\% \pm 16.95$ SD) than on loss trials [average = $79.55\% \pm 25.83$; $t(35) = 2.74$, $P = 0.01$] (Fig. 2A). Consistent with these results, participants rated their desire to observe informative cues significantly higher on gain trials (mean rating = 117.5 ± 101.1) than on loss trials [mean rating = 88.77 ± 106.3 ; $t(35) = 3.16$, $P = 0.003$; Fig. 2B].

Importantly, the more likely participants were to win on gain trials the more they wanted to know the outcome [mean slope between probability of winning and information choice = 0.216 ± 0.42 , $t(35) = 3.06$, $P = 0.004$; Fig. 2C]. The more likely participants were to lose on loss trials, the less they wanted to know the outcome [mean slope between probability of losing and information choice = -0.266 ± 0.49 , $t(35) = -3.28$, $P = 0.002$, Fig. 2C]. In other words, preference for knowledge was valence-dependent (Fig. 2E): The greater the EV of the lottery the more likely participants were to select knowledge over ignorance [mean slope between EV and information choice = 0.119 ± 0.19 , $t(35) = 3.71$, $P = 0.001$]. The desire to know, as measured by participants' ratings, was similarly related to the lottery EV (Fig. 2D). Ratings increased with probability of winning on gain trials [mean slope between probability of winning and rating = 185.4 ± 234.9 , $t(35) = 4.73$, $P = 0.00004$] and decreased with probability of losing on loss trials [mean slope between probability of losing




Fig. 1. Experimental design. (A and B) On each trial participants observed a pie representing the likelihood of (A) winning or (B) losing \$1 on that trial. They also observed two offers (blue bars) representing the likelihood of the outcome being revealed (order of pie and bars counterbalanced across trials). Participants selected between these two offers, revealing their preference to observe the outcome cue on that trial. Then, either a knowledge cue (presented here in green) appeared, indicating that the informative outcome cue would be revealed in a few seconds (either win/zero/loss), or an ignorance cue (presented here in red) appeared, indicating that a noninformative outcome cue would be observed (XXXX), followed by a 1-s fixation cross. At the end of the study participants received the accumulation of gains/losses. (C) Postscan rating task in which participants indicated their desire to know following each pie, by moving a cursor on a scale from -300 ("Not at all") to $+300$ ("Extremely"). Color associations were counterbalanced across individuals.

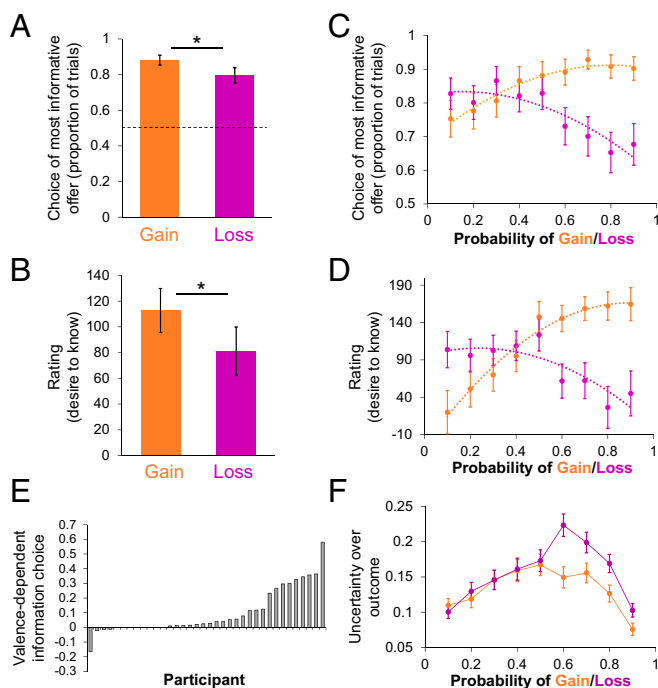


Fig. 2. Preference for knowledge over ignorance is valence-dependent. Participants (A) selected the most informative offer more often during gain than loss trials and (B) rated their desire for knowledge higher during gain than loss trials. Participants (C) were more likely to choose the most informative offer and (D) rated their desire for knowledge higher on gain trials when they were more likely to win (orange curve) and on loss trials when they were more likely to lose (purple curve). Trendlines in C and D represent second-order polynomial fit (C: gain trials: $R^2 = 0.962$; loss trials: $R^2 = 0.823$; D: gain trials: $R^2 = 0.969$; loss trials: $R^2 = 0.726$). (E) Effect of lottery EV on knowledge preference, calculated for each individual participant as the slope between EV and the proportion of times they selected the most informative offer. (F) Uncertainty over outcomes following the delivery of informative or noninformative outcome cues, calculated for each trial as 0 when information is delivered and as the SD of the lottery when information is denied, then averaged separately for gain and loss trials across participants. Error bars represent ± 1 SEM. Two-tailed t test: $*P < 0.05$.

and rating = -98.88 ± 263.8 , $t(35) = -2.25$, $P = 0.031$], thus giving rise to a valence-dependent effect [mean slope between lottery EV and rating = 52.60 ± 71.73 , $t(35) = 4.40$, $P = 0.0001$]. These results were replicated in an independent behavioral pilot sample (*SI Appendix* and *SI Appendix*, Fig. S2).

As displayed in Fig. 2 C and D, the data were fit by a polynomial trendline that included both linear and quadratic components. This nonlinear relationship likely reflects the additional effect of uncertainty on preference for knowledge. Uncertainty is maximal at intermediate probabilities of winning/losing (i.e., 50%). We formally tested this by running a general linear mixed-effect model predicting choices from EV and uncertainty (defined as the SD of the outcome distribution). This revealed the strong positive effect of EV on knowledge choice [estimate = 0.706 ± 0.196 (SE), $t(4,176) = 3.60$, $P = 0.0003$], as well as a smaller but significant effect of uncertainty [estimate = 0.196 ± 0.067 (SE), $t(4,176) = 2.93$, $P = 0.0034$] (see *SI Appendix* and *SI Appendix*, Fig. S3 and Table S1 for additional results and other variable included in the model). In other words, participants were more likely to seek knowledge when the likelihood of winning was high and losing was low, and there was a boost in information seeking when outcome probability was most uncertain (i.e., close to 0.5).

One potential explanation for our results would be simple Pavlovian conditioning: The knowledge cue (green bar) could acquire differential value in the gain vs. loss blocks due to trial-by-trial reinforcement. However, our control analyses show this

was not the case (*SI Appendix*, Fig. S4). As detailed in *SI Appendix*, our data are also inconsistent with the possibility that participants misunderstood the instructions, believing that no information meant no outcome.

Valence-dependent information seeking would inevitably lead to differences in the degree of uncertainty about desirable and undesirable outcomes. To quantify and illustrate the amount of uncertainty that remained about each trial's outcome after the informative or noninformative outcome cue was presented we defined uncertainty as 0 when information was obtained and as the SD of the lottery when information was denied. We then plotted remaining uncertainty averaged over participants and trials for gain and loss blocks for each of the outcome probabilities. This graph illustrates that participants' valence-dependent information-seeking strategy granted greater certainty about positive outcomes, while leaving greater wiggle room for beliefs about negative outcomes, especially when outcome probability was 50% or higher. Indeed, entering these data into an ANOVA confirms not only greater uncertainty about losses than gains [main effect of valence: $F(1,35) = 9.18$, $P = 0.005$] but also an interaction between outcome probability and valence [$F(8,280) = 2.09$, $P = 0.037$; Fig. 2F].

Thus far, the pattern of information preference we observed is consistent with the notion that valence plays a key role in how people value the opportunity to gain knowledge; participants selected knowledge over ignorance more often when information was expected to confirm desirable beliefs than when it was expected to confirm undesirable beliefs. Next, we turned to our fMRI data to ask whether the opportunity to gain knowledge that can support desirable beliefs is also represented by the same neural architecture and code as traditional rewards, and differently from knowledge that can support undesirable beliefs.

Hypotheses for Neural Representation of the Opportunity to Gain Knowledge. If beliefs have utility, then the opportunity to gain knowledge about good outcomes (and thus form certain positive beliefs), but not bad outcomes, may be coded as a primary reward. Electrophysiological studies in nonhuman primates have identified neural signals that encode errors in predicting the opportunity to gain knowledge, which may provide reinforcement for seeking information (22–24). These have been named information prediction errors (IPEs) and are suggested to be analogous to reward prediction errors (RPEs), which encode errors in predicting rewards (22). We hypothesized here that IPEs are valence-dependent (VD-IPEs), coded differently in the domain of gains and losses, and should predict behavioral patterns of preference for knowledge and ignorance.

Alternatively, if the opportunity to gain knowledge is coded as a primary reward regardless of the motivational value of the outcome, then IPEs may be present in reward-related brain areas equally for information about gains and losses, because in both cases knowledge would be gained. These possibilities are not mutually exclusive; for instance, they could be represented by different neuronal populations. Finally, if reward regions only compute the likelihood of receiving material rewards in this task, then IPEs should not be observed. Importantly, our task dissociated the probability of monetary outcomes from the probability of receiving knowledge about those outcomes. This enabled us to examine whether reward regions compute errors in predicting the opportunity to gain knowledge about monetary rewards and losses (IPEs or VD-IPEs) irrespective of errors in predicting those rewards and losses themselves (RPEs).

To test these different predictions, regions of interest (ROIs) were defined by generating a map in Neurosynth (25), a meta-analysis based on 11,406 studies, reflecting the likelihood that the term “reward” was used in a study given the presence of reported activation in a particular voxel. The map reflected the relative selectivity with which voxels activate in relation to “reward,” by comparing all of the studies in the database that contained the term (671 studies for the term “reward”) and all those that did not. This map revealed three peaks: one in the VTA/SN and two in the NAc. A 4-mm sphere was drawn around

each peak to create one bilateral NAc ROI (Fig. 3A) and one VTA/SN ROI (Fig. 3B). A comparison with an anatomical map (32) confirmed that the VTA/SN ROI included voxels both in the VTA and in the SN (*SI Appendix, Fig. S6*).

Neural Representation of Prediction Errors in Gaining Knowledge Are Valence-Dependent in Reward ROIs. We examined whether our ROIs tracked IPEs and VD-IPEs. As portrayed in *SI Appendix, Fig. S1B* the critical time where we should observe IPEs and VD-IPEs but not RPEs is when the knowledge cue (green bar) and ignorance cue (red bar) appear. Both IPEs and VD-IPEs were computed as described in *SI Appendix, Fig. S1B* and entered as parametric modulators of the blood oxygen level-dependent (BOLD) signal (SPM GLM 1; see *Materials and Methods* for details). Specifically, IPE was quantified as “actual opportunity to gain knowledge” (coded as 1 for the knowledge cue and 0 for the ignorance cue) minus “expected opportunity to gain knowledge” (the chosen probability of receiving knowledge). VD-IPE was calculated as IPE multiplied by EV. VD-IPEs represent errors in the expected information gain as a function of probability of gains and losses and thus depend both on the likelihood of receiving information and on the desirability of the expected outcome. Betas were then averaged over all voxels in each ROI. Importantly, EV was also added as a regressor in the model, ensuring that VD-IPE and IPE signals did not simply reflect EV coding. EV, VD-IPE, and IPE were not correlated with each other as observed by correlating the regressors in each participant and then comparing the resulting coefficients to zero [EV and VD-IPE: mean correlation coefficient = -0.027 ± 0.125 (SD), $t(32) = 1.24$, $P = 0.22$; EV and IPE: mean correlation coefficient = 0.017 ± 0.068 (SD), $t(31) = 1.41$, $P = 0.17$; IPE and VD-IPE: mean correlation coefficient = 0.013 ± 0.089 (SD), $t(32) = 0.84$, $P = 0.41$].

The results revealed a significant effect of VD-IPE in the VTA/SN [mean beta = 0.111 ± 0.244 (SD), $t(32) = 2.63$, $P = 0.013$; positive effect observed in 76% of participants; Fig. 4A]. We confirmed that the VD-IPE signal observed in the VTA/SN contained the two key components of a true prediction error (33), tracking the difference between actual and predicted events (SPM GLM 2; see *Materials and Methods* for details): BOLD signal in the VTA/SN ROI tracked the “actual opportunity to gain knowledge” component of the signal positively [knowledge vs. ignorance*EV; mean beta = 0.248 ± 0.51 , $t(32) = 2.79$, $P = 0.009$] and the “expected opportunity to gain knowledge” component negatively [expected knowledge*EV; mean beta = -0.231 ± 0.46 , $t(32) = -2.86$, $P = 0.007$; Fig. 4B]. This suggests that the VTA/SN tracks IPEs as a function of the EV of the outcome.

Valence-independent IPEs were not observed in the ROIs [mean beta in VTA/SN = 0.0397 ± 0.234 , $t(32) = 0.975$, $P = 0.34$; mean beta in NAc = 0.0417 ± 0.242 , $t(32) = 0.991$, $P = 0.33$; SPM GLM 1; see section below for results elsewhere in the brain]. Examining tracking of IPEs separately in the loss and gain do-

main in the VTA/SN revealed an interaction between valence and IPE. Specifically, we computed the BOLD response in the VTA/SN ROI for each participant and trial and entered these betas into a general linear mixed-effect model with fixed and random (participant) effects of IPE, valence (gain vs. loss), and their interaction, as well as fixed and random intercepts. Since behavior shows greater preference for knowledge for gain than for loss trials, combined with a general desire to know (Fig. 2A), we expected IPE to be encoded more positively for gains than for losses. Indeed, this analysis revealed a significant interaction between IPE and valence [$t(3,824) = 1.98$, $P = 0.047$], with no main effect of IPE [$t(3,824) = 1.55$, $P = 0.12$] and no main effect of valence [$t(3,824) = 0.688$, $P = 0.49$]. To visualize this effect, we plotted the average IPE-related activity from this model for each valence (Fig. 4C), showing that, as predicted, the interaction is such that IPEs in VTA/SN are coded more positively for gain trials than for loss trials.

Finally, we asked whether VD-IPE tracking in our ROIs was related to the valence-dependent preference for knowledge over ignorance across participants. In the NAc, a significant correlation was observed between the VD-IPE parameter betas (SPM GLM 1) and how sensitive participants’ choices were to EV [$R(33) = 0.433$, $P = 0.012$; Fig. 4D]. In other words, individuals exhibiting strong tracking of VD-IPE in the NAc also select knowledge more when the chances of winning are high and select ignorance more when the chances of losing are high (see example participant’s data in Fig. 4D), while individuals who select knowledge regardless of EV do not exhibit a VD-IPE signal in the NAc. We confirmed that each component of the VD-IPE signal in NAc was related to behavior: BOLD tracking of the “actual opportunity to gain knowledge” (knowledge or ignorance*EV) component of the VD-IPE in NAc was positively correlated with behavior [$R(33) = 0.419$, $P = 0.015$, *SI Appendix, Fig. S5A*], while the “expected opportunity to gain knowledge” (expected knowledge*EV) component was negatively correlated with behavior [$R(33) = -0.466$, $P = 0.006$, *SI Appendix, Fig. S5B*], with a significant difference between the two correlations (Steiger’s $Z = 2.80$, $P = 0.005$).

EV at the time of knowledge/ignorance cue (SPM GLM 1) was not observed in the VTA/SN [mean beta = 0.0123 ± 0.27 (SD), $t(32) = 0.26$, $P = 0.80$], nor did EV tracking in NAc correlate with behavior [$R(33) = 0.041$, $P = 0.82$], confirming that the results cannot be explained by a signal encoding EV. Control analyses also confirmed that these findings cannot be explained by choice per se (*SI Appendix*). It is of interest that while VD-IPEs were tracked significantly across the group in the VTA/SN, it is the variability across individuals in this signal in the NAc, a major target of VTA/SN projections (34), which was associated with differences in behavior.

Greater Response to Knowledge over Ignorance in the Orbitofrontal Cortex. Behaviorally, our results also showed a general preference for knowledge over ignorance. However, a whole brain-corrected exploratory analysis [familywise error (FWE) $P < 0.05$ cluster-level correction after thresholding at $P < 0.001$ uncorrected] did not reveal tracking of valence-independent IPEs. We thus considered the possibility that the general preference for knowledge vs. ignorance is coded in a more simplistic manner rather than an error signal. To that end, we conducted a whole-brain exploratory analysis comparing BOLD response to knowledge cue with ignorance cue [FWE $P < 0.05$ cluster-level correction after thresholding at $P < 0.001$ uncorrected (35, 36), SPM GLM 3; see *Materials and Methods* for details]. A significant effect was observed in the medial orbitofrontal cortex (OFC) with greater BOLD response to the knowledge cue than to the ignorance cue (Montreal Neurological Institute [MNI] coordinates: [3,59,-5], $T = 4.86$, $k = 49$ voxels, $P = 0.046$; Fig. 5).

Closer examination of this cluster revealed that its activity was significantly tracking valence-independent IPEs [average mean IPE beta from SPM GLM 1 averaged over all voxels in the OFC region identified above = 0.385 ± 0.57 (SD), $t(32) = 3.88$, $P = 0.0005$]. Further tests indicated that it was not, however, a “true”




Fig. 3. ROIs. ROIs were defined as 4-mm-radius spheres around the peaks of the Neurosynth “reward” map (threshold $Z > 10$). Specifically, based on a meta-analysis of 11,406 studies, this map reflects the relative selectivity with which voxels activate in relation to “reward,” by comparing all of the studies in the database that contained the term (671 studies for the term “reward”) to all those that did not. Three peaks were identified on the map, resulting in two ROIs shown in red: (A) bilateral NAc (peaks at $[-10,8,-6]$ and $[12,10,-8]$) and (B) VTA/SN (peak at $[4,-16,-12]$).




Fig. 4. Neural coding for the opportunity to obtain knowledge and remain ignorant is valence-dependent. (A) Individual participants' parameter estimates of VD-IPE (equal to EV times IPE) during knowledge/ignorance cue in the VTA/SN ROI. This effect was significant at the group level, suggesting that IPEs are coded as a function of the EV of the outcome. (B) Individual components of the VD-IPE signal in VTA/SN showed positive tracking of the "actual opportunity to receive knowledge" component and negative tracking of the "expected opportunity to receive knowledge" component. (C) BOLD at the time of knowledge/ignorance cue was extracted for each trial and participant and analyzed in a mixed model with fixed and random effects of valence (binary coding: loss/gain), IPE, and their interaction. Resulting trendlines corresponding to fixed effect estimates of IPE tracking in VTA/SN are shown for gain and loss trials. (D) Individual differences in VD-IPE signal in the NAc ROI correlated with the effect of EV on information choice behavior as measured by the slope of the regression between EV and the proportion of trials in which knowledge is selected over ignorance. Choice of knowledge as a function of EV is plotted for an example participant.

prediction error (SPM GLM 4; see *Materials and Methods* for details). Specifically, only one component of the IPE was tracked, the actual opportunity to gain knowledge [mean beta = 0.48 ± 0.67 (SD), $t(32) = 4.12$, $P = 0.0003$, which is obvious as the voxels were selected as ones where response was greater for knowledge cue than ignorance cue]. However, it did not track the second component of the IPE, the expected opportunity to gain knowledge [mean beta = -0.026 ± 0.67 , $t(32) = -0.21$, $P = 0.83$]. This suggests that the general preference for knowledge over ignorance is likely coded in the OFC using a simple heuristic (knowledge > ignorance) rather than a more sophisticated error term. Additional analyses confirmed that this binary response coded the opportunity to gain knowledge rather than other variables such as EV or choice (*SI Appendix*).

This result accords with previous findings that the OFC codes for the opportunity to increase knowledge (24) for both rewards and punishments (37). We speculated that the NAc integrates a signal from the OFC about the opportunity to increase knowledge together with a valence-dependent signal from the VTA/SN, to produce valenced evaluation of information. This is possible given the known anatomical connectivity between these regions (38, 39). Furthermore, using psychophysiological interaction (PPI) analysis with the NAc ROI used as seed region (see *Materials and Methods* for details), we observed significant enhancement in functional connectivity between the OFC and the NAc [mean PPI beta = 1.22 ± 1.17 (SD), $t(32) = 5.96$, $P < 0.001$], as well as between the NAc and the VTA/SN [mean PPI beta = 0.28 ± 0.52 (SD), $t(32) = 3.10$, $P = 0.004$] during the time of knowledge/ignorance cue. Together, the findings suggest these regions may form a network important for information-seeking decisions.

Tracking of Reward Prediction Errors During Informative Outcome Cue and Expected Value During Pie Presentation. As expected, traditional RPEs were coded in the NAc ROI when informative outcome cue was revealed [i.e., WIN/LOSS/ZERO cue; SPM GLM 1; mean beta in NAc ROI = 0.187 ± 0.313 (SD), $t(32) = 3.425$, $P = 0.002$; *SI Appendix*, Fig. S5C]. We confirmed that this signal contained the key components of a prediction error (SPM GLM 2): It positively tracked actual outcome value [mean beta = 0.248 ± 0.51 , $t(32) = 2.79$, $P = 0.009$] and negatively tracked expected outcome value [mean beta = -0.206 ± 0.52 , $t(32) = -2.28$, $P = 0.03$]. Importantly, this RPE signal in the NAc, in contrast to the VD-IPE signal, was not associated with valence-dependent information-seeking behavior [correlation across participants between strength of NAc RPE signal and behavioral beta: $R(33) = 0.06$, $P = 0.74$; *SI Appendix*, Fig. S5D].

The NAc also coded information about EV. This was observed at the time the lottery (pie) was presented, not at the time the knowledge/ignorance cue was presented. This is sensible as EV information is provided to the participants when they observe the pie, and no new information about EV is provided by the knowledge/ignorance cue. Interestingly, while our initial analysis (SPM GLM 1) did not reveal EV coding per se in the NAc ROI at the time of the pie, exploratory analysis revealed EV was coded in the NAc relative to the average lottery EV in each block, similar to RPE [mean beta = 0.093 ± 0.23 (SD), $t(32) = 2.38$, $P = 0.024$; SPM GLM 5; see *Materials and Methods* for details].

Willingness to Pay for Knowledge and Ignorance. Thus far our results provide evidence that mesolimbic reward systems selectively treat knowledge about favorable outcomes, but not about unfavorable outcomes, as a reward that should be approached. If our interpretation is correct then we expect people should be willing to forgo monetary rewards to gain knowledge that can support positive beliefs, even when information cannot be used to alter outcomes. In addition, participants may at times forgo monetary rewards to avoid knowledge that can support negative beliefs. We tested these predictions in Experiment 2.

Forty-two participants were given £10 at the beginning of the experiment to invest in two out of five stocks in a simulated stock market. On each trial participants observed the evolution of the market (i.e., whether the market was going up or down; Fig. 6A). We confirmed that when the global market was trending upward participants expected their stocks had increased in value and when the market was trending downward they expected their stocks had decreased in value (*Materials and Methods*). Participants then bid

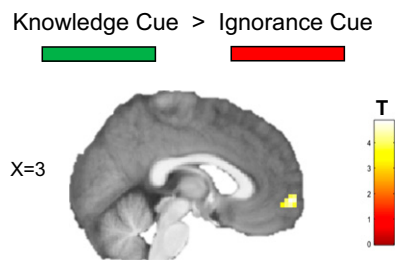


Fig. 5. Knowledge signal in OFC. Medial OFC cluster showing increased BOLD response during delivery of knowledge cues relative to ignorance cues (peak at [3,59,-5]; FWE whole-brain cluster-level corrected $P < 0.05$; threshold $P < 0.001$ uncorrected; overlaid on mean anatomical scan).




Fig. 6. Factors influencing the amount paid for knowledge and ignorance. (A) Participants observed the evolution of a financial market after investing in two of its five companies then decided how much they were willing to pay to receive or avoid information about their portfolio value. (B) Example participant's data showing a positive correlation between market change from the previous trial and signed WTP [coded positively when participants paid for knowledge and negatively when they paid for ignorance; $R(200) = 0.401$, $P < 0.001$]. Each dot represents one trial. (C) A mixed-effects model was run to predict signed WTP across all participants. Estimated coefficients, depicting the fixed effect of two significant factors, signed and absolute market change, are plotted. Additional control variables were added to the model and detailed in *SI Appendix*. (D) The effect of signed market change on signed WTP for knowledge over ignorance, extracted for each individual participant from the mixed-effects model. (E) Plotted is the difference in the total amount participants were willing to pay across all trials in which the market went up minus all trials in which the market went down, separately for when payment was for knowledge and ignorance. Error bars represent \pm SE; * $P < 0.05$.

for a chance to know (or remain ignorant about) the value of their portfolio. Specifically, they indicated how much they were willing to pay to receive or avoid this information on a scale ranging from 99p ("p" indicates pence) to gain knowledge through 0p (no preference) to 99p to remain ignorant. The more they were willing to pay, the more likely their choice was to be honored. Knowledge was noninstrumental: It could not be used to increase rewards, avoid losses, or make changes to their portfolio.

On average participants were willing to pay on 47.7% of trials. On 62.8% of those trials they selected to pay for knowledge (averaging 16.8p) and on 37.1% to remain ignorant (averaging 14.3p). Importantly, willingness to pay (WTP) to receive or avoid knowledge was tied to participants' expectations on whether information would be positive or negative (see example participant in Fig. 6B). Specifically, we used a mixed-effects model to estimate the effect of signed market change on signed WTP (that is, WTP coded positively if participants indicated they wanted to know and negatively if they wanted to avoid knowing), controlling for a host of other factors and with participants as a random factor (*Materials and Methods*). This revealed a significant fixed effect of signed market change [estimate = 1.217 ± 0.54 (SE), $t(8,090) = 2.24$, $P = 0.025$, Fig. 6C; for individual effect estimates see Fig. 6D]. In other words, participants placed greater value on knowledge when the market was promising than when it was

ominous and greater value on ignorance when the market was ominous than when it was promising. There was also an effect of absolute market change, reflecting higher valuation of knowledge in volatile markets, while all other factors had no significant effect (Fig. 6C; see *SI Appendix* for all predictors and statistics).

We confirmed our conclusions with an additional analysis in which the total amount participants indicated they were willing to pay was entered into an ANOVA with choice (knowledge vs. ignorance) and market change (up vs. down) as within-subjects factors. Indeed, this revealed a significant interaction [$F(1,41) = 5.26$, $P = 0.027$; Fig. 6E], characterized by greater WTP for ignorance when the market was down (when participants expected bad news) vs. up (when participants expected good news) and the opposite pattern for WTP for knowledge.

Discussion

Understanding what drives the pursuit and evasion of knowledge is crucial, since information significantly impacts our economy and well-being. Our current findings offer important insight into how humans value the opportunity to gain knowledge. We provide unique empirical evidence that when all else is held constant people value knowledge more about desirable future outcomes than about undesirable ones and value ignorance more about undesirable future outcomes than about desirable ones. Not only was this evident in people's decisions to gain knowledge or remain ignorant but also in the amount they were willing to pay for knowledge or its avoidance.

By recording brain activity of individuals who were making choices between knowledge and ignorance, we show that the opportunity to receive information that strengthens beliefs about future gains, but not losses, is coded in mesolimbic reward regions similarly to primary rewards that should be approached. Specifically, mesolimbic reward regions compute, on a trial-by-trial basis, errors in predicting the opportunity to gain knowledge which scale with the expected valence of that knowledge. The strength of this neural signal in the NAc predicted people's preference for information that can strengthen positive beliefs over negative beliefs. In the OFC, however, we observe a larger response to the opportunity to receive knowledge over ignorance regardless of the likely valence of the information. We speculate that the NAc may integrate a signal from the OFC about the opportunity to increase knowledge together with a value signal in the VTA/SN, to produce a general preference for knowledge that is modulated by valence. Because information is a crucial ingredient for the formation of beliefs, this neural principle has the potential to bias beliefs.

Our findings provide evidence that beliefs have utility in and of themselves. This idea, which is fundamental to theories of belief formation, has roots in philosophy and psychology and has been expressed formally in recent behavioral economic models (12–17). According to these models, positive beliefs generate positive utility and negative beliefs negative utility. These utilities are increased when beliefs are confirmed with certainty. Thus, people will be motivated to boost the positive utility of desirable beliefs by gaining knowledge that can confirm them but motivated to reduce the negative utility of undesirable beliefs by remaining oblivious to information that can confirm them. This bias sits on top a general preference to increase knowledge. It is also consistent with the finding that the longer people have to wait for a reward, the more they prefer advance knowledge (40), presumably because knowing allows boosting the anticipation of the reward. Indeed, real-world observations reveal that people monitor investments more during bull than during bear markets (i.e., when they expect good news) (17, 18).

Notably, by using this information-seeking strategy people are left with a greater degree of uncertainty about future losses than gains, in other words, revealing greater tolerance for uncertainty in the domain of losses. We speculate that this allows people greater wiggle room in their beliefs about losses, affording them the opportunity to bias their beliefs in a desired direction. It would be interesting to examine whether ambiguity-averse individuals, who assign negative utility to outcome uncertainty (41), are less likely to

select ignorance in the domain of loss than other individuals and whether they display a stronger general preference for knowledge regardless of valence. It has also been suggested that ambiguity aversion may be specific to the domain of gain and reduced, or nonexistent, for losses (42, 43), raising the possibility that the observed valence-dependent information seeking may be linked to a valence-dependent bias in ambiguity aversion.

Our study focused on how people value knowledge about valenced outcomes and its underlying neural representation. However, the results also show a strong general preference for knowledge over ignorance. Behaviorally, participants select knowledge at a higher rate than ignorance, and neurally we observe a larger response in the medial OFC to knowledge cues than to ignorance cues. This accords with findings that the OFC codes for the opportunity to increase knowledge (24) for both rewards and punishments (37) and responds to curiosity relief (8). Two previous papers also found that human curiosity (rather than its relief) for nonvalenced trivia questions was associated with enhanced activation in reward regions (21, 44). Curiosity may be the common thread tying our findings, in the sense that people may be more curious about positive outcomes. We speculate that the NAc may integrate a signal from the OFC about the opportunity to obtain knowledge with a valence-dependent value signal in the VTA/SN, to produce greater information seeking when content is expected to be positive than negative. Indeed, there is known anatomical connectivity between these regions (38, 39) and we also find increased functional connectivity between them at the time knowledge and ignorance cues are presented. The findings suggest these regions together form a network important for encoding the value of information (see related suggestions in refs. 8, 22–24, 37, 45, and 46).

Information in our study did not have instrumental value. A future question is whether IPEs in mesolimbic reward circuitry can bias behavior when information has instrumental utility. It has been observed that people's decisions to seek information are likely influenced by valence even when information can inform action. For example, in one study (47), 396 women who gave blood samples were later told that those samples had been analyzed to identify genes that predispose for breast cancer and were asked whether they would like to receive the results. Even though individuals at risk for breast cancer can take precautionary actions to reduce the likelihood of developing the disease, 42% of the participants chose not to know.

Taken together, our findings reveal a biological mechanism that underlies the valuation of knowledge and ignorance, providing insights into this integral part of human behavior. Our work demonstrates that a basic neural computation represented in mesolimbic reward circuitry—errors in predicting the opportunity to gain knowledge—is modulated by valence and ties to a tendency to seek knowledge that can produce desirable beliefs over undesirable ones. In daily life, this valence-modulated information search has considerable implications for many domains including politics, finance, and health behavior. It can lead to suboptimal outcomes if not properly managed, such as when individuals fail to attend medical screenings in an attempt to shun bad news.

Materials and Methods

Participants (Experiment 1). Thirty-nine healthy volunteers were recruited via an advertisement. One participant fell asleep in the scanner, one participant aborted the study because of claustrophobia, and behavioral data files were lost for one participant. Behavioral data are thus reported for the remaining 36 participants [16 males, 20 females; mean age 25.41 \pm 4.59 (SD); age range 18–35 y]. Additionally, data from three participants were excluded from the fMRI analysis due to movement greater than 3 mm; fMRI data are thus reported for 33 participants (14 males, 19 females; mean age 25.61 \pm 4.67). All participants were right-handed, free from past or present psychiatric or neurological disorders, and MRI-safe. The study was approved by the Massachusetts Institute of Technology (MIT) Committee on the Use of Humans as Experimental Subjects and the data collected at the Athinoula A. Martinos Imaging

Center at McGovern Institute for Brain Research, MIT. All participants gave written informed consent and were paid for their participation.

Replication of the main behavioral result was obtained on a pilot sample of 26 participants [11 males, 15 females; mean age 22.84 \pm 4.25 (SD); age range 18–32 y]. Data from these participants were collected at University College London (UCL), and that study was approved by the departmental ethics committee at UCL.

Procedure and Task Design. A phone screening was conducted to ensure participants met eligibility criteria (age between 18 and 35 y, right-handed, no past or present psychiatric or neurological disorder, no alcohol or substance dependence or abuse, no medication or recreational drug use in the week preceding the study, as well as MRI safety criteria). When arriving to the laboratory participants were given instructions and five practice trials. They then completed the information-seeking task while their BOLD signal was recorded. The task was programmed using Cogent Graphics (www.vislab.ucl.ac.uk/cogent_graphics.php) running under MATLAB (<https://www.mathworks.com/>).

On each trial (Fig. 1 A and B) participants were presented with three items sequentially for 3 s each: a pie indicating the likelihood of winning or losing \$1 and two vertical blue bar offers representing the probability of viewing the outcome cue. One bar was presented on the left and one on the right. The order of these stimuli (pie first or bars first) was counterbalanced across trials. All three items were then presented on screen together and participants had up to 3 s to select between the two offers using a button press. Then, either a knowledge cue appeared in the form of a green bar or an ignorance cue appeared in the form of a red bar (colors counterbalanced across participants). The horizontal bar gradually disappeared, wiping out of the screen from left to right, for a jittered duration of 3 s to 7 s. These cues were deterministic; a knowledge cue was always followed by an informative outcome cue for 3 s revealing the lottery outcome (i.e., Win, Zero, or Lose). An ignorance cue was always followed by a noninformative outcome cue ("XXXX") for 3 s. An intertrial fixation cross was shown for 1 s.

Participants were instructed that regardless of whether they viewed an informative or noninformative outcome cue on a given trial the lottery would be played and they would still receive the same money outcome from that trial at the end of the session. They were told there was no right or wrong answer and that they should choose whether they preferred a higher or lower probability of receiving information about their outcome.

The task consisted of two gain blocks (Fig. 1A), in which the outcome of each trial was either win \$1 ("WIN") or nothing ("ZERO"), and two loss blocks (Fig. 1B), in which the outcome was lose \$1 ("LOSE") or nothing ("ZERO"). Each block contained 30 trials and lasted ~11 min. Gain and loss blocks alternated and whether the first block was a gain or loss was counterbalanced across participants. Participants were informed at the beginning of each block whether they were about to play a gain block or a loss block. Block type was also indicated with a different pie color (orange or purple, counterbalanced across participants). The probability of winning or losing varied parametrically from 0.1 to 0.9 (in 0.1 increments) and was independent from the offered probabilities of receiving information, which varied from 0 to 1 (in 0.1 increments). These probabilities of receiving information for the two offers were set on each trial such that they were never equal to each other and such that the difference between the two offers varied uniformly from 0.1 to 1 (in 0.1 increments).

If participants missed a response, the words "Too late! Wait until next trial" were displayed on the screen for 8.5 s (equivalent to the maximum duration of a trial, so that participants would not complete the task quicker by missing responses). Participants were also instructed that they would receive the worst possible outcome when they missed a trial. On average, participants only missed responses on 3.42% of trials.

Immediately after scanning, participants completed a follow-up rating task in which they were presented with the pies used in the main task, but instead of making choices they rated the extent to which they wanted to know the outcome of that lottery (Fig. 1C). The task had two blocks (one for gain and one for loss, in the same order as in the main fMRI task). On each trial they were presented with a pie for 3 s, followed by the question "Would you like to know the outcome?" They responded using the left and right arrow to move a cursor along a rating scale from "Not at all" (300 pixels left from screen center) to "Extremely" (300 pixels right from screen center), pressing the space bar to confirm their ratings. The cursor starting position was randomized around the middle of the scale (between -100 and +100 on the scale from -300 to +300). There was no response time limit. In each block, participants were presented with 18 pies twice in a random order. For consistency with the main task, a knowledge or an ignorance cue was presented for a jittered duration of 3 s to 7 s, followed by the informative outcome cue or noninformative outcome cue for 3 s, and a 1-s fixation cross between trials. The probability of observing the knowledge or

ignorance cue was random (50% probability) and independent from the participant's rating. Outcomes were added to the participant's payment.

Finally, participants completed a comprehension task in which they were shown six stimuli used in the main task (pies and blue bars) and asked to enter the corresponding probability in percent. All participants reported correct percentages with less than 8% error from the true value [mean error across six trials = $1.15\% \pm 2.72$ (SD), range = 0–7.17% error], indicating they understood the mathematical probabilities associated with pies and bars during the main task. They were also given a debriefing questionnaire. Participants received \$60 for completing the study. In addition, they were endowed \$10 and the accumulated outcome from all trials was added or removed from this amount.

Behavioral Data Analysis. Data were extracted using MATLAB and statistical tests were performed using IBM SPSS Statistics (version 22). The proportion of trials in which participants chose the most informative offer (i.e., highest blue bar, referred to as "information choice") was calculated (i) across all trials and compared with a random choice propensity of 0.5 using a one-sample *t* test, (ii) separately for gain and loss blocks and compared using a paired *t* test (Fig. 2A), and (iii) separately for each probability of winning/losing and analyzed by calculating the slope of the best-fit regression line between probability of winning/losing and information choice for each individual. For illustrative purposes, we also plotted a trendline from a second-order polynomial fit (Fig. 2C). Information preference ratings collected in the postscan follow-up task were analyzed in a similar manner as above (Fig. 2B and D).

Individual differences in valence-dependent information seeking were assessed by calculating for each individual the slope of the best-fit regression line between EV (ranging from –0.9 to +0.9) and probability of choosing the most informative target (Fig. 2E).

We also assessed the remaining uncertainty over outcomes following the presentation of informative or noninformative cues. Uncertainty was defined as 0 when informative outcome cue was revealed and as the SD of the lottery when the noninformative outcome cue was revealed (e.g., if the probability of a win on a trial is 0.8, then uncertainty is equal to the SD of a distribution with an 0.8 chance of \$1 and 0.2 chance of \$0). Uncertainty was averaged for each participant separately for gain and loss trials and for each outcome probability (from 0.1 to 0.9) and analyzed in a two- (valence: gain/loss) by-nine (outcome probability) repeated-measures ANOVA (Fig. 2F).

General Linear Mixed-Effect Models of Information Choice. A general linear mixed-effects model of choice was run to assess the effect of several factors on choice, while controlling for others. The dependent variable was choice, coded as 1 if knowledge (higher blue bar) was chosen and 0 if ignorance (lower blue bar) was chosen. Three predictors were included in the model:

- Difference between the information probabilities of the two offers (blue bars):

$$P(\text{info})_{\text{diff}} = P(\text{info})_{\text{high}} - P(\text{info})_{\text{low}}$$

- EV. Because the possible outcomes were \$1 or 0 for gains and –\$1 or 0 for losses this is equal to the probability of gain in gain trials and the negative probability of loss on loss trials:

$$EV = \begin{cases} P(\text{win}) & \text{for gain trials} \\ -P(\text{lose}) & \text{for loss trials} \end{cases}$$

- Outcome uncertainty, defined as the SD of the outcome distribution (i.e., if the probability of a win on a trial is 0.8 then uncertainty is equal to the SD of a distribution with a 0.8 chance of \$1 and 0.2 chance of \$0):

$$\text{Uncertainty} = \begin{cases} \sqrt{P(\text{win})(1 - EV)^2 + (1 - P(\text{win}))(0 - EV)^2} & \text{for gain trials} \\ \sqrt{P(\text{lose})(-1 - EV)^2 + (1 - P(\text{lose}))(0 - EV)^2} & \text{for loss trials} \end{cases}$$

Uncertainty was maximum when the outcome probability was 0.5 (uncertainty = 0.5) and was at its minimum when the outcome probability was at its extremes of 0.1 or 0.9 (uncertainty = 0.3).

These three factors were z-scored and included in the model as fixed effects and random effects (varying between participants). A fixed intercept and a random intercept were also included. This model (model 1) was the full model.

Comparison models were defined as follows (see *SI Appendix, Table S1* for details of model performance):

- Model 2: fixed and random effect of EV; fixed and random intercept
- Model 3: fixed and random effect of uncertainty; fixed and random intercept
- Model 4: fixed and random effect of $P(\text{info})$ difference; fixed and random intercept
- Model 5: null model: fixed and random intercept

All models were run on MATLAB using the fitglm function with a binomial response distribution and maximum likelihood estimated using the Laplace method. For each model, Bayesian information criterion (BIC), Akaike information criterion (AIC), and adjusted R^2 were computed. Given that the models differed in their number of parameters, BIC and AIC (rather than R^2), which penalize models with additional parameters, were used to compare models.

fMRI Data Acquisition. Neuroimaging data were collected on a Siemens Trio 3T MRI scanner using a 32-channel head coil at the Athinoula A. Martinos Imaging Center at McGovern Institute for Brain Research, MIT. Four functional scanning sessions, starting with four discarded dummy volumes and manually stopped 5–10 s after the end of the task block, were acquired using a gradient echo-planar imaging (EPI) sequence with the following parameters: volume repetition time = 2.2 s, echo time = 30 ms, flip angle = 90°, matrix = 64×64 , voxel size = $3.13 \times 3.13 \times 3$ mm³, 36 axial slices automatically placed for whole-brain coverage. To correct for inhomogeneities of the static magnetic field, a field-map sequence was then acquired and used in the unwarping stage of data preprocessing. A T1-weighted MPRAGE anatomical scan was acquired at the end of the session (176 sagittal slices, repetition time = 2.53 s, echo time = 3.48 ms, flip angle = 7°, matrix = 256×256 , voxel size = $1 \times 1 \times 1$ mm³).

fMRI Data Preprocessing. MRI data preprocessing and analysis was performed using SPM8 software (Wellcome Trust Centre for Neuroimaging, www.fil.ion.ucl.ac.uk/spm) in MATLAB. A field map was created for each functional session using the SPM FieldMap toolbox. Using this field-map file for phase correction, images were realigned to the first functional volume of each session and unwarped using seventh-degree B-spline interpolation. Movement was checked using the artifact detection toolbox (https://www.nitrc.org/projects/artifact_detect) to ensure that any scan-to-scan translations greater than one-half of a voxel (1.5 mm) or rotations greater than 1° did not cause artifacts in the corresponding scan(s). The anatomical scan was coregistered to the unwarped mean functional image. All images were then reoriented such that the anterior commissure lay at coordinates [$x = 0, y = 0, z = 0$]. Functional images were spatially normalized to the standard MNI EPI template (with voxels resized to $3 \times 3 \times 3$ mm³) using seventh-degree B-spline interpolation, followed by smoothing using a 6-mm³ FWHM Gaussian kernel.

fMRI Data Analysis. For each participant, GLMs were used to model BOLD signal during the task, incorporating an AR (1) model of serial correlations and a high-pass filter at 1/128 s. Below is a description of the variables used as predictors in the GLM analyses (see also *SI Appendix, Fig. S1B*).

RPE was calculated for each trial in which an informative outcome cue was delivered as the difference between the outcome and its EV:

$$\text{RPE} = V - EV,$$

where V is +\$1 for WIN, 0 for ZERO, and –\$1 for LOSE. EV is equal to the size of the pie (0.1–0.9) multiplied by \$1 on gain trials and multiplied by –\$1 on loss trials.

IPE is equal to the difference between actual and expected opportunity to gain knowledge:

$$\text{IPE} = I - EI,$$

where I is coded as 1 for knowledge cue and 0 for ignorance cue. EI is equal to the height of the chosen blue bar, representing the probability of observing the knowledge cue between 0 and 1.

VD-IPE is calculated as

$$\text{VD-IPE} = \text{IPE} * EV.$$

For all GLMs described below onsets were modeled as stick functions (duration = 0 s). Trials of all four blocks were collapsed into one session in the design matrix. Regressors of no interest included missed choice onset (if any), block type (1: gain block volumes; 0: loss block volumes), and block number (1: first block of each type; 0: second block of each type)—those two regressors accounted for the transitions between blocks after collapsing them into a single session in the model—and six movement parameter regressors in each block. All parametric modulators associated with the same

onset regressor were allowed to compete for variance (the default serial orthogonalization of parametric modulators in SPM8 was turned off).

First-level contrasts were created through linear combinations of the resulting beta images and analyzed at the group level with one-sample *t* tests using the standard summary statistics approach to random-effects analysis implemented in SPM. Exploratory whole-brain analyses were performed with a cluster-forming threshold of $P < 0.001$ uncorrected, followed by cluster-level FWE correction at $P < 0.05$ (35, 36).

Five GLMs were run, with additional control models reported in *SI Appendix*:

GLM 1. The main GLM used included the following predictors: (i) onsets of the knowledge/ignorance cue (i.e., green and red bars) parametrically modulated by (ii) IPE, (iii) EV, and (iv) VD-IPE; (v) onsets of the informative outcome cue (i.e., WIN/LOSS/ZERO) parametrically modulated by (vi) RPE; (vii) onsets of the noninformative outcome cue (i.e., XXXX); (viii) onsets of the pie (lottery), with (ix) EV as a parametric modulator; (x) onsets of each information offer (blue bars) parametrically modulated by (xi) the associated probability of receiving information.

GLM 2. The goal of this model was to conduct a rigorous test of whether the VD-IPE-related activity fulfills the key requirement of a prediction error signal; that is, whether it contains both components of a prediction error, tracking one positively (“actual outcome”) and the other negatively (“predicted outcome”). This GLM was the same as GLM 1 except that the VD-IPE parameter was replaced with its two components: (i) the “actual opportunity to gain knowledge” (the presence of the knowledge or ignorance cue, coded as 1 or 0) multiplied by the trial’s EV (I*EV) and (ii) the “expected opportunity to gain knowledge” (the height of the chosen blue bar) multiplied by EV (Ei*EV) (Fig. 4B). The RPE parameter from GLM 1 was also replaced with its two components: (i) outcome (coded as 1, 0, -1) and (ii) and EV.

GLM 3. The goal of this model was to test whether the general preference for information is represented simply as differential BOLD response to the knowledge cue and ignorance cue. This GLM was the same as GLM 1 except that the IPE parametric modulator was replaced by a binary parametric modulator coding knowledge cue as 1 and ignorance cue as 0.

GLM 4. Similar to GLM 2, the goal of this model was to conduct the analogous test of prediction error coding for IPE. Thus, the IPE parameter was replaced with its two components: (i) the “actual opportunity to gain knowledge” (the presence of the knowledge or ignorance cue, coded as 1 or 0) and (ii) the “expected opportunity to gain knowledge” (the height of the chosen blue bar).

GLM 5. The goal of this model was to test whether EV was coded relative to the average EV of the current context. This was calculated as the difference between the lottery EV (varying from +0.1 to +0.9 for gains and from -0.1 to -0.9 for losses) and the average EV of the current block (average EV = +0.5 for the gain blocks and -0.5 for the loss blocks). This GLM was the same as GLM 1 except that the EV parametric modulator (both at the time of knowledge/ignorance cue and at the time of the pie) was replaced by “relative EV.” In this GLM VD-IPE signal in VTA/SN ROI remained significant [mean beta = 0.138 ± 0.25 (SD), $t(32) = 3.17$, $P = 0.003$], as did the correlation between VD-IPE signal in NAc ROI and behavioral bias [$R(33) = 0.40$, $P = 0.021$]. Relative EV was coded at the time of pie presentation [mean beta in NAc ROI = 0.093 ± 0.23 (SD), $t(32) = 2.38$, $P = 0.024$], not during knowledge/ignorance cue presentation [mean beta in NAc ROI = 0.016 ± 0.20 (SD), $t(32) = 0.48$, $P = 0.64$; mean beta in VTA/SN ROI = 0.014 ± 0.22 (SD), $t(32) = 0.37$, $P = 0.72$].

ROIs. ROIs were built using the Neurosynth (neurosynth.org/) meta-analysis map of the term “reward.” The three highest peaks in this map were located in the left NAc (MNI coordinates [-10,8,6]), the right NAc ([12,10,-8]), and the midbrain VTA/SN ([4,-16,-12]). We built 4-mm-radius spheres around each of these peaks and combined the left and right NAc spheres into one bilateral mask for lack of an a priori lateralization hypothesis, thus leading to two ROIs [bilateral NAc (Fig. 3A) and midbrain VTA/SN (Fig. 3B)]. Overlaying the midbrain ROI with an anatomical atlas of the midbrain (32) indicated that the midbrain ROI contains voxels in both VTA and SN (*SI Appendix*, Fig. S6). We therefore refer to this ROI as VTA/SN throughout the paper. Parameter estimates for contrasts of interest (e.g., VD-IPE signal; Fig. 4A) were calculated for each ROI by averaging across the voxels for each participant and were also correlated with individual differences in the behavioral bias. Bonferroni correction was applied to account for two ROIs wherever effects were tested in both, leading to a two-tailed *t* test threshold of $P < 0.025$.

Trial-by-Trial Extraction of BOLD Signal and Correlation with IPE Values. A trial-by-trial model was built in SPM. For each participant, we created a design matrix in which each presentation of knowledge or ignorance cue (30 per block) was modeled as a separate event (without parametric regressors attached to any of these events). Such a procedure has been used many times in the past (e.g., refs. 26, 48, and 49). Other onsets and regressors (initial stimuli presentation, informative outcome cues, noninformative outcome cues, and movement regressors) were added similarly to the main model (GLM 1) described above. BOLD signal during knowledge/ignorance cue presentation was then extracted from the VTA/SN ROI separately for each trial. A general linear mixed-effects model was run (outside of SPM, using fitglm in MATLAB) to predict BOLD in VTA/SN on each trial from the following predictors: fixed and random (varying between participants) effects of IPE, fixed and random effects of valence (coded as 1 for gain trials and -1 for loss trials), and fixed and random IPE*valence interaction, as well as fixed and random intercepts. The output of interest consisted of the fixed effect *F* statistic for the IPE*valence interaction in predicting BOLD in VTA/SN (Fig. 4C).

Functional Connectivity Analyses. Functional connectivity between the NAc and OFC and between the NAc and VTA/SN was analyzed using PPI in SPM8. The NAc ROI was used as the seed region: BOLD time series was extracted across all voxels in this mask using the Volume of Interest utility of SPM. A first-level model was created for each participant including the deconvolved NAc BOLD time series (physiological regressor), the onsets of the knowledge/ignorance cues (psychological regressor), and their cross-product (PPI regressor). This PPI model also included all other regressors from GLM 1. Two nuisance time series were also added: from a white-matter voxel (corpus callosum body, MNI coordinates [0,14,19]) and from a cerebrospinal fluid voxel (center of the right lateral ventricle, MNI coordinates [4,14,18]). Contrasts were defined on the PPI regressor, reflecting increased functional connectivity with the NAc during presentation of the knowledge/ignorance cues. Functional connectivity betas were extracted over all voxels in each ROI of interest (OFC functional cluster identified in Fig. 5 and the VTA/SN ROI) for each participant and mean betas were compared with zero using a one-sample *t* test.

Stock Market Task (Experiment 2). If beliefs have utility people may be willing to pay to gain knowledge that will induce certainty in positive beliefs and may at times pay to avoid knowledge that will induce certainty in negative beliefs. To test this hypothesis, we ran a separate study.

Forty-five participants were recruited via the UCL subject pool. Data from participants who missed more than one-fourth of all trials were excluded ($N_{\text{excluded}} = 3$). Final data are thus reported on 42 participants [14 males, 28 females; mean age = $28.05 \text{ y} \pm 11.99$ (SD); age range 18–66 y]. The study was approved by the departmental ethics committee at UCL.

The task consisted of four blocks of 50 trials each. At the beginning of each block, participants were endowed with 100 points, worth £10, which they were to invest in two of five fictitious stocks that compose a “global market.” On each trial participants observed the global market evolution (a dynamic increase or decrease in the curve lasting 2.3 s). Given that the participant’s portfolio consisted of two companies out of the five, the global market was a partial indicator of the change to the participant’s own portfolio value. Unbeknown to the participants, on each trial there was a 65% likelihood that their portfolio would follow the market trend; otherwise, the portfolio would vary in the opposite direction than the market with a randomly generated magnitude.

On each trial, after observing the change to the global market (2.3 s) participants were given the opportunity to find out their portfolio value (Fig. 6A). In particular, they had up to 8 s to indicate how much they were willing to pay to know or to remain ignorant. They did so using a scale ranging from 99p to avoid information (“NO”), through 0, to 99p to receive information (“YES”). Left/right positions of YES and NO were counterbalanced across participants. We refer to this scale as the WTP scale (50). The more participants were willing to pay the greater the probability their wish would be honored. If they selected 0p, then information was delivered at random (50%). If they selected to pay to receive/avoid information, then their selection determined the probability their wish was honored. If they selected an amount between 1p and 20p, their wish was honored with 55% probability; between 21p and 40p, their wish was honored with 65% probability, and so on up to 95% probability. Participants were not told the exact mathematical relationship between their payment and their likelihood of receiving/avoiding information. However, they were clearly instructed (both by the experimenter and on the instruction screens) that the more they paid toward YES the more likely they would be to find out their portfolio value and the more they paid toward NO the more likely they would be to NOT find out.

Next, their portfolio value in points was presented on screen or hidden (“XX points” was shown) for 3 s.

At the end of the study one trial was selected randomly (regardless of whether information was delivered) and participants were paid that outcome (e.g., portfolio value of 110 points = £11). In addition, if they paid money to gain or avoid knowledge on that trial and their choice was respected (e.g., they paid 50p to receive information and received it), that amount was deducted from their payment (e.g., £11 – £0.50 = £10.50).

The hypothesis was that when the global market went up participants would expect their portfolio value to increase and would be more eager to have their outcome revealed, and vice versa when the global market went down. This hypothesis assumes that participants’ expectations regarding the value of their own portfolio were related to global market trend. To test this assumption in the last two blocks participants were asked on each trial, after observing the global market, to guess whether their portfolio value increased or decreased relative to the previous trial [from –4 (“decreased a lot”) to +4 (“increased a lot”)] and indicate confidence in their judgment [from 1 (“not confident at all”) to 9 (“extremely confident”). Each rating had an 8-s time limit. The rest of the trial proceeded as before (i.e., WTP scale followed by delivery or denial of information).

Indeed, there was a tight correlation between the market trend and participants’ expectations on whether their value had increased [mean regression slope between market change and expectation rating = 0.168 ± 0.088 (SD), $t(41) = 12.34$, $P < 0.001$]. Furthermore, as one would predict, the stronger the market trend (i.e., the greater the absolute change in market value on current trial relative to the last) the greater the participant’s confidence in their prediction [mean regression slope between absolute market change and confidence rating = 0.017 ± 0.033 (SD), $t(41) = 3.40$, $P = 0.001$].

1. Stigler G (1961) The economics of information. *J Polit Econ* 69:213–225.
2. Hirshleifer J, Riley J (1979) The analytics of uncertainty and information—An expository survey. *J Econ Lit* 17:1375–1421.
3. Berlyne DE (1957) Uncertainty and conflict: A point of contact between information-theory and behavior-theory concepts. *Psychol Rev* 64:329–339.
4. Kreps DM, Porteus EL (1978) Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46:185–200.
5. Grant S, Kajiji A, Polak B (1998) Intrinsic preference for information. *J Econ Theory* 83: 233–259.
6. Eliaz K, Schotter A (2007) Experimental testing of intrinsic preferences for non-instrumental information. *Am Econ Rev* 97:166–169.
7. Eliaz K, Schotter A (2010) Paying for confidence: An experimental study of the demand for non-instrumental information. *Games Econ Behav* 70:304–324.
8. van Lieshout LLF, Vandenbroucke ARE, Müller NCI, Cools R, de Lange FP (2018) Induction and relief of curiosity elicit parietal and frontal activity. *J Neurosci* 38:2816–2817.
9. Thornton RL (2008) The demand for, and impact of, learning HIV status. *Am Econ Rev* 98:1829–1863.
10. Persoskie A, Ferrer RA, Klein WMP (2014) Association of cancer worry and perceived risk with doctor avoidance: An analysis of information avoidance in a nationally representative US sample. *J Behav Med* 37:977–987.
11. Dwyer LA, Shepperd JA, Stock ML (2015) Predicting avoidance of skin damage feedback among college students. *Ann Behav Med* 49:685–695.
12. Caplin A, Leahy J (2001) Psychological expected utility theory and anticipatory feelings. *Q J Econ* 116:55–79.
13. Golman R, Loewenstein G (2016) Information gaps: A theory of preferences regarding the presence and absence of information. *Decision*, 10.1037/dec0000068.
14. Koszegi B (2003) Health anxiety and patient behavior. *J Health Econ* 22:1073–1084.
15. Koszegi B (2010) Utility from anticipation and personal equilibrium. *Econ Theory* 44: 415–444.
16. Golman R, Hagman D, Loewenstein G (2017) Information avoidance. *J Econ Lit* 55:96–135.
17. Karlsson N, Loewenstein G, Seppi D (2009) The ostrich effect: Selective attention to information. *J Risk Uncertainty* 38:95–115.
18. Sicherman N, Loewenstein G, Seppi DJ, Utkus SP (2016) Financial attention. *Rev Financ Stud* 29:863–897.
19. Eil D, Rao JM (2011) The good news-bad news effect: Asymmetric processing of objective information about yourself. *Am Econ J Microecon* 3:114–138.
20. Ganguly A, Tasoff J (2016) Fantasy and dread: The demand for information and the consumption utility of the future. *Manage Sci* 63:4037–4060.
21. Gruber MJ, Gelman BD, Ranganath C (2014) States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron* 84:486–496.
22. Bromberg-Martin ES, Hikosaka O (2009) Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron* 63:119–126.
23. Bromberg-Martin ES, Hikosaka O (2011) Lateral habenula neurons signal errors in the prediction of reward information. *Nat Neurosci* 14:1209–1216.
24. Blanchard TC, Hayden BY, Bromberg-Martin ES (2015) Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron* 85:602–614.
25. Yarkoni T, Poldrack RA, Nichols TE, Van Essen DC, Wager TD (2011) Large-scale automated synthesis of human functional neuroimaging data. *Nat Methods* 8:665–670.
26. Garrett N, Lazzaro SC, Arieli D, Sharot T (2016) The brain adapts to dishonesty. *Nat Neurosci* 19:1727–1732.
27. Knutson B, Taylor J, Kaufman M, Peterson R, Glover G (2005) Distributed neural representation of expected value. *J Neurosci* 25:4806–4812.
28. Kable JW, Glimcher PW (2007) The neural correlates of subjective value during intertemporal choice. *Nat Neurosci* 10:1625–1633.
29. Rigoli F, Chew B, Dayan P, Dolan RJ (2016) Multiple value signals in dopaminergic midbrain and their role in avoidance contexts. *Neuroimage* 135:197–203.
30. Cromwell HC, Schultz W (2003) Effects of expectations for different reward magnitudes on neuronal activity in primate striatum. *J Neurophysiol* 89:2823–2838.
31. Tobler PN, Fiorillo CD, Schultz W (2005) Adaptive coding of reward value by dopamine neurons. *Science* 306:1642–1646.
32. Murty VP, et al. (2014) Resting state networks distinguish human ventral tegmental area from substantia nigra. *Neuroimage* 100:580–589.
33. Rutledge RB, Dean M, Caplin A, Glimcher PW (2010) Testing the reward prediction error hypothesis with an axiomatic model. *J Neurosci* 30:13525–13536.
34. Haber SN, Fudge JL, McFarland NR (2000) Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. *J Neurosci* 20:2369–2382.
35. Eklund A, Nichols TE, Knutsson H (2016) Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. *Proc Natl Acad Sci USA* 113:7900–7905.
36. Flandin G, Friston KJ (2016) Analysis of family-wise error rates in statistical parametric mapping using random field theory. arXiv:1606.08199.
37. Jessup RK, O’Doherty JP (2014) Distinguishing informational from value-related encoding of rewarding and punishing outcomes in the human brain. *Eur J Neurosci* 39:2014–2026.
38. Haber SN, Knutson B (2010) The reward circuit: Linking primate anatomy and human imaging. *Neuropsychopharmacology* 35:4–26.
39. Marquand AF, Haak KV, Beckmann CF (2017) Functional corticostriatal connection topographies predict goal directed behaviour in humans. *Nat Hum Behav* 1:0146.
40. Iigaya K, Story GW, Kurth-Nelson Z, Dolan RJ, Dayan P (2016) The modulation of savouring by prediction error and its effects on choice. *eLife* 5:e13747.
41. Levy I, Snell J, Nelson AJ, Rustichini A, Glimcher PW (2010) Neural representation of subjective value under risk and ambiguity. *J Neurophysiol* 103:1036–1047.
42. Cohen M, Jaffray J-Y, Said T (1987) Experimental comparison of individual behavior under risk and under uncertainty for gains and for losses. *Organ Behav Hum Decis Processes* 39:1–22.
43. Tymula A, Rosenberg Belmaker LA, Ruderman L, Glimcher PW, Levy I (2013) Like cognitive function, decision making across the life span shows profound age-related changes. *Proc Natl Acad Sci USA* 110:17143–17148.
44. Kang MJ, et al. (2009) The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychol Sci* 20:963–973.
45. Smith DV, Rigney AE, Delgado MR (2016) Distinct reward properties are encoded via corticostriatal interactions. *Sci Rep* 6:20093.
46. Tricomi E, Fiez JA (2012) Information content and reward processing in the human striatum during performance of a declarative memory task. *Cogn Affect Behav Neurosci* 12:361–372.
47. Lerman C, et al. (1998) What you don’t know can hurt you: Adverse psychological effects in members of BRCA1-linked and BRCA2-linked families who decline genetic testing. *J Clin Oncol* 16:1650–1654.
48. Edelson MG, Dudai Y, Dolan RJ, Sharot T (2014) Brain substrates of recovery from misleading influence. *J Neurosci* 34:7744–7753.
49. Charpentier CJ, Moutsiana C, Garrett N, Sharot T (2014) The brain’s temporal dynamics from a collective decision to individual action. *J Neurosci* 34:5816–5823.
50. Becker GM, DeGroot MH, Marschak J (1964) Measuring utility by a single-response sequential method. *Behav Sci* 9:226–232.