

PNAS

www.pnas.org

Supplementary Information for

The Ventral Striatum Dissociates Information Expectation, Reward Anticipation, and Reward Receipt

Flavia Filimon, Jonathan D. Nelson, Terrence J. Sejnowski, Martin I. Sereno,
Garrison W. Cottrell

Address correspondence to: Terrence J. Sejnowski or Flavia Filimon
Email: terry@salk.edu, f.filimon@surrey.ac.uk or flavia.filimon@gmail.com

This PDF file includes:

- Supplementary Methods
- Supplementary Results
- References for SI
- Figs. S1 to S6
- Tables S1 to S3

Supplementary Methods

Number of training trials. Because people vary widely in the number of trials required to learn probabilistic classification tasks, the number of training trials was adapted automatically by the software according to each subject's performance.

Subjects required between 205 and 1079 training trials (mean=563.4, median=656, s.d.=275.2) in the first behavioral training session, 1-4 days in advance of scanning; and between 201 and 829 training trials in the refresh training session, immediately before scanning (mean=418.2, median=325, s.d.=242.4). Note that due to technical limitations in the software, the number of training trials for 1 subject, and the number of refresh trials for 9 subjects, could only be computed to a precision of ± 12 trials following achievement of the learning criterion (which required a minimum of 200 trials).

Most (9 of 10) subjects were given initial behavioral training on the probabilistic plankton classification stimuli at least one day (and up to four days) prior to the fMRI experiment. To ensure that subjects had not forgotten the probabilities, they were given refresh training immediately prior to the fMRI session. The refresh training was identical to the initial behavioral training, with the same stringent performance criterion, to ensure knowledge of the probabilistic classification task structure. One subject was unable to come on two separate days. This subject received initial behavioral training on the immediately prior to the fMRI session, and thus was not given additional refresh training before scanning.

Calculating the informativeness of each cue. The expected information value of each feature can be calculated in the following way. Various Optimal Experimental Design (OED) models could be used. We use probability gain, since Nelson, McKenzie, Cottrell and Sejnowski (2010) found it to be the best psychological model on a related task. Other OED models (Nelson, 2005; Crupi, Nelson, Meder, Cevolani & Tentori, 2018) agree that the 85% accuracy feature is much more useful than the 60% accuracy feature.

A feature's expected probability gain is the expected reduction in classification error (Bayes's error) that could be obtained on average if that feature were viewed, assuming an optimal (always guess what is most probable, given everything known) response strategy.

Let K be the class (category) variable, taking value k_1 if the specimen is species A, and value k_2 if the specimen is species B. Let D be a feature that can be observed, with its possible values d_1 and d_2 . We use $u_{pg}(d_1)$ to denote the probability gain of viewing feature d_1 , $u_{pg}(d_2)$ the probability gain of viewing feature d_2 , and $eu_{pg}(D)$ to denote the expected probability gain of viewing feature D , before whether it is known whether the feature takes version d_1 or d_2 .

Before viewing feature D , the more probable of k_1 and k_2 would be guessed (if category A and B were not equally likely, a priori). The probability of guessing correctly would be $\max(P(k_1), P(k_2))$, and the probability of error would be $1 - \max(P(k_1), P(k_2))$.

After viewing feature D , the more probable of k_1 and k_2 , given the observed value of feature D , would be guessed. Suppose that d_1 is observed. The probability of correct classification, given this new information, is $\max(P(k_1|d_1), P(k_2|d_1))$, and the probability of error is $1 - \max(P(k_1|d_1), P(k_2|d_1))$. The probability gain, given that d_1 is observed, is the prior probability of classification error minus the posterior probability of classification error given d_1 , namely $u_{pg}(d_1) = [1 - \max(P(k_1), P(k_2))] - [1 - \max(P(k_1|d_1), P(k_2|d_1))]$.

Similarly, if d2 were observed, the probability gain would be
 $u_{pg}(d2) = [1 - \max(P(k1), P(k2))] - [1 - \max(P(k1|d2), P(k2|d2))]$.

Before a feature is viewed, the value that it will take is not known, so the expected probability gain, $eu_{pg}(D)$, is the average of the probability gain given each feature value, weighted by each feature value's probability of occurrence:

$$eu_{pg}(D) = P(d1) * u_{pg}(d1) + P(d2) * u_{pg}(d2).$$

For the 85% accuracy feature,

$$\begin{aligned} u_{pg}(d1) &= [1 - \max(P(k1), P(k2))] - [1 - \max(P(k1|d1), P(k2|d1))] \\ &= [1 - 0.5] - [1 - 0.85] \\ &= 0.35 \end{aligned}$$

$$u_{pg}(d2) \text{ is similarly } 0.35. \text{ Thus, } eu_{pg}(D) = 0.5 * 0.35 + 0.5 * 0.35 = 0.35$$

We could similarly calculate that the expected probability gain of the 60% accuracy feature is 0.10.

To make the above calculations more concrete, and to see how the probability gain of each individual feature can be seen from the joint probabilities of the feature combinations, consider a case in which the Eye feature is the HI information feature, and the Claw feature is the LO information feature. In this case, the frequencies of the four possible stimulus combinations, and the associated probabilities of the more probable category, were:

$$\begin{aligned} P(\text{eye1} \ \& \ \text{claw1}) &= 28.5\%; \ P(\text{species A} \ | \ \text{eye1} \ \& \ \text{claw1}) = 89.47\%; \\ P(\text{eye1} \ \& \ \text{claw2}) &= 21.5\%; \ P(\text{species A} \ | \ \text{eye1} \ \& \ \text{claw2}) = 79.07\%; \\ P(\text{eye2} \ \& \ \text{claw1}) &= 21.5\%; \ P(\text{species B} \ | \ \text{eye2} \ \& \ \text{claw1}) = 79.07\%; \\ P(\text{eye2} \ \& \ \text{claw2}) &= 28.5\%; \ P(\text{species B} \ | \ \text{eye2} \ \& \ \text{claw2}) = 89.47\%. \end{aligned}$$

The accuracy that can be achieved by looking at both features is 85% on average ($2 * .285 * .8947 + 2 * .215 * .7907$). This is equal to the accuracy that can be achieved solely by looking at the more-informative (Eye) feature:

$$\begin{aligned} P(\text{eye1}) &= P(\text{eye2}) = 0.5; \\ P(\text{species A} \ | \ \text{eye1}) &= 85\%, \ P(\text{species B} \ | \ \text{eye1}) = 15\%; \\ P(\text{species A} \ | \ \text{eye2}) &= 15\%, \ P(\text{species B} \ | \ \text{eye2}) = 85\%. \end{aligned}$$

The less informative (claw) feature leads to only 60% classification accuracy on its own:

$$\begin{aligned} P(\text{claw1}) &= P(\text{claw2}) = 0.5; \\ P(\text{species A} \ | \ \text{claw1}) &= 60\%, \ P(\text{species B} \ | \ \text{claw1}) = 40\%; \\ P(\text{species A} \ | \ \text{claw2}) &= 40\%, \ P(\text{species B} \ | \ \text{claw2}) = 60\%. \end{aligned}$$

The probabilities in the experiment were designed so that the best choice would be an A classification in half of trials, and a B classification in half of trials, such that in each trial there would be maximal uncertainty about both the motor response and the true category, up until the point that the specific feature value was revealed. For each subject, for each trial, the stimulus was drawn at random (without replacement) from the true probabilities. Therefore, the experienced probabilities (averaging over at least 200 learning trials, per subject) would have been very similar to, but not exactly the same as, the theoretical probabilities.

Supplementary Results

Calculations of Theoretical Response Profiles to each condition in the event-related fMRI trials under Information Gain (IG) and Reward Prediction Error (RPE) Models

The calculations underlying the hypothesized amount of activation (height of each bar in the bar graphs in Figures 3c(I), 3C(II), and 3D) are as follows:

A) Predicted response profile for a brain area that treats expected information as a predictor for reward (Figure 3C(I))

1) Information Expectation stage:

At the beginning of a trial, the subject knows that s/he will receive a blurry cue that predicts information leading to either 85% or 60% categorization success. Half the trials contain 85% cues, the other half 60% cues. Thus, on average a categorization success of $(85\%+60\%)/2 = 72.5\%$ is expected before the trial has even started. Once an actual blurry cue arrives, a prediction error can be calculated. The prediction error and predicted activation profile during the blurry cue (Information Expectation) stage is thus:

$85 \text{ (actual)} - 72.5 \text{ (expected)} = 12.5 \text{ units of "excitement" during the HI Information Expectation phase}$

$60 \text{ (actual)} - 72.5 \text{ (expected)} = -12.5 \text{ units of "excitement" during the LO Information Expectation phase.}$

The first two bars in Figure R1A show this hypothetical response profile.

2) Information Revelation and Outcome Anticipation stage:

During the information revelation stage (Information Revelation and Outcome Anticipation), there is no surprise and hence no prediction error, since the blurry cue predicted exactly whether the cue would be an 85% or 60% cue, even if the exact category (A or B) could not yet be chosen during the Information Expectation stage. Instead of indicating "0 activation", which is physiologically unrealistic, we plot a small positive activation (+3 units) so that the bars can be seen in the graph. The third and fourth bars in Figure R1A show this hypothetical response profile.

3) Reward Receipt stage:

Finally, what is the average surprise (reward prediction error) a subject experiences upon receiving smiley feedback, aggregating across HI and LO Information trials? In HI Information trials, the subject expects to receive a smiley (reward) with 85% probability; in other words, the mean expectation is for 85 smile (reward) units on each smile trial. If a smile occurs, the surprise is:

$100 \text{ (actual)} - 85 \text{ (expected)} = 15 \text{ smile (reward) units; thus, the prediction error is 15 units. Similarly, during LO Information trials, there is 60\% probability of receiving a smiley (reward); thus, the mean expectation is 60 smiley (reward) units. On these trials, the prediction error upon receipt of a smile (reward) is:}$

$100 \text{ (actual)} - 60 \text{ (expected)} = 40 \text{ smiley (reward) prediction error units during the Feedback Stage.}$

What is the average surprise (reward prediction error), aggregating among all smile trials? Among all trials in the experiment, half are HI Information and half are LO Information. However, because the HI Information trials have a greater tendency to result in correct

classification, they comprise about 59% of the smile trials¹. Thus, the surprise associated with HI Information trials has greater weight in the aggregate smile feedback predictions. LO Information trials are less likely to result in correct classification, comprising about 41% of all smiley trials. Therefore, the average surprise or prediction error across all smile trials is $(0.5862*15 + 0.4138*40) = 25.34$ units. This is how surprised subjects are to receive smile feedback, aggregating across both HI and LO Information trials. This value is plotted as a bar above the smiley symbol in Figure R1A.

An analogous rationale applies in the case of the frown feedback trials. In HI Information frown trials, failure to receive a reward is associated with surprise (reward prediction error) of $0 - 85 = -85$ units (0 received, 85 were expected). In LO Information frown trials, failure to receive a reward is associated with surprise (reward prediction error) of $0 - 60 = -60$ units. In HI Information trials, frown feedback is rare; HI Information trials comprise only 27.27% of frown trials.² LO Information trials more frequently result in frown feedback, comprising 72.72% of frown trials. Therefore, the average surprise (reward prediction error) in frown trials is $(0.2727 * (-85)) + (0.7272 * (-60)) = -66.82$. This value is plotted as a bar above the frowny symbol in Figure R1A.

B) Predicted response profile for a brain area that is sensitive to expectation of information

What is the predicted response pattern for a region that is sensitive to expected information, as quantified in Optimal Experimental Design theory, but that is not sensitive to obtained information or to reward? The expected usefulness of a query according to OED theory can be quantified using information gain (expected reduction in Shannon entropy), probability gain (expected reduction in classification error), or other models (Nelson, 2005). We plot predicted response profiles according to information gain (Figure 3D). For deriving the response profile, we take as a starting point the 72.5% classification accuracy that one achieves on average in the task $((85\%+60\%)/2 = 72.5\%$ average accuracy a subject expects to obtain), and its associated entropy. We then calculate the difference in expected entropy that is occasioned by seeing a particular (HI or LO information) blurry feature.

¹ The percent of rewarded (smile feedback) trials belonging to HI vs. LO Information cue trials is calculated as follows: The probability of obtaining smile feedback is 85% given a HI Information cue, but only 60% given a LO Information cue, assuming optimal responding with consistent choice of the optimal (most probable, given the feature shown) category. The subjects achieved near-ceiling performance behaviorally, so this is a reasonable simplification. Thus, the proportion of all trials that are HI Information (85% cue) and receive a smile feedback is: 50% of all trials * 85% probability of obtaining a smile = 42.5% of all trials. The proportion of all trials that are LO Information (60% cue) and smile feedback is: 50% of all trials * 60% probability of obtaining a smile = 30% of all trials. Thus, the total proportion of trials that receive a smile is 42.5%+30%=72.5%. Given the above, we can calculate that among all *smile* trials, $42.5/72.5 = 58.62\%$ are HI Information trials, and $30/72.5 = 41.38\%$ are LO Information trials.

² Among all trials, 7.5% are HI Information trials that result in frown feedback. [50% of all trials are HI Information, and 15% of these result in frown feedback; $50\% * 0.15 = 7.5\%$.] Among all trials, 20% are LO Information trials that result in frown feedback. [50% of all trials are LO Information, and 40% of these result in frown feedback; $50\% * 0.4 = 20\%$.] Thus, 27.5% of all trials result in frown feedback. Among frown trials, $7.5/27.5 = 27.27\%$ are HI Information, and $20/27.5 = 72.72\%$ are LO Information.

If the HI information cue is presented, the prior entropy of 0.8485 bits (the Shannon entropy corresponding to a [0.725, 0.275] probability distribution) decreases to 0.6098 bits (the entropy corresponding to a [0.85, 0.15] probability distribution), for a gain of 0.2387 bits of information, plotted as 23.87 units of activation in Figure 3D. If however the LO information cue is presented, entropy increases from 0.8485 bits to 0.9710 bits (the entropy in a [0.60, 0.40] probability distribution), for a decrease of 0.1224 bits of information, plotted as -12.24 units in Figure 3D. (For equations and numeric examples, see Nelson, 2005, Appendix A).

If, rather than Shannon entropy, Error entropy (Bayes's error) is used to quantify uncertainty, a predicted response profile can be calculated based on expected probability gain, which is expected reduction in Error entropy (an inverted V curve; see Crupi et al., 2018). If the HI information cue is presented, the prior Error entropy of 0.275 (the Error entropy corresponding to a [0.725, 0.275] probability distribution) decreases to 0.150 (the Error entropy corresponding to a [0.85, 0.15] probability distribution), for a gain of 0.125. If however the LO information cue is presented, Error entropy increases from 0.275 to 0.400 bits (the Error entropy in a [0.60, 0.40] probability distribution), for a decrease of 0.125.

Note that information could also be calculated according to expected reduction in other entropy models; the pattern of results does not depend strongly on whether Shannon entropy or other psychologically plausible (for instance high-order Arimoto entropy) entropy models (Crupi, Nelson, Meder, Cevolani & Tentori, 2018) are used.

Additional explanation of correlation and bootstrap analyses

Results suggest that VS subregions respond differentially to expectation of information, versus to expectation and receipt of reward. A further question is whether, or the extent to which, the response patterns across trial stages and event types match the theoretical expected information gain (IG) or reward prediction error (RPE) models.

To address this, we conducted analyses to explore how well the aggregate and/or individual-subject data — across the six trial stages and event types (InfoExpect_HI, InfoExpect_LO, RevealAnticip_HI, RevealAnticip_LO, Feedback_POS and Feedback_NEG) — might be explained by the theoretical IG or RPE models. We considered various methods for quantifying these relationships. We used Pearson correlation in this and subsequent analyses; other correlation coefficients (for instance Kendall or Spearman rank correlations) produce similar results, as does a measure based on signed percent of variance explained. Due to variability in magnitudes of individual subjects' BOLD (% signal change) data we focused on subjects' rank-based data. We used Matlab's tiedrank function, within each subject's data across the trial stages and event types, before aggregating across subjects. Aggregate results (Figure 3E) suggest that left lateral VS is better explained by the IG model, whereas right VS/NAcc is better explained by the RPE model.

We also conducted additional analyses to better understand variability among subjects and its implications, again using the within-subject-ranked data in these analyses. We first addressed whether, across the six trial stages and event types, the IG or RPE model provided a "better" explanation of a particular ROI. As an interpretable measure of how much better a particular model fits data from a particular ROI, we took those data's correlation to the IG model, minus those data's correlation with the RPE model. Positive values of this correlation difference statistic therefore imply that the IG model provides a better explanation, whereas negative values imply that the RPE model provides a better explanation. We report counts of individual subjects for which a particular model correlates more highly with each ROI. We also report bootstrap

confidence intervals (simple bootstrap, based on 1 million bootstrap samples) for the difference in correlation to the theoretical IG model, minus to the theoretical RPE model, for each ROI. The input to these analyses is based on the individual subjects' differences in correlation to the IG minus to the RPE model.

Additionally, we further used the whole dataset, in each case taking a bootstrap sample of 10 subjects with replacement, across the six trial stages and event types, to check the tendency (using simple bootstrap sampling of the whole dataset, with 1 million bootstrap samples) for a particular model to better explain (more highly correlate with) a particular ROI. Figure S6, left and center panels, provides histograms of the results of these bootstrap analyses for the left lateral VS and for the right VS/NAcc, respectively. Finally, we considered a relative information sensitivity statistic, based on taking the IG-RPE correlation difference in the left lateral VS ROI, minus the IG-RPE correlation difference in the right VS/NAcc ROI. We conducted analogous analyses for this statistic, namely reporting the proportion of the subjects for which it was positive, and a bootstrap confidence interval (based on the individual subjects' relative information sensitivity statistic) for the statistic. We also checked the proportion, across 1 million bootstrap samples of the whole dataset, for this statistic to be positive. Figure S6, right panel, provides a histogram of the results of this bootstrap analysis. Matlab code (`corrAnalVoi.m`) for these analyses is also provided.

Supplementary Figures

Fig. S1. Percent signal change in ventral striatal BOLD activations in individual subjects.

Fig. S2. Group-level GLM contrasts versus baseline. BOLD activations ($p < 0.005$) for each regressor are shown both above (red-to-yellow) and below (dark-blue-to-light-blue) baseline. Cross hairs are centered on the ventral striatum (VS; $x,y,z = -14,8,-12$). Red arrow indicates above-baseline, positive BOLD activation in the left lateral VS, during InfoExpect_HI only. Blue arrows show deactivation (negative BOLD) in the nucleus accumbens (NAcc) during InfoExpect_LO and Feedback_NEG, suggesting a functional difference between NAcc and left lateral VS.

Fig. S3. Activation overlap between contrasts in ventral striatum. BOLD activation overlaps in ventral striatum for contrasts InfoExpect_HI vs. baseline, InfoExpect_HI vs _LO, and InfoExpect_LO vs. baseline, shown at different p-values. Top row shows substantial overlap in left ventral striatum between InfoExpect_HI>LO (red, $p < 0.05$, small-volume corrected) and InfoExpect_HI vs. baseline (blue, $p < 0.01$, basal-ganglia-volume corrected). The overlap is indicated in purple. Middle row: the same contrasts are shown at more stringent thresholds. The amount of overlap is more visible in the axial view (right column, middle). Bottom row: InfoExpect_LO vs. baseline yields only voxels that are deactivated vs. baseline, and only in nucleus accumbens, not left lateral ventral striatum. Thus, lateral parts of the left ventral striatum show positive BOLD activation in response to HI Information Expectation, whereas medially, bilateral nucleus accumbens is deactivated in response to LO Information Expectation.

Fig. S4. Percent BOLD signal change across the different trial stages (HI and LO InfoExpect, HI and LO RevealAnticip, and Feedback_POS and _NEG) in ROIs from Fig. 2. The shaded gray rectangle in each panel indicates the contrast based on which the ROIs were selected. The percent signal change bars are for illustrative purposes; no additional t-tests were performed on extracted BOLD signals, since statistical testing was already performed in the selection of the ROIs ($\alpha < 0.005$, $k=25$; see Figs 2 and S5). Error bars represent standard error of the mean.

Fig. S5. Group BOLD activation contrasts across the different stages of the trial: (A) Information Expectation; (B) Information Revelation & Outcome Anticipation; (C) Feedback. (A) red-to-yellow: Brain areas more active for LO than HI Info Expectation. (B) green-to-light-green: Areas more active for revelation of LO than HI Info features and concomitant Outcome Anticipation. (C) blue-to-light-blue: Areas more active for receipt of negative than positive feedback. Voxels shown are significant at $\alpha < 0.005$, cluster-corrected with a 25-voxel cluster threshold, except for the InfoExpect LO>HI contrast (A), which did not survive the α threshold and is shown at an uncorrected α of 0.05. Abbreviations: SFG=superior frontal gyrus; gyr=gyrus; IFS/IFG=inferior frontal sulcus/gyrus; STS=superior temporal sulcus; IPL=inferior parietal lobule; lat.=lateral; PFC=prefrontal cortex; ant.=anterior; see also Fig. 2. For full list of activations, see Tables S1-3. Colors represent different stages of the trial, not positive vs. negative BOLD signal.

Fig. S6. Histograms from bootstrap sampling of whole dataset within-subject-ranked activations. Left panel shows left lateral VS ROI, correlation to Information Gain (IG) model minus correlation to Reward Prediction Error (RPE) model. Center panel shows right VS/NAcc ROI, correlation to IG minus correlation to RPE model. Right panel shows differential sensitivity to IG minus RPE, subtracting the right VS/NAcc (IG - RPE) score minus the left lateral VS (IG - RPE) score.

Supplementary Tables

Table S1. InfoExpect HI>LO

Average MNI coordinates (in mm) for brain regions showing stronger activations for High than Low Information Expectation (see also Fig. 2). All activations $p < 0.005$, $k = 25$. Abbreviations: HI=High information; LO=Low information; L=left; R=right.

Note: No activations survived the opposite contrast, InfoExpect LO>HI at this threshold. For uncorrected $\alpha = 0.05$ activations, see Fig. S5A.

Table S2. RevealAnticip HI>LO; RevealAnticip LO>HI

Average MNI coordinates for brain regions showing stronger activations for RevealAnticip_HI>LO as well as the reverse contrast (Figs 2B and S5B, respectively). All activations $p < 0.005$, $k = 25$.

Table S3. Feedback POS>NEG; Feedback NEG>POS

Average MNI coordinates for brain regions showing stronger BOLD activations for positive Feedback (smile emoticon) than negative Feedback (frown emoticon), as well as the reverse contrast (Feedback_NEG>POS). All activations $p < 0.005$, $k = 25$, except where noted.

Supplemental References

Crupi V, Nelson JD, Meder B, Cevolani G & Tentori K (2018) Generalized information theory meets human cognition: Introducing a unified framework to model uncertainty and information search. *Cogn Sci* 42:1410-1456.

Nelson JD (2005) Finding useful questions: on Bayesian diagnosticity, probability, impact, and information gain. *Psychol Rev* 112:979–999.

Nelson JD, McKenzie CRM, Cottrell GW & Sejnowski TJ (2010) Experience matters: information acquisition optimizes probability gain. *Psychol Sci* 21:960–969.

LEFT (lateral) VENTRAL STRIATUM

RIGHT VENTRAL STRIATUM (NAcc)

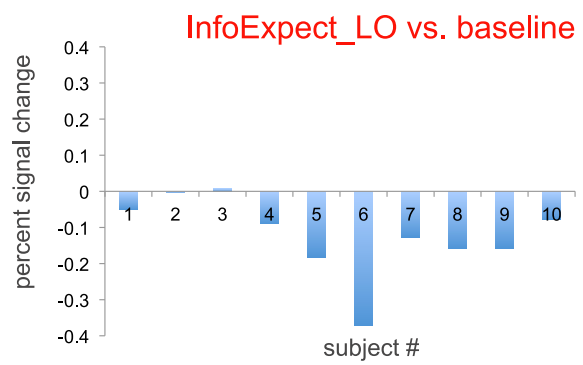



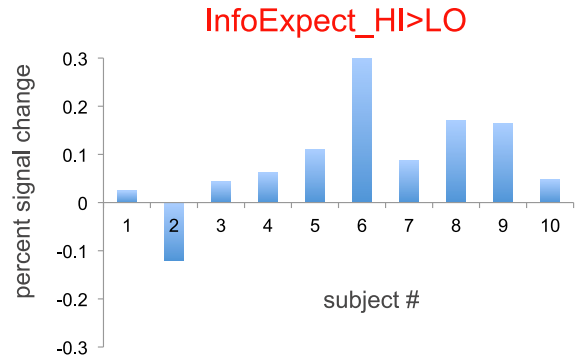



Figure S1. Percent signal change in ventral striatal BOLD activations in individual subjects.

ACTIVATION OVERLAP FOR GLM CONTRASTS InfoExpect_HI vs. InfoExpect_LO vs. baseline

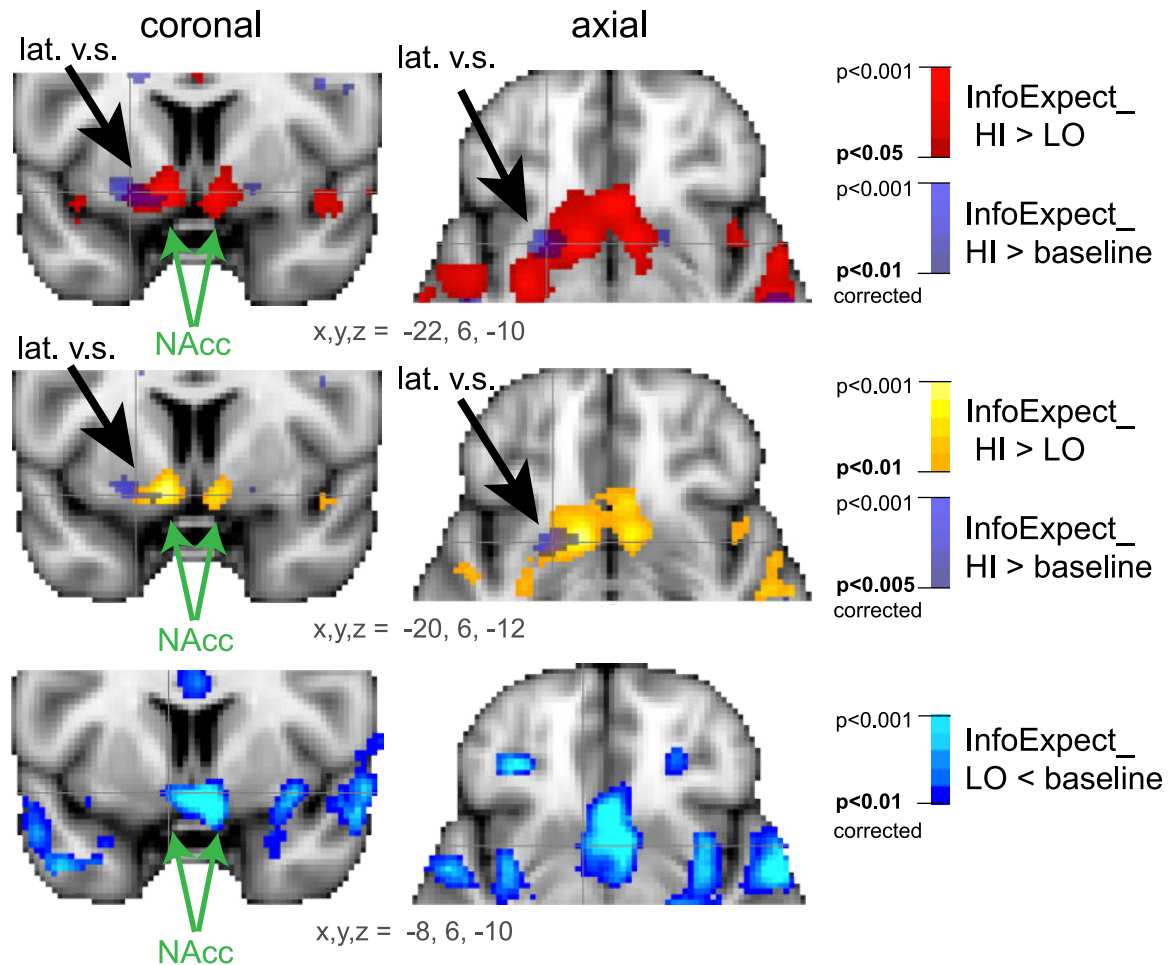


Figure S3. Activation overlap between contrasts in ventral striatum. BOLD activation overlaps in ventral striatum for contrasts InfoExpect_HI vs. baseline, InfoExpect_HI vs. LO, and InfoExpect_LO vs. baseline, shown at different p-values. Top row shows substantial overlap in left ventral striatum between InfoExpect_HI > LO (red, $p < 0.05$, small-volume corrected) and InfoExpect_HI vs. baseline (blue, $p < 0.01$, basal-ganglia-volume corrected). The overlap is indicated in purple. Middle row: the same contrasts are shown at more stringent thresholds. The amount of overlap is more visible in the axial view (right column, middle). Bottom row: InfoExpect_LO vs. baseline yields only voxels that are deactivated vs. baseline, and only in nucleus accumbens, not left lateral ventral striatum. Thus, lateral parts of the left ventral striatum show positive BOLD activation in response to HI Information Expectation, whereas medially, bilateral nucleus accumbens is deactivated in response to LO Information Expectation.

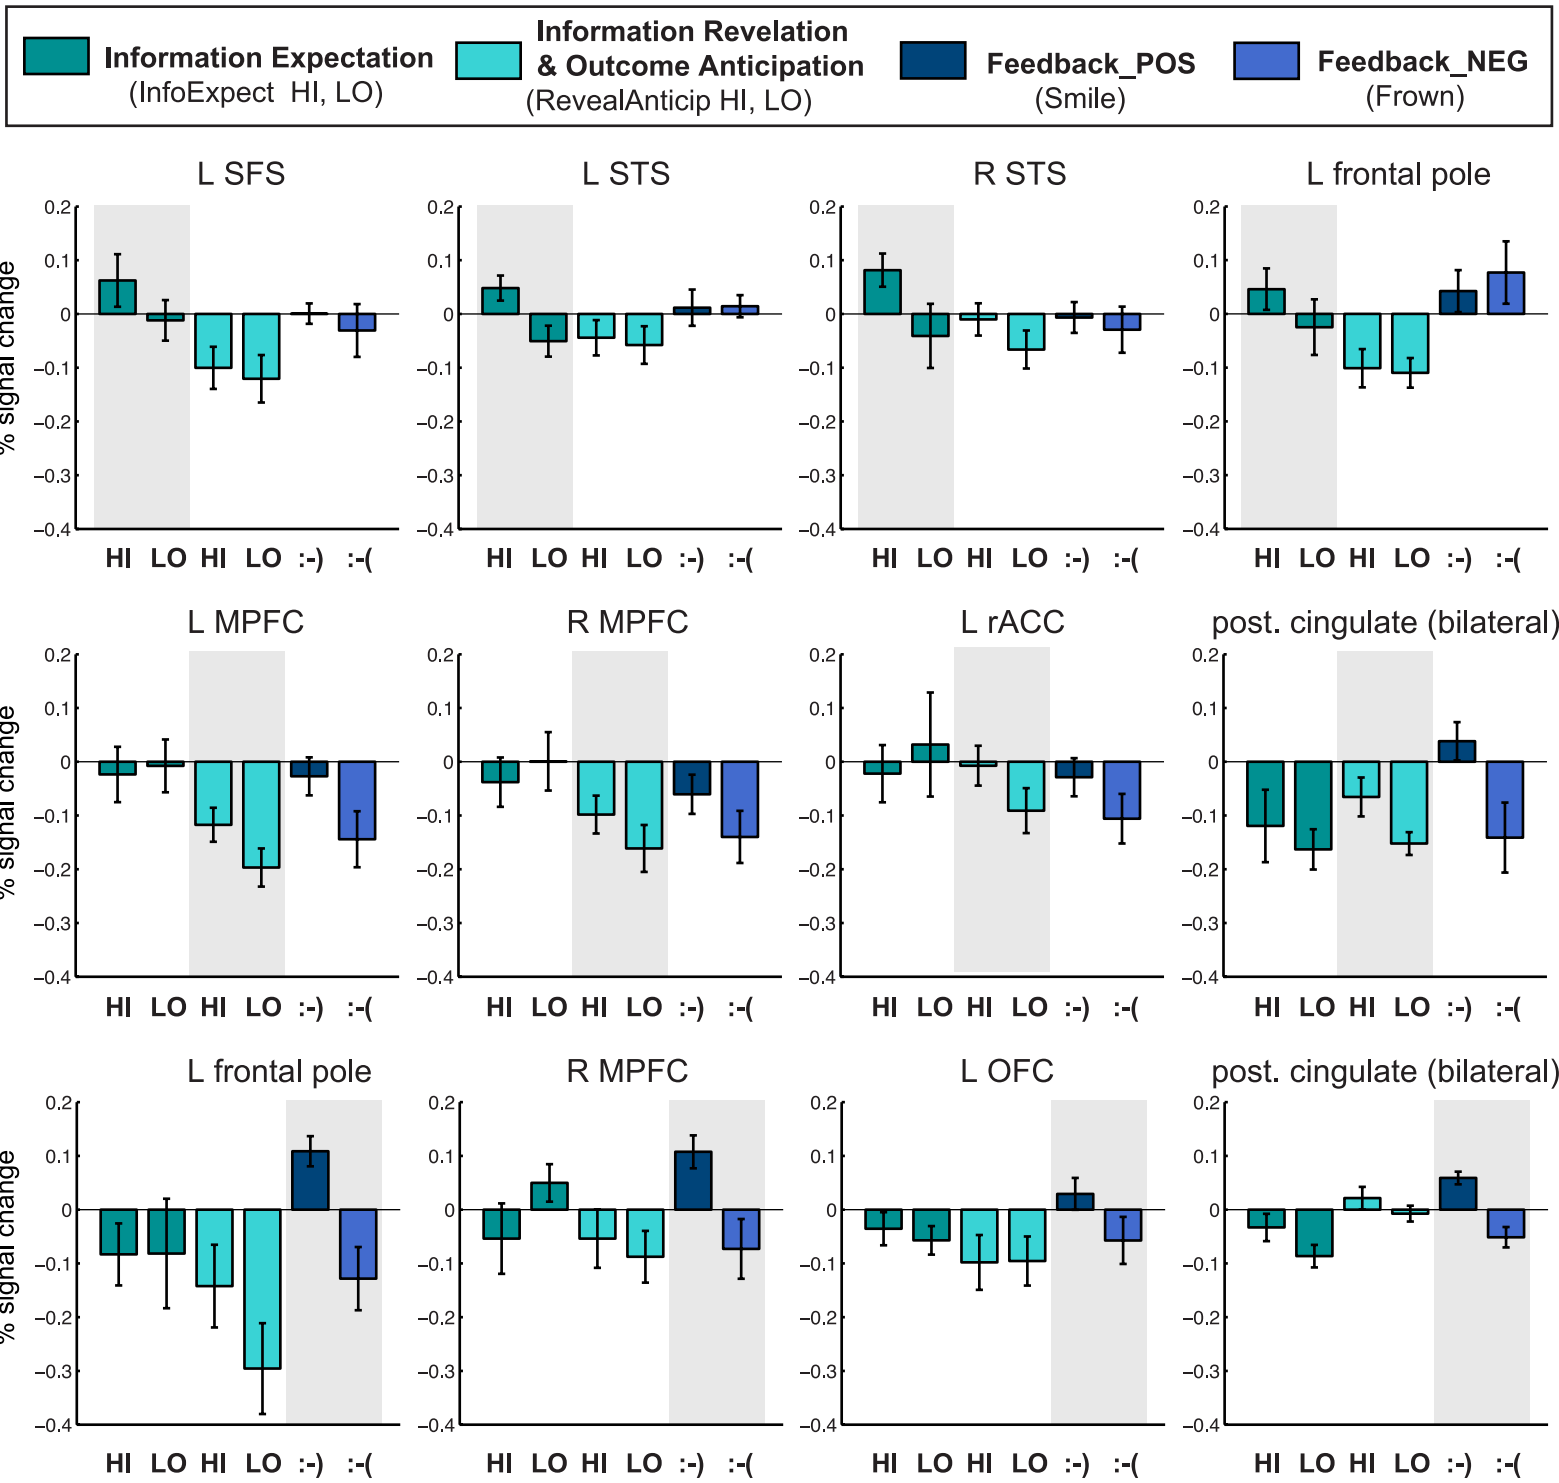


Figure S4. Percent BOLD signal change across the different trial stages (HI and LO InfoExpect, HI and LO RevealAnticip, and Feedback_POS and _NEG) in ROIs from Fig. 2. The shaded gray rectangle in each panel indicates the contrast based on which the ROIs were selected. The percent signal change bars are for illustrative purposes; no additional t-tests were performed on extracted BOLD signals, since statistical testing was already performed in the selection of the ROIs ($\alpha < 0.005$, $k = 25$; see Figs 2 and S5). Error bars represent standard error of the mean.

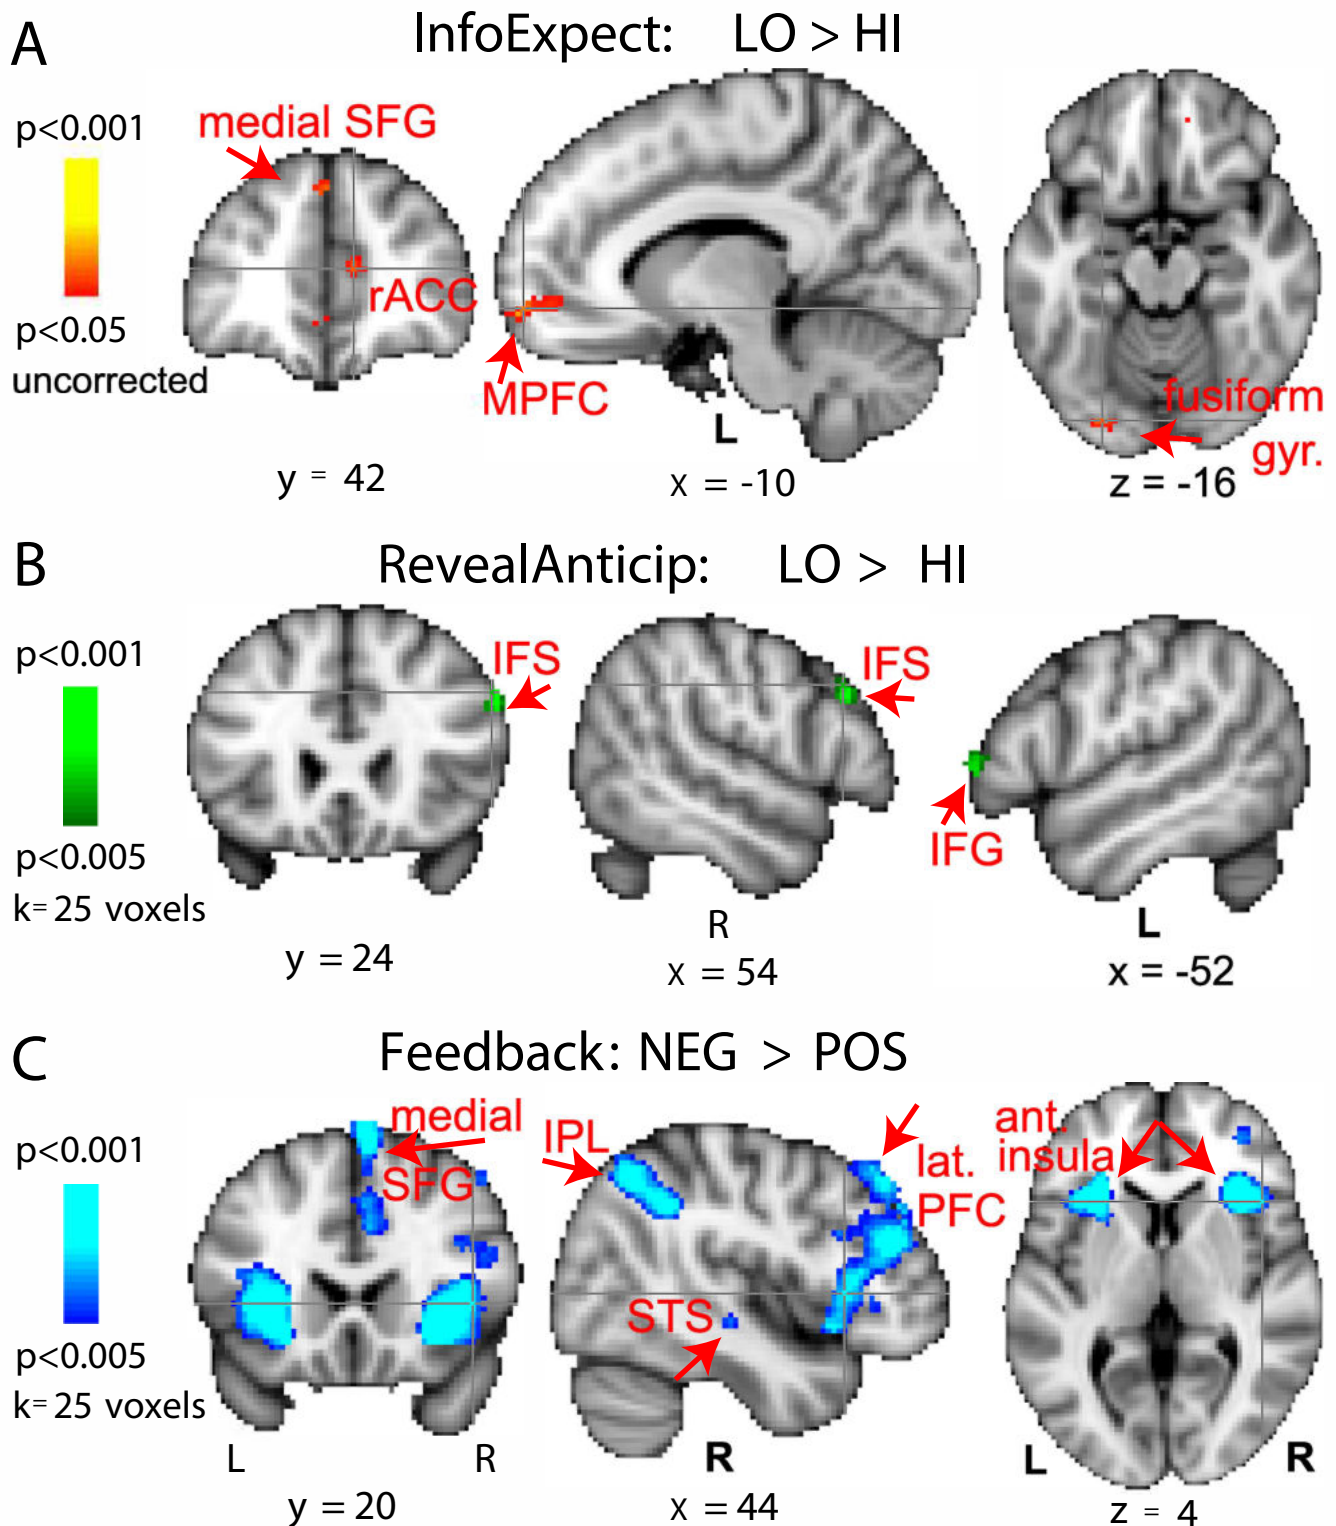


Figure S5. Group BOLD activation contrasts across the different stages of the trial: (A) Information Expectation; (B) Information Revelation & Outcome Anticipation; (C) Feedback. (A) red-to-yellow: Brain areas more active for LO than HI Info Expectation. (B) green-to-light-green: Areas more active for revelation of LO than HI Info features and concomitant Outcome Anticipation. (C) blue-to-light-blue: Areas more active for receipt of negative than positive feedback. Voxels shown are significant at $p < 0.005$, cluster-corrected with a 25-voxel cluster threshold, except for the InfoExpect LO>HI contrast (A), which did not survive the a threshold and is shown at an uncorrected a of 0.05. Abbreviations: SFG = superior frontal gyrus; gyr=gyrus; IFS/IFG = inferior frontal sulcus/gyrus; STS = superior temporal sulcus; IPL = inferior parietal lobule; lat.=lateral; PFC = prefrontal cortex; ant.=anterior; see also Fig. 2. For full list of activations, see Tables S1-3. Colors represent different stages of the trial, not positive vs. negative BOLD signal.

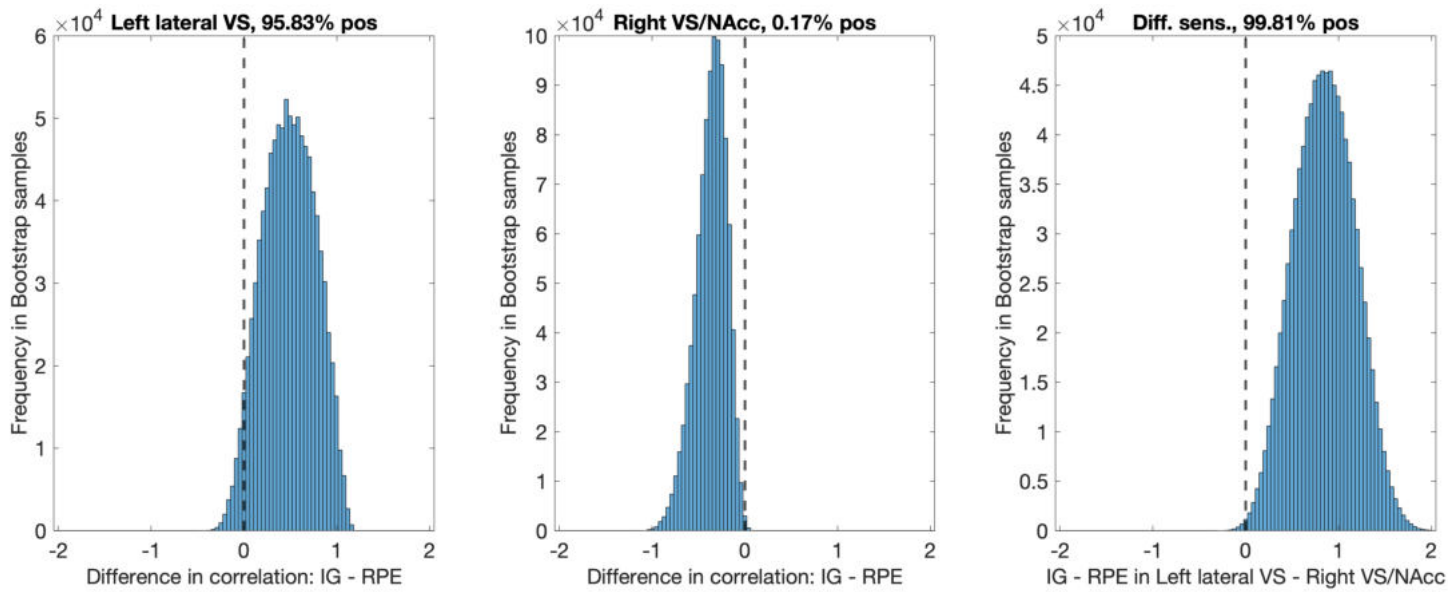


Fig. S6. Histograms from bootstrap sampling of whole dataset within-subject-ranked activations. Left panel shows left lateral VS ROI, correlation to Information Gain (IG) model minus correlation to Reward Prediction Error (RPE) model. Center panel shows right VS/NAcc ROI, correlation to IG minus correlation to RPE model. Right panel shows differential sensitivity to IG minus RPE, subtracting the right VS/NAcc (IG - RPE) score minus the left lateral VS (IG - RPE) score.

Table S1. Information Expectation HI >LO

Brain region	Nr. of voxels	peak Z-value	x	y	z
L dorsal central/postcentral sulcus	660	3.51	-26	-30	50
cerebellar vermis	270	3.66	4	-78	-8
L nucleus accumbens (NAcc)	172	3.82	-10	8	-12
R nucleus accumbens/anterior cingulate	128	3.41	2	20	-10
L paracentral lobule	119	3.33	-2	-14	64
L anterior superior frontal gyrus	98	3.16	-10	58	36
L superior frontal sulcus/gyrus	95	3.52	-22	12	64
R dorsal central sulcus	86	3.5	22	-28	70
L superior temporal sulcus	69	3.32	-50	-14	-4
L posterior putamen	52	3.15	-30	-8	-10
R superior occipital gyrus	48	3.67	30	-90	28
L superior frontal sulcus	43	3.05	-26	26	34
L posterior parahippocampal gyrus	40	3.13	-12	-44	-8
R superior temporal sulcus	29	2.91	56	-12	-10
R cerebellum	25	2.88	20	-50	-28

Average MNI coordinates (in mm) for brain regions showing stronger activations for High than Low Information Expectation (see also Fig. 2). All activations $p < 0.005$, $k = 25$. Abbreviations: HI=High information; LO=Low information; L=left; R=right.

Note: No activations survived the opposite contrast, InfoExpect LO>HI at this threshold. For uncorrected $\alpha = 0.05$ activations, see Fig. S5A.

Table S2. Information Revelation & Outcome Anticipation HI > LO

Brain region	Nr. of voxels	peak Z-value	MNI coordinates		
			x	y	z
L MPFC	296	4.14	-10	50	6
L posterior cingulate gyrus	252	3.85	-2	-50	22
L middle frontal gyrus/ SFS	249	3.41	-34	34	44
R posterior cingulate gyrus	128	4.13	2	-50	22
L posterior insula	175	3.27	-46	-18	18
R MPFC	173	3.32	10	50	8
R cerebellar vermis	142	3.85	-4	-54	-2
L caudate	120	3.26	-16	28	8
L posterior insula	30	3.01	-42	-32	4
L anterior cingulate gyrus	27	3.08	-8	36	-12
R superior temporal sulcus	26	3.03	52	-4	-16
R precuneus/ post. cingulate sulcus	25	3.44	16	-42	44

Information Revelation & Outcome Anticipation LO > HI

Brain region	Nr. of voxels	peak Z-value	MNI coordinates		
			x	y	z
R posterior IFS	43	3.59	54	26	30
L anterior IFG	25	3.26	-48	46	6

RevealAnticip HI>LO; RevealAnticip LO>HI

Average MNI coordinates for brain regions showing stronger activations for RevealAnticip_HI>LO as well as the reverse contrast (Figs 2B and S5B, respectively). All activations survive $\alpha < 0.005$, $k=25$.

Table S3. Feedback POS>NEG; Feedback NEG>POS**Feedback POS>NEG**

Brain region	Nr. of voxels	peak Z-value	MNI coordinates		
			x	y	z
L posterior cingulate gyrus	1025*	5.09	-16	-48	10
R posterior cingulate gyrus, extending onto lingual gyrus	722*	4.51	18	-46	14
L dorsal central sulcus, extending onto postcentral gyrus	615	4.64	-18	-30	58
L fusiform gyrus	550	4.76	-36	-82	-12
R dorsal central sulcus, extending onto postcentral gyrus/sulcus	514*	3.82	28	-28	58
R fusiform gyrus	378*	5.14	30	-78	-14
R calcarine fissure	362*	3.94	20	-78	12
R superior occipital gyrus	147*	3.7	24	-86	26
L frontal pole	138	4.05	-14	64	-6
L inferior frontopolar gyrus	135	3.32	-40	62	-4
L superior precuneus	125	3.99	-4	-62	58
L orbitofrontal cortex (arcuate orbital sulcus)	84	3.66	-34	38	-6
L nucleus accumbens	68*	3.59	-14	-2	-14
R posterior insula	55*	3.68	38	-18	20
R MPFC	50*	3.67	4	74	8
L superior occipital gyrus	41*	3.41	-12	-90	22
L hippocampus	40	3.14	-30	-18	-14
R nucleus accumbens	45	2.97	6	8	-10
R hippocampus	38	2.87	28	-22	-22
R post. cingulate sulcus	34	3.01	6	-48	68

* threshold raised above $z=3.1$ to separate large activations into separate clusters; activations merge into one large cluster at $z=2.6$

Feedback NEG>POS

Brain region	Nr. of voxels	peak Z-value	MNI coordinates		
			x	y	z
R angular gyrus (inferior parietal lobule)	1280	5.08	52	-46	38
R medial SFG, extending into cingulate sulcus/gyrus	909	4.01	6	18	66
L anterior insula	854	4.44	-28	18	-2
R anterior insula	524*	5.2	32	24	-2
R inferior frontal sulcus/MFG	188*	4.22	48	40	28
L supramarginal gyrus	152	3.3	-64	-46	18
R STS	122	3.99	50	-16	-6
R superior frontopolar gyrus/anterior SFS	118	3.26	20	56	28
R MFG (DLPFC)	56*	3.68	40	40	42

* threshold raised above $z=3.1$ to separate large activations into separate clusters; activations merge into one large cluster at $z=2.6$

Average MNI coordinates for brain regions showing stronger BOLD activations for positive Feedback (smile emoticon) than negative Feedback (frown emoticon), as well as the reverse contrast (Feedback_NEG>POS). All activations are $p<0.005$, $k=25$, except where noted.