Supporting Information

O'Reilly et al. 10.1073/pnas.1305373110

SI Methods

Paradigm. Participants completed two 300-trial sessions of the task as described in the main text, with the first one taking place in a behavioral laboratory and the second in the functional MRI (fMRI) scanner. Before each session, participants were given explicit instructions about the meaning of different target colors and the distributions from which target locations were drawn.

The mean of the target location distribution for each run was generated randomly with the constraint that the new mean should be more than 2 SDs from the old mean; the variance randomly took one of four values (SD of 2°, 6°, 14°, or 30°). Actual target locations from the experiment are illustrated in Fig. 1.

Saccadic Reaction Times. We used the rate of gaze displacement to parse eye data into fixations and saccades. We then defined our response measures as follows. Onset time was defined as the end time of the fixation immediately preceding the saccade in which gaze first moved outside a 50-pixel window around the fixation cross (target locations were 300 pixels from the fixation cross). Arrival time was defined as the time of onset of the first fixation within a 50-pixel radius of the target location. Dwell time was defined as the duration of the first fixation within a 50-pixel radius of the target. To remove outliers, onset and arrival times less than 0 ms from the target onset and more than 800 ms (double the average arrival time) were disregarded; values more than 3 SDs from the mean for each measure were also discarded.

Pupillometry. For the general linear model (GLM) analyses of pupil size, pupil data were divided into epochs beginning 600 ms before the onset of the warning cue and ending 2,000 ms after the target onset. Each epoch was baseline-corrected by subtracting the mean pupil area in the first 200 ms of the epoch. We used a multiple regression on the pupil size at each time point to determine the effects of four regressors: the main effect (which captured target-related changes in pupil diameter over time across all trial types), the entropy of the model, surprise [Shannon information (I_S)], and updating [Kullback–Leibler divergence (D_{KL})]. We plotted the mean effect size (beta value) across subjects ± SEM. Effects were considered significant for time points when the Z-score of the mean (mean/SEM) was greater than ±2.3, corresponding to a P value of 0.01 uncorrected.

For the raw pupil analysis presented in Fig. S2, each epoch was baseline-corrected by dividing by the mean pupil area in the first 200 ms of the epoch.

Whole-Brain fMRI Analysis. fMRI data processing and preprocessing were carried out using Centre for Functional MRI of the Brain (FMRIB) Expert Analysis Tool (FEAT), version 6.0, which is part of the FMRIB's Software Library (FSL; www.fmrib.ox.ac.uk/fsl). The following prestatistics processing was applied: motion correction using MCFLIRT (Motion Correction with FMRIB Linear Image Registration Tool), slice timing correction using Fourier-space time-series phase-shifting, nonbrain removal using BET (FMRIB Brain Extraction Tool), spatial smoothing using an 8.0-mm FWHM Gaussian kernel, grand-mean intensity normalization of the entire 4D dataset by a single multiplicative factor, and high-pass temporal filtering (Gaussian-weighted, least-squares, straight-line fitting; $\sigma = 50$ s).

At the first level (individual subject analysis), time-series statistical analysis was carried out using FILM with local autocorrelation correction. A GLM analysis was carried out with four task regressors (main effect of task, entropy of the model, I_S, and D_{KL}) as described above, plus six regressors of no interest capturing the effects of head motion. Z-statistics (Gaussianized t statistics) were passed to the second (group)-level analysis.

At the group level, individual participants' Z-statistic images were normalized to the Montreal Neurological Institute brain using nonlinear registration [FNRIT (FMRIB Non-linear Image Registration Tool)] and entered into a random effects analysis in which the average effect of each regressor across the group was tested against zero. Higher level analysis was carried out using the FMRIB's Local Analysis of Mixed Effects (FLAME) stage 1 with automatic outlier detections and deweighting.

Group statistics were corrected for multiple comparisons using a cluster size-based approach (as implemented in the FSL), with a cluster-forming threshold of $Z > 2.3$ and a cluster significance threshold of $P < 0.05$ (corrected).

Additional Pupillometry Experiment. To replicate the unexpected finding that updating negatively modulates pupil diameter, we conducted an additional behavioral experiment with 18 participants to check potential confounding factors that could have driven the pupillometric effect. The additional experiment was similar in design to the main experiment, except that instead of making a saccade to a colored dot on a circular perimeter as in the main experiment, participants maintained central fixation while a Gabor grating stimulus was presented on the circular perimeter. The task was to discriminate the grating's orientation. Trial types (expected, one-off, and update) were indicated by the color of a warning stimulus presented 500 ms before the Gabor patches.

This additional experiment addresses two potential concerns about the pupillometric results: first, that the effect could be driven by the different colors of the dots on update and one-off trials (because all targets were isoluminant Gabor patches in the additional experiment) and, second, that the effect could be contaminated by eye movements (because there were no eye movements in the additional experiment). We observed a very similar pattern of opposite effects of surprise and updating as in the main experiment (Fig. S3*C*).

Although this additional experiment demonstrates that the pupillometric effects reported in the main text were not an artifact of target color or eye movements, it is worth noting that even within the main experiment, these confounds were not likely to have driven the observed effects because:

i) In the original experiment, although the target dots were different colors on one-off and update trials, the targets on update trials were the same colors as the targets on expected trials (i.e., if an update trial had a red target, all the subsequent expected trials in that run also had red targets). Therefore, any small luminance differences could not account for the difference between update and expected trials shown in Fig. S2*B*.

ii) Although participants did make eye movements during the task, these eye movements were complete by the time point at which the effect of updating became significant (Fig. 4).

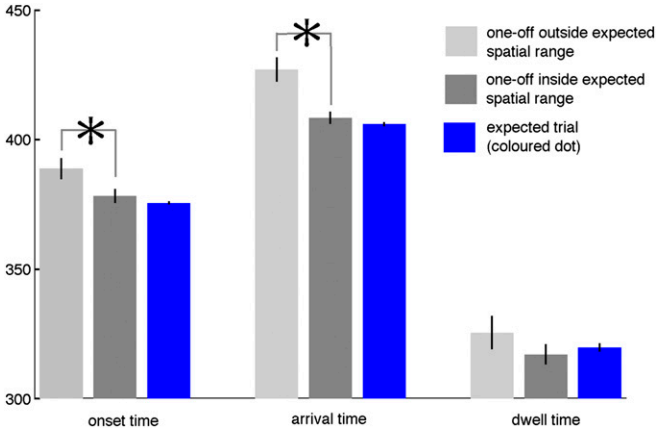

Fig. S1. Behavior on one-off trials is determined by spatial surprise, not feature-based surprise. One-off trial targets could fall anywhere on the circle; hence, they sometimes fell within the expected spatial range by chance. We divided the one-off trials into those with locations within 2 SDs of the mean of the expected distribution and those outside that range. Onset times were significantly slower on one-off trials outside the expected spatial range than inside it ($t_{16} = 2.53$, $P = 0.011$, paired samples t test for the group of 18 participants comparing mean onset time in each condition within subjects), but there was no evidence of a difference between one-off trials within the expected spatial range and expected trials ($t_{16} = 0.11$, $P = 0.46$). Similarly, arrival times were significantly slower on one-off trials outside the expected spatial range than on one-off trials within the expected spatial range ($t_{16} = 3.63$, $P = 0.0010$), but there was no evidence of a difference between one-off trials within the expected spatial range and expected trials ($t_{16} = 0.15$, $P = 0.44$). There was no evidence for a difference between conditions in dwell time (t_{16}= 0.37, $P = 0.36$; $t_{16} = 0.016$, $P = 0.51$); note that in the main behavioral analysis, dwell times were not reported to be longer on one-off trials than on expected trials. In summary, behavior on the one-off trials that fell within the expected spatial range was similar to that on expected trials, whereas one-off trials outside the expected spatial range showed slower responses. This suggests that behavioral costs were determined by spatial reorienting/surprise rather than the nonspatial features (e.g., color) of the targets. *$P < 0.05$.

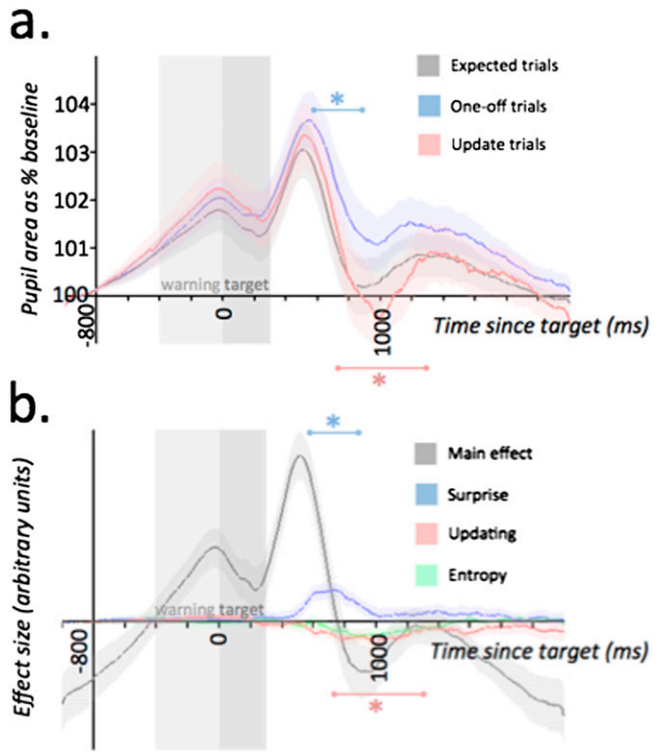

Fig. S2. Raw pupil response on update, one-off, and regular trials. (*A*) Data on which the GLM analysis in Fig. 4 is based: the pupil area on each of three trial types (update, one-off, and expected). Pupil area is expressed as a percentage (%) of the area on a trial-by-trial baseline, made up of the first 200 ms of the trial (this was the same baseline correction approach used in the GLM analysis). The data lines are the mean pupil area across participants for each trial type, and the shaded surrounds are the SEM. The gray-shaded boxes indicate the period in which the warning signal and the target dot were on the screen. The blue and red lines with an asterisk indicate the period in which the GLM effects of surprise and updating were significant, and they are shown in this figure to aid comparison with Fig. 4. There is a large pupil response to the presentation of the warning signal and the target that ramps up before the warning signal (as in the behavioral session, the timing of stimuli was predictable). Following the target, there is a differential response to update and one-off trials. (*B*) GLM results show all four regressors (i.e. the red and blue lines for effect of surprise and updating are the same as the light- and dark-grey lines indicating the effects of surprise and updating in Fig. 4). Note that the "main effect," target-evoked effects that do not vary across trial types or with the computational regressors, dwarf the computationally specific effects. The fact that these nonspecific evoked effects are modeled out makes the GLM approach particularly sensitive.

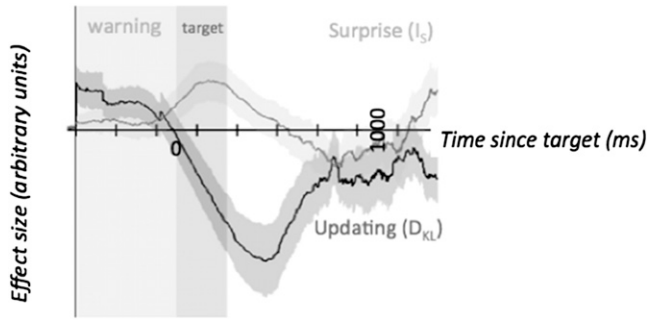

Fig. S3. Pupil area effects from a separate behavioral experiment. Pupil effects for updating and surprise in an additional behavioral experiment, as reported in the main text, are shown. Note that the pattern of effects is similar to those seen in the main experiment, with opposite effect directions for surprise and updating and an earlier time of effect for surprise compared with updating.

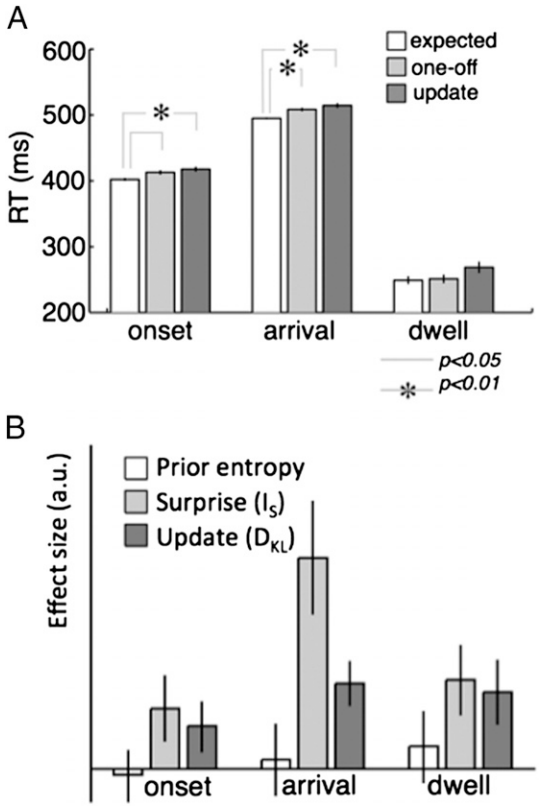

Fig. S4. Behavioral eye data from the fMRI session. (*A*) Breakdown of saccadic reaction times (RTs) into saccade onset, arrival, and dwell times for data recorded during the fMRI session. These are defined in *Methods*. The plotting conventions are exactly as in Fig. 2; this figure shows that the saccadic behavior in the fMRI session was similar to that recorded in the separate eye-tracking session. Onset time was significantly longer on both update trials ($t_{16} = 5.8$, $P = 1.1 \times 10^{-5}$, paired samples *t* test) and one-off trials ($t_{16} = 5.3$, $P = 2.8 \times 10^{-5}$) compared with expected trials; there was no significant difference in onset time between update and one-off trials ($t_{16} = 0.94$, $P = 0.18$). Arrival time was significantly longer on both update trials ($t_{16} = 3.7$, $P = 8.2 \times 10^{-4}$) and one-off trials ($t_{16} = 2.6$, $P = 0.010$) compared with expected trials; there was no significant difference in arrival time between update and one-off trials ($t_{16} = 0.46$, $P = 0.32$). Dwell time was significantly longer in the update trials compared with both the one-off trials ($t_{16} = 2.7$, $P = 0.0073$, paired samples *t* test) and expected trials ($t_{16} = 3.5$, $P = 0.0013$). (*B*) Results of a GLM on saccadic onset, arrival, and dwell times for data recorded during the fMRI session. The analysis and plotting conventions are as in Fig. 3.

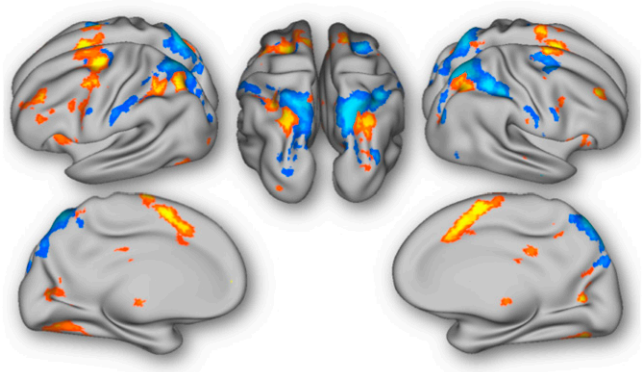

Fig. S5. Uncorrected whole-brain activity patterns for surprise (I_S) and updating (D_{KL}). These maps are not corrected for multiple comparisons, and hence are shown only to give an impression of the overall pattern of activity that underlies the multiple comparisons-corrected effects reported in the main text. The maps are thresholded at $Z > 2.0$ uncorrected. Blue colors indicate surprise, and red/yellow colors indicate updating.

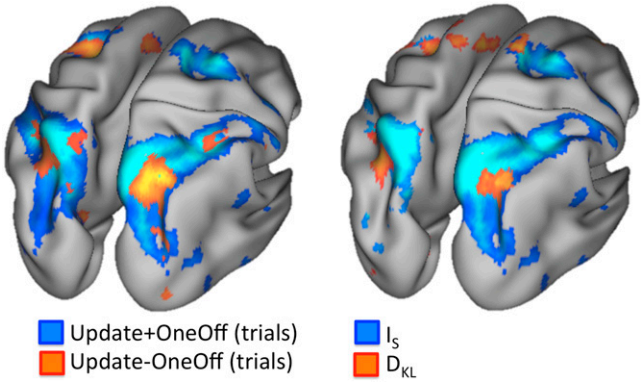

Fig. S6. fMRI analysis on trial type (update ± one-off trials). Although the parametric regressors I_S and D_{KL} provide a specific model to compare against brain activity, it is useful to check that the activations evoked by different trial types (update and one-off trials) are consistent with these predictions. Because surprise/behavioral reorienting is present on both trial types but updating is present only on update trials, the relevant contrasts are (update trials + one-off trials) and (update trials − one-off trials). Maps are thresholded at $P < 0.01$ uncorrected to illustrate the similarities between maps for the different contrasts; note that only corrected statistics are presented in the main text.

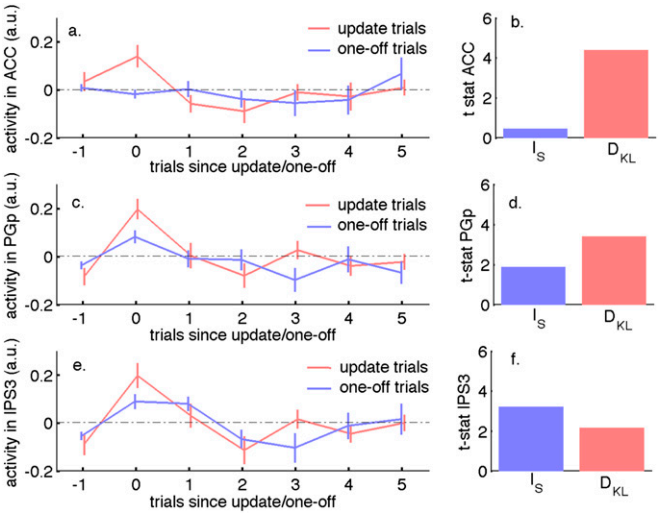

Fig. S7. Activity in cingulate and parietal regions of interest (ROIs). (*A*, *C*, and *E*) Raw activity level on update and one-off trials and the trials around update and one-off trials. Trials are numbered so that the update or one-off trial is 0, the preceding trial is −1, and trials after the update/one-off trial are numbered + 1, +2, etc. Note that in both parietal ROIs [parietal region PGp and interparietal sulcus 3 (IPS3)], activity on both update trials and one-off trials is above baseline (baseline is the average activity across expected trials). In the anterior cingulate cortex (ACC), activity is above baseline on update trials but not on one-off trials. (*B*, *D*, and *F*) Group *t*-statistics for I_S and D_{KL} in each ROI, from a GLM analysis of ROI activity. Note that there is a parametric effect of surprise/behavioral reorienting (I_S) in both parietal ROIs but not in the ACC. In all three ROIs, there is an effect of updating (D_{KL}). Within the parietal cortex, IPS3 is more activated in relation to the parametric effect of surprise/behavioral reorienting, whereas PGp is more activated in relation to the parametric effect of updating. a.u., arbitrary units.

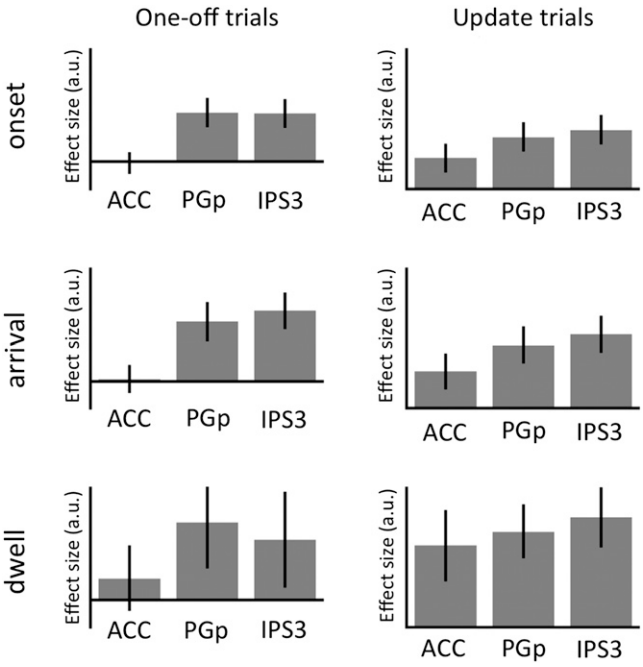

Fig. S8. Regression of onset, arrival, and dwell times on activity in the ACC, IPS3, and PGp. Bars show the effect size (beta value) ± SEM from a multiple regression in which activity in the three ROIs (ACC, PGp, and IPS3) was used to predict each behavioral measure, as in Fig. 7. Note that the parietal ROIs are more strongly associated with all behavioral measures; the relationship between ACC activity and behavior does not reach statistical significance in any case.

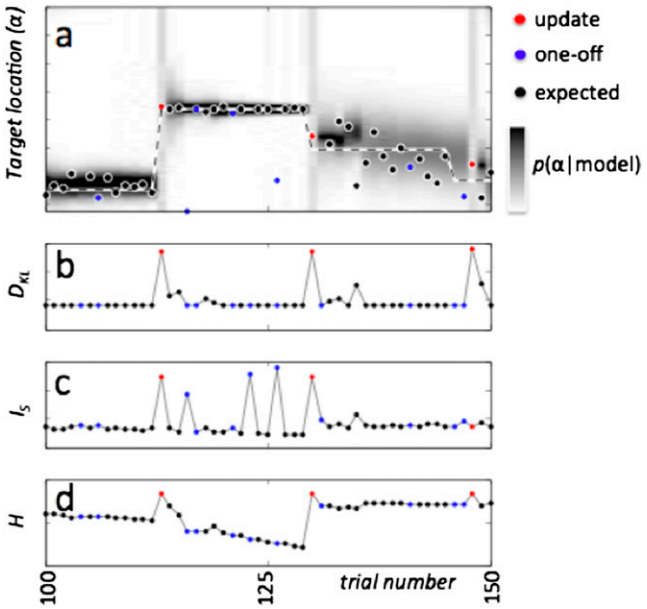

Fig. S9. Model and regressors. All panels show the data from 50 trials, with trials on the *x* axis. Trial types are indicated by dot color as in the key. (*A*) Plot of the state of the learning model over 50 trials. On the *y* axis is target location (α = 0–360°). The dashed line indicates the mean of the actual (generative) Gaussian from which expected and update trials were drawn. Dots indicate the actual values of target location on each trial. Shading indicates the probability that a target will appear at the location, given the state of the model on that trial. Note that the distribution of probability is broadest at the start of each run and that it is broader when the generative distribution has a higher variance [increased spread of data points about the mean (e.g., toward the right of the plot)]. (*B–D*) Computational regressors used in data analysis and based on the model in *A*. The equations by which these regressors were derived are given in *Methods*. Again, the colored dots are the actual values (of the regressors) on each trial. (*B*) Update regressor: D_{KL} between the state of the model before and after observing the target location on each trial. (*C*) Surprise regressor: I_S of each target location, given the model. (*D*) Entropy (H) of the model is roughly equivalent to the spread of probability density over possible values of α, as shown by the spread of the shaded area in *A*.

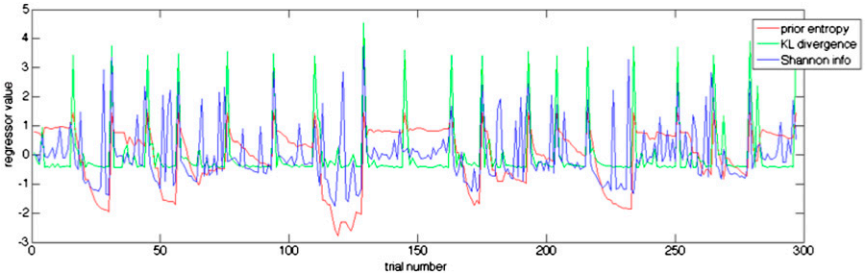

Fig. S10. Regressors over all 300 trials. This plot illustrates the values of the three information theoretic regressors (prior entropy, I_S, and D_{KL}) over all 300 trials of the task (this was one of three different schedules used). Each regressor was normalized such that its mean value across all 300 trials was 0 and its variance across all 300 trials was 1. Note that although there is some correlation between regressors (e.g., on update trials, I_S, D_{KL}, and entropy are all high), the regressors were not strongly correlated overall. The shared variance (R^2) between prior entropy and I_S was 0.14, that between prior entropy and D_{KL} was 0.14, and that between D_{KL} divergence and I_S was 0.15.