

Supplemental data

Figure S1, related to Fig. 2: Supplemental behavioral analyses.

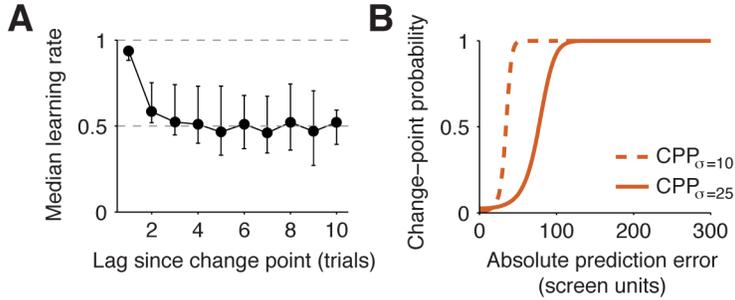


Figure S1. Additional behavioral analyses further supported our conclusion that participants flexibly adapted their learning rates as predicted by the normative model.

A: Direct estimation of trial-wise learning rate (LR). In addition to the regression-based behavioral analyses described in the main text, we also performed descriptive analyses directly on empirical trial-wise learning rate (cf. Nassar et al., 2010). The learning rate on trial t can be estimated as the subsequent update in bucket position ($B_{t+1} - B_t$) expressed as a fraction of the current spatial prediction error (δ_t). This estimate will be affected by motor noise in joystick placement and will be less precise when δ_t is small; we therefore omitted trials with $\delta_t < 5$ screen units (see *Experimental Procedures*).

The overall trial-wise learning rate (median of medians from individual participants) was 0.55 (IQR 0.43 to 0.76). Learning rate was greater after change points (median=0.94, IQR 0.88 to 0.95) than otherwise (median=0.51, IQR 0.40 to 0.73; median difference=0.34, IQR 0.15 to 0.46, signed-rank $p < 0.001$), consistent with increased learning from surprisingly large errors. Learning rates also tended to decrease across successive non-change-point trials, mirroring the decrease in belief uncertainty (see Fig. 1C; the median within-participant Spearman correlation between lag and learning rate across non-change-point trials $\rho = -0.11$, IQR -0.19 to 0.00, signed-rank $p < 0.001$). Plotted points and error bars represent median and IQR.

B: Context-sensitive estimation of change-point probability (CPP). Theoretical CPP increases as an approximately sigmoidal function of $|\delta_t|$ but also depends on the level of observation noise. Our experimental task dissociated CPP from physical values of δ_t by manipulating the noise level across runs. A moderate-sized prediction error might imply high CPP if it occurred in a low-noise context (dashed line) but low CPP if it occurred in a high-noise context (solid line). We used additional regression-based analyses to demonstrate that behavior showed the expected noise sensitivity.

For these analyses, we estimated approximate CPP as a function of δ_t (holding RU fixed at its median value for simplicity), separately for the low-noise context ($CPP_{\sigma=10}[\delta_t]$) and the high-noise context ($CPP_{\sigma=25}[\delta_t]$). We then defined a new participant-wise regression model with 3 terms of interest. The first regressor was based on the average of these functions: $\delta_t \times [CPP_{\sigma=10}(\delta_t) + CPP_{\sigma=25}(\delta_t)]/2$. Coefficients were significantly positive for this term (median=0.48, IQR 0.28 to 0.67, signed-rank

$p < 0.001$), reflecting the general tendency for learning rate to rise with $|\delta_t|$. The other two terms were based on the difference between the two functions: $\delta_t \times [\text{CPP}_{\sigma=10}(\delta_t) - \text{CPP}_{\sigma=25}(\delta_t)]$, defined separately for low-noise and high-noise runs only. These difference terms have values near zero for small prediction errors (which do not signify a change point in either noise condition) and for large prediction errors (which signify a change point in both conditions). They have positive values for intermediate prediction errors (approximately 25–100 screen units), which are more indicative of a change point in the low-noise condition than in the high-noise condition. The two difference regressors should receive similar coefficients if learning rate depends solely on δ_t and coefficients of opposite sign if learning rate is influenced by the noise context. Our results showed that learning rate demonstrated the appropriate noise sensitivity: coefficients were significantly positive for the second term (median=0.11, IQR 0.05 to 0.20, signed-rank $p < 0.001$) and significantly negative for the third term (median=-0.20, IQR -0.32 to -0.14, signed-rank $p < 0.001$), implying that participants reacted differently to physically equivalent stimuli in the two noise contexts, consistent with normative estimation of CPP.

Figure S2, related to Fig. 4: Theoretical BOLD time courses.

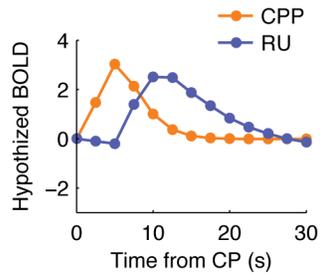


Figure S2. Theoretically predicted BOLD time courses for CPP and RU as a function of time after a change point, derived from the approximately Bayesian model. To create this plot we extracted the mean model-derived CPP and RU for sequences of trials aligned to large change points ($CPP > 0.5$) and convolved these values with a canonical HRF.

Figure S3, related to Fig. 2: Followup-up session behavior.

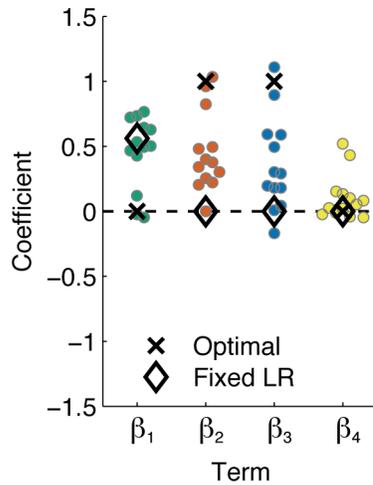


Figure S3. Behavioral results from the follow-up session. A subset of the original participants ($n=13$) was re-tested behaviorally with eye-tracking to assess the association between task variables and saccadic eye movements. Behavioral effects were evaluated using the same regression framework described in the main text (see Fig. 2), and results were comparable to the original data set. Participant-wise coefficients significantly exceeded zero for the fixed-learning-rate term (β_1 ; median=0.50, IQR 0.35 to 0.66, signed-rank $p=0.001$), and for modulation of learning rate by CPP (β_2 ; median=0.38, IQR 0.25 to 0.58, signed-rank $p<0.001$), RU (β_3 ; median=0.29, IQR 0.15 to 0.59, signed-rank $p=0.001$), and reward value (β_4 ; median=0.01, IQR 0 to 0.03, signed-rank $p=0.033$). As in Fig. 2, estimates of β_4 are scaled by a factor of 5 for visibility. Black markers represent simulated data as in Fig. 2. See Fig. S4 for eye-tracking results.

Figure S4, related to Fig. 5: Saccadic eye movements.

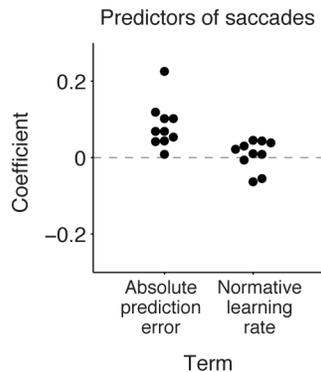


Figure S4. Given the visuospatial nature of our experimental task, a potential concern is that learning might be associated with increased oculomotor activity, which in turn might provide a more proximate explanation for the observed BOLD effects. To address this issue, we conducted a follow-up eye-tracking session outside of the scanner. Eye position was recorded at 60 Hz using a Tobii T60XL eye tracker (Tobii Technology, Stockholm, Sweden). Blinks and other unreliable signals were identified using a custom algorithm and not analyzed. Trials in which less than 50% of the data were deemed reliable were excluded from analysis. Three participants were removed from analysis due to insufficient reliable data (<120 trials). Saccades were identified using a custom algorithm implemented in Matlab. Saccades were counted for each trial and modeled with two GLMs designed to probe any relationship between learning factors and eye movements. Coefficients for terms in each model were estimated using `glmfit` in Matlab with a Poisson link function.

The first model looked for a relationship between eye movements and learning factors using regressors analogous to the primary fMRI GLM: (1) the outcome's left/right position on the screen, (2) model-derived CPP, (3) model-derived RU, (4) reward value, and (5) the difference between the outcome position and the center of the screen. This model also contained two nuisance variables reflecting: (1) the proportion of reliable eye data measured during the outcome phase, and (2) the proportion of reliable eye data measured during the decision phase. Reward value coefficients were slightly negative on average and not significantly different from zero (signed-rank $p=0.105$). In contrast, coefficients were significantly greater than zero for CPP (signed-rank $p=0.002$) and RU (signed-rank $p=0.037$), indicating a potential confound with eye movement.

However, these relationships between eye movements and either CPP or RU were driven by the absolute magnitude of the prediction error, not adaptive learning. We estimated a second GLM that included: (1) absolute prediction error magnitude; and (2) normative learning rates that combined RU and CPP, orthogonalized on absolute prediction-error magnitude. Across participants, the first term tended to be greater than zero (left-hand set of points; signed-rank $p=0.002$), but the second term, reflecting normative learning, did not differ from zero (right-hand set of points; $p=0.492$).

Moreover, this relationship between saccadic activity and the absolute magnitude of the prediction error could not by itself explain the relationships between BOLD signals

and adaptive learning. We conducted a corresponding fMRI analysis in which the trial-related BOLD response was modulated by: (1) the absolute value of the prediction error; and (2) normative learning rates derived from the approximate Bayesian model, orthogonalized on absolute prediction-error magnitude. Treating the seven common adaptive learning-rate regions as regions of interest (ROIs), each showed significant effects for both terms (all $ps < 0.01$ after Bonferroni correction for 7 comparisons). This result suggests that the observed BOLD effects in these regions were not attributable solely to a simple factor such as increased oculomotor activity for large errors.

Figure S5, related to Fig. 4: Functional connectivity.

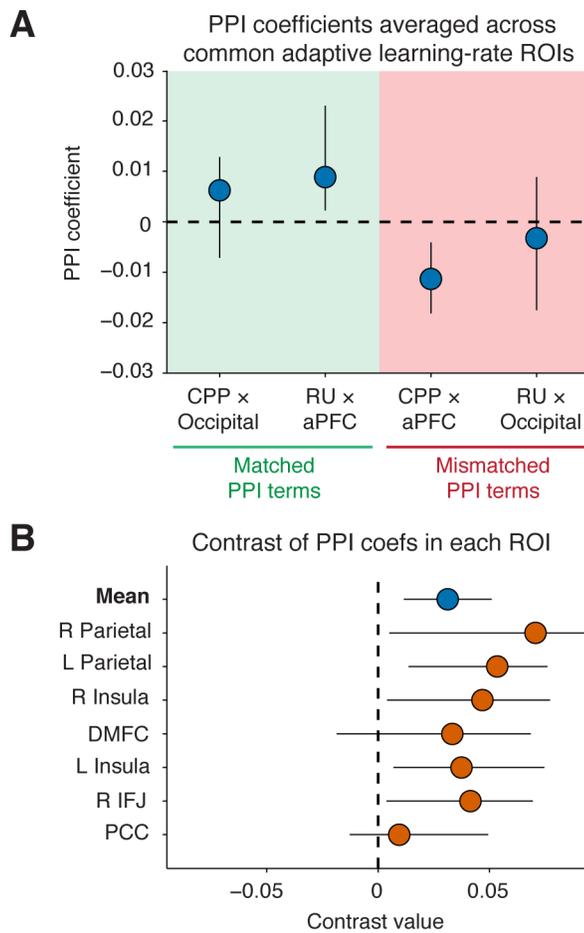


Figure S5. Factor-specific regions selectively explain variance in common adaptive learning-rate regions (LR regions). To test whether regions reflecting a specific computational factor might communicate with LR regions according to the extent to which the represented factor is driving learning on a given trial we conducted a psychophysiological interaction (PPI) analysis. In particular, we designed a regression model to examine whether physiological variables reflecting CPP and RU might be more related to LR regions when the reflected variable should be driving more of the learning on a given trial. The psychological variables were trial-by-trial estimates of RU and CPP estimated by our computational model. The physiological variables were trial-wise BOLD amplitudes extracted from the occipital cluster identified as a CPP region and the right aPFC cluster identified as an RU region (Fig. 4). The two physiological and two psychological variables were used to create 4 PPI variables from the interactions of z-scored psychological and physiological terms (CPP×Occipital, CPP×aPFC, RU×Occipital, RU×aPFC). Two of these PPI terms were “matched,” corresponding to the interaction between a learning factor and the area reflecting that factor (CPP×Occipital,

RU×aPFC). The other two were “mismatched”, corresponding to the interaction between a learning factor and an area representing the other factor (CPP×aPFC, RU×Occipital). We hypothesized that matched interaction terms would take positive coefficients indicating an increase in shared variance between factor-specific and LR regions when the factor was contributing to learning. We hypothesized that mismatched terms would take smaller or even negative values, which would represent a decrease in shared variance between factor-specific and LR regions when that factor contributed less to the prescribed learning rate.

Psychological, physiological, and PPI variables were included in a regression model that also included an intercept term and several other nuisance variables (outcome value, outcome location, outcome distance from center of screen, signed prediction error). This model was applied to the series of trial-wise BOLD amplitudes extracted from each of the LR regions (Fig. 5).

A: Individual PPI coefficients. PPI coefficients were averaged across the seven LR regions for each participant. The resulting coefficient for each term is plotted on the ordinate (points = across-participant medians, lines = bootstrapped 95% CIs). Terms on the left half of the abscissa (shaded green) are based on matched psychological and physiological variables, and terms on the right half of the abscissa (shaded red) are based on mismatched psychological and physiological variables. The contrast of matched versus mismatched terms was significant (median=0.031, IQR 0.001 to 0.068, signed-rank $p=0.002$). While all PPI coefficients followed the predicted pattern, the trends were stronger in the aPFC PPI terms (RU×aPFC: median=0.009, IQR -0.010 to 0.031, signed-rank $p=0.043$; CPP×aPFC: median=-0.011, IQR -0.021 to -0.001, signed-rank $p=0.002$).

The overall contrast result was not significant in alternative versions of the PPI analysis that substituted other RU-selective regions in place of right aPFC. Results were non-significant for the RU-selective regions in parietal cortex (median=0.018, IQR -0.030 to 0.042; signed-rank $p=0.184$) and cerebellum (median=0.033, IQR -0.042 to 0.094; signed-rank $p=0.067$). Comparing the effect across RU-selective regions yielded modest evidence that the effect was greater for aPFC than the parietal region (signed-rank $p=0.047$) but no evidence of regional specificity in the other two pairwise comparisons.

B: Contrast of matched versus mismatched PPI terms in individual LR regions. Points and lines represent across-participant medians with bootstrapped 95% CIs. Contrast results for individual LR regions are plotted in brown and the aggregate result across LR regions (reported above and in the main text) is plotted in blue. The contrast was significantly greater than zero in the majority of individual regions (signed-rank $p<0.05$, uncorrected, for all LR regions except DMFC [$p=0.197$] and PCC [$p=0.304$]). Only the right parietal region was significant after Bonferroni correction for 7 comparisons (corrected $p=0.009$).

Movie S1, related to Fig. 4: Change-point-aligned BOLD time courses.

Supplemental file: postCP_10fps.avi

Movie S1. Animated surface-rendering of change-point-aligned BOLD time courses. Time courses were estimated for each participant using a finite-impulse-response model in the framework of a GLM. The GLM included the same baseline and single-trial nuisance regressors as the primary analysis (see *Experimental procedures*) plus a total of 26 piecewise linear spline (“tent”) basis functions corresponding to points in the event-related time series. Basis functions were centered every 2.5s from 0 to 30s after change points; a second set of basis functions was placed at the same lags relative to non-change-points. To estimate change-point-related effects we took the difference between each participant’s change-point-aligned and non-change-point-aligned timecourses. The animation displays a simple estimate of across-participant effect size (ES; mean/SD) at each location and time point. Results are shown on an inflated brain and spatiotemporally interpolated for visualization. The animation was created using Freesurfer (<http://surfer.nmr.mgh.harvard.edu>) and Paraview (<http://www.paraview.org>).

Table S1, related to Fig. 3. BOLD effects for individual GLM terms.

#Voxels	Region	Peak <i>t</i>	Peak <i>x</i>	Peak <i>y</i>	Peak <i>z</i>
<i>Change-point probability: Positive effects</i>					
6137	R parieto-occipital sulcus	12.83	21	-60	24
	L parieto-occipital sulcus	11.19	-18	-63	21
	R inferior temporal cortex	12.10	30	-48	-9
	L inferior temporal cortex	9.94	-27	-48	-9
	R lateral occipital cortex	10.76	33	-78	24
	L lateral occipital cortex	8.65	-27	-84	33
	posterior cingulate cortex	8.06	0	-27	27
1230	R superior frontal gyrus	8.96	24	3	51
	DMFC	6.85	0	18	42
	R dorsolateral PFC	5.88	30	36	45
396	L superior frontal gyrus	8.78	-24	6	57
142	R anterior insula	8.71	33	21	9
105	L dorsolateral PFC	5.36	-36	36	36
103	L anterior insula	8.39	-33	18	9
61	L intraparietal sulcus	6.21	-33	-48	39
<i>Change-point probability: Negative effects</i>					
889	R posterior peri-ventricular	-7.29	24	-45	18
	L medial temporal lobe	-6.91	-21	-18	-15
	R medial temporal lobe	-7.20	18	-12	-15
	L ventral striatum	-6.52	-9	9	-6
	R ventral striatum	-6.14	9	12	-3
250	L inferior frontal gyrus	-8.18	-51	33	6
249	L superior temporal sulcus	-6.24	-51	-39	0
246	R superior temporal gyrus	-6.69	63	-3	6
231	L posterior peri-ventricular	-6.65	-21	-48	15
166	R cerebellum	-5.20	15	-72	-45
62	L lateral occipital cortex	-6.61	-30	-96	-9
48	L superior PFC	-5.10	-15	36	54
47	R lateral occipital cortex	-6.20	33	-93	-6
<i>Relative uncertainty: Positive effects</i>					
7293	posterior parietal cortex	11.56	9	-72	57
	L intraparietal sulcus	10.55	-33	-48	39
	R intraparietal sulcus	9.80	33	-45	42
	L cerebellum	10.90	-33	-45	-39
	R cerebellum	10.17	30	-45	-48
	L lateral occipital cortex	6.56	-45	-63	-3
	R lateral occipital cortex	8.22	54	-60	-9
1613	R superior frontal gyrus	11.13	30	3	66
	R anterior PFC	7.70	33	57	24
	R anterior insula	8.27	33	24	6
	R inferior frontal junction	7.69	54	15	33
315	L superior frontal gyrus	7.96	-27	3	63

174	R thalamus	7.48	18	-30	12
123	DMFC	6.86	3	30	45
116	L thalamus	6.00	-12	-18	9
110	L dorsolateral PFC	6.04	-39	36	36
84	L inferior frontal junction	5.81	-48	6	18
73	L anterior insula	6.97	-30	18	9
46	posterior cingulate cortex	7.26	3	-30	27

Relative uncertainty: Negative effects

4250	L medial temporal lobe	-12.38	-21	-21	-15
	ventromedial PFC	-9.47	0	60	0
	R posterior peri-ventricular	-10.60	21	-39	21
	L posterior peri-ventricular	-10.26	-18	-42	18
	R medial temporal lobe	-10.73	24	-21	-15
	L superior temporal sulcus	-7.94	-51	-39	0
	posterior cingulate cortex	-8.06	-9	-54	30
	subcallosal gyrus	-9.26	0	9	-9
	R posterior insula	-7.72	39	-15	21
140	L posterior insula	-7.70	-39	-15	18
139	R central sulcus	-5.42	27	-27	57
131	L lateral occipitoparietal	-6.53	-54	-63	24
69	R cerebellum	-5.20	30	-81	-36

Reward value: Positive effects

666	R intraparietal sulcus	6.11	54	-30	54
143	DMFC	5.37	3	27	45
115	L lateral occipitoparietal	5.13	-30	-93	15
109	R inferior frontal gyrus	6.24	51	9	24
102	R anterior PFC	5.51	45	48	9
66	R lateral occipital	4.71	48	-72	-15
66	R ventral striatum	7.79	9	12	0
61	R anterior insula	6.49	33	21	6
51	L ventral striatum	7.15	-6	12	0
51	posterior cingulate cortex	5.05	3	-27	27
45	brainstem	6.02	3	-18	-18
41	L anterior insula	5.74	-33	21	0

Reward value: Negative effects

93	R posterior insula	-5.58	48	-6	-6
40	R posterior insula	-5.04	42	-15	18

Residual update: Positive effects

70	anterior cingulate cortex	4.74	0	33	21
44	DMFC	5.55	3	18	66