

Supplementary Information for: A unique adolescent response to reward prediction errors

Jessica R. Cohen^{1,4}, Robert F. Asarnow², Fred W. Sabb^{2,3}, Robert M. Bilder^{1,2,3}, Susan Y. Bookheimer^{1,2,3}, Barbara J. Knowlton^{1,3} & Russell A. Poldrack^{1,2,3,5}

¹Department of Psychology, University of California, Los Angeles

²Department of Psychiatry & Biobehavioral Sciences, University of California, Los Angeles

³Brain Research Institute, University of California, Los Angeles

⁴Current Affiliation: University of California, Berkeley

⁵Current Affiliation: University of Texas at Austin

Overall accuracy and response times by group

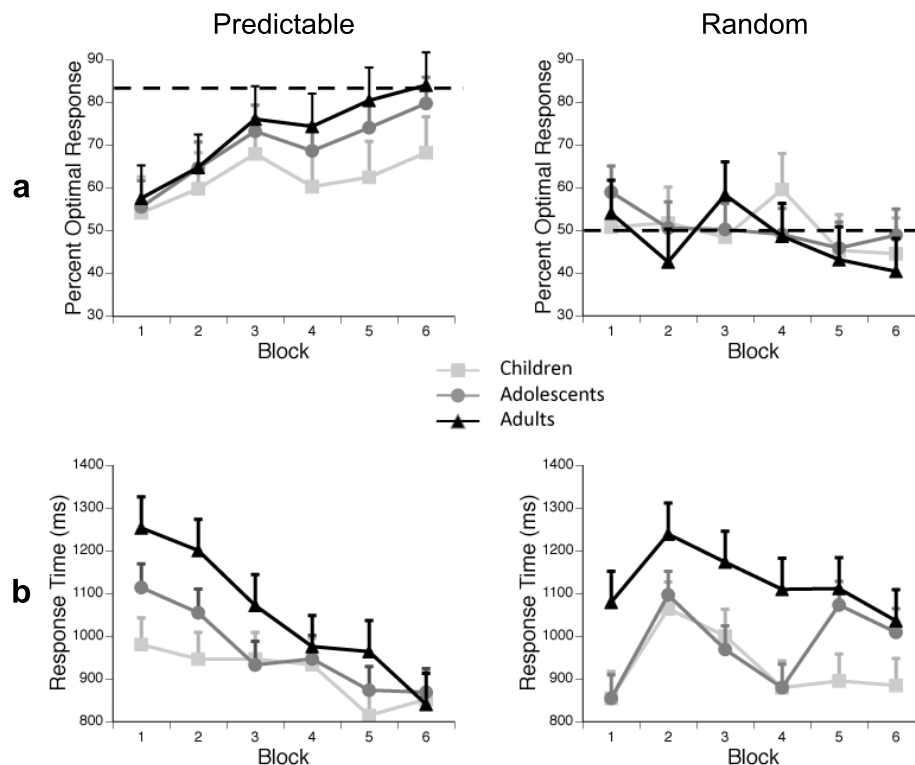


Figure 1: Behavioral performance. (a) Accuracy, or percent optimal response, across blocks of the task for predictable and random stimuli. Dashed lines are the actual probabilities associated with a specific outcome (83% for predictable stimuli; 50% for random stimuli). (b) Response time across blocks of the task for predictable and random stimuli. All error bars indicate s.e.m.

Relationship between decision value, prediction error, and learning rate (alpha)

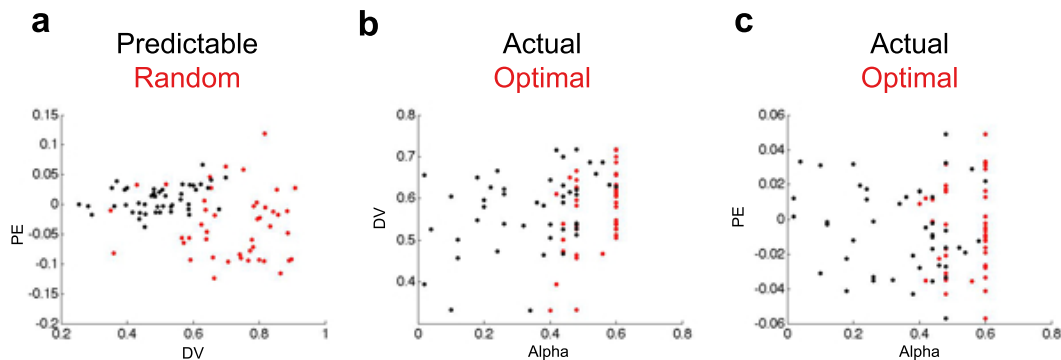


Figure 2: Plots depicting the relationship between decision value, prediction error, and learning rate (alpha). (a) The relationship between the average decision value and prediction error of each participant. There was a positive linear relationship between decision value and prediction error for predictable stimuli (in black; $r = 0.36$, $P = 0.01$) but not for random stimuli (in red; $r = -0.04$, $P = 0.78$). We also examined the trial-by-trial relationship between decision value and prediction error separately for each participant and found no consistent relationship between the two variables at the level of individual trials. (b) The relationship between decision value and both the optimal and actual alphas. (c) The relationship between prediction error and both the optimal and actual alphas. The optimal alphas were used in the model. When examining the difference between the optimal alpha (maximum rate within 10% of the best fit) and the actual alpha (absolute best fit), we found that they were highly correlated ($r = 0.49$, $P < 0.001$). Additionally, we found that decision value was similarly correlated with both the optimal alpha (in red; $r = 0.46$, $P = 0.001$) and the actual alpha (in black; $r = 0.36$, $P = 0.01$) and that prediction error was not correlated with either ($r = 0.05$, $P = 0.76$ for the optimal alpha and $r = -0.12$, $P = 0.41$ for the actual alpha). There was not a significant difference between predictable and random stimuli, thus we collapsed across them for this analysis.

Whole-brain main effects by group

Here we displayed the whole-brain responses for each age group to stimulus onset as compared to baseline (Supplementary Fig. 3; Supplementary Table 2) and positive as compared to negative feedback (Supplementary Fig. 4; Supplementary Table 3), separated by age group. The neural response to stimulus onset included neural regions commonly associated with probabilistic learning^{1,2}, including the left middle frontal gyrus and bilateral anterior cingulate gyrus, paracingulate gyrus, pre-supplementary motor area, supplementary motor area, insula, dorsal and ventral striatum, thalamus, midbrain, supramarginal gyrus, and occipital cortex, although to different degrees for the different age groups (Supplementary Fig. 3). The pattern of neural activity appeared to be changing linearly with age.

When comparing the neural response to positive as compared to negative feedback, a set of regions commonly associated with receiving rewards was active bilaterally, such as the medial prefrontal cortex, orbitofrontal gyrus, inferior frontal gyrus, anterior cingulate gyrus, ventral striatum, amygdala, posterior cingulate cortex, supramarginal gyrus, and occipital cortex (Supplementary Fig. 4)^{3–8}. Again, there seemed to be different degrees of activation across the different age groups, with adolescents appearing to have an enhanced neural response to feedback. These images parallel the whole-brain correlations reported in the main text of the paper, in that there were linear effects with age and stimulus decision value, but an enhanced adolescent response to feedback prediction error.

Stimulus onset > baseline by group

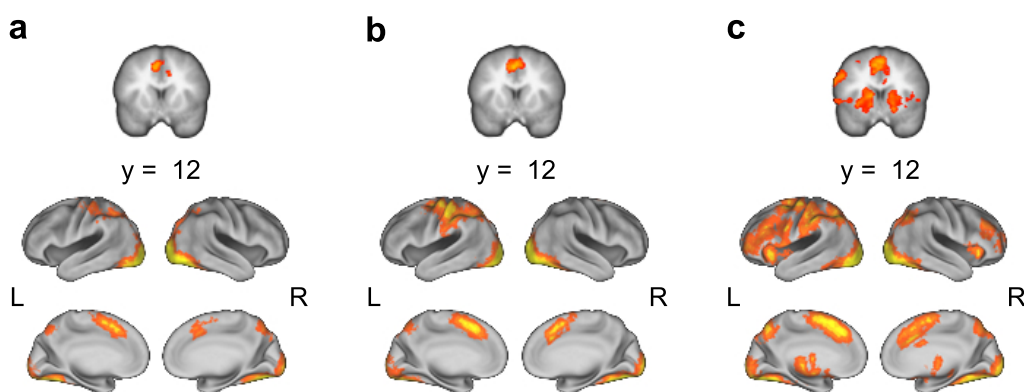


Figure 3: Neural regions that showed increased activity when viewing all stimuli as compared to baseline for (a) children ($n = 18$, aged 8–12), (b) adolescents ($n = 16$, aged 14–19), and (c) adults ($n = 11$, aged 25–30). All clusters survived whole-brain correction at $z > 2.3$, $p < .05$. For a list of clusters of activity, see Supplementary Table 2.

Positive > negative feedback by group

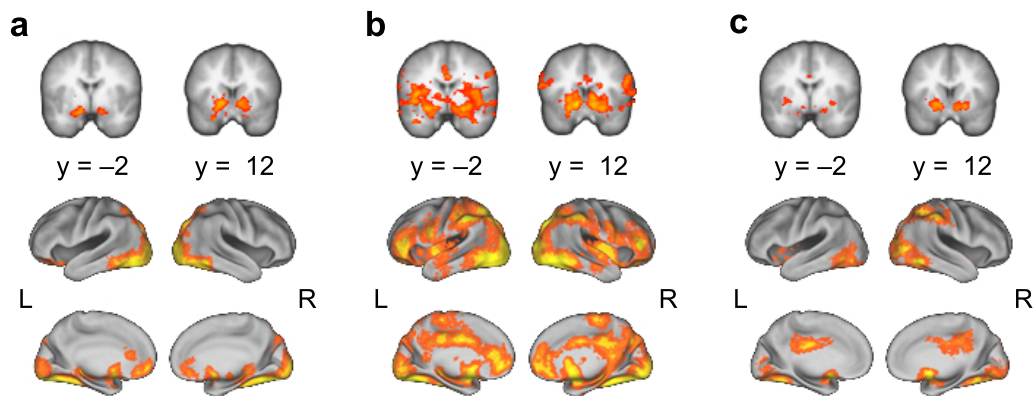


Figure 4: Neural regions that showed increased activity when viewing positive as compared to negative feedback for (a) children (n = 18, aged 8–12), (b) adolescents (n = 16, aged 14–19), and (c) adults (n = 11, aged 25–30). All clusters survived whole-brain correction at $z > 2.3$, $P < 0.05$. For a list of clusters of activity, see Supplementary Table 3.

Region of interest analysis as a function of age

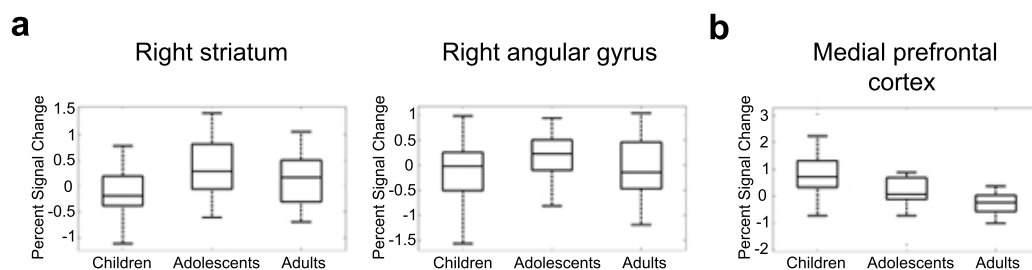


Figure 5: Boxplots for each of the regions that showed relationships with age in whole-brain analyses corrected for multiple comparisons at a threshold of $z > 2.3$, $P < 0.05$. Note that these are non-independent analyses, and are shown only to demonstrate the functional form of the observed whole-brain results. (a) Negative correlation between positive prediction error and age² in the right striatum and right angular gyrus. (b) Negative correlation between stimulus decision value and age in the medial prefrontal cortex. All boxplot lines show the lower quartile, median, and upper quartile values. Whiskers show the extent of the remaining data. Outliers are marked by crosshairs past the extent of the whiskers. See Figure 2a in the main text for images and Supplementary Table 4 for coordinates.

Whole-brain correlations between age and stimulus decision value/feedback prediction error in the secondary model

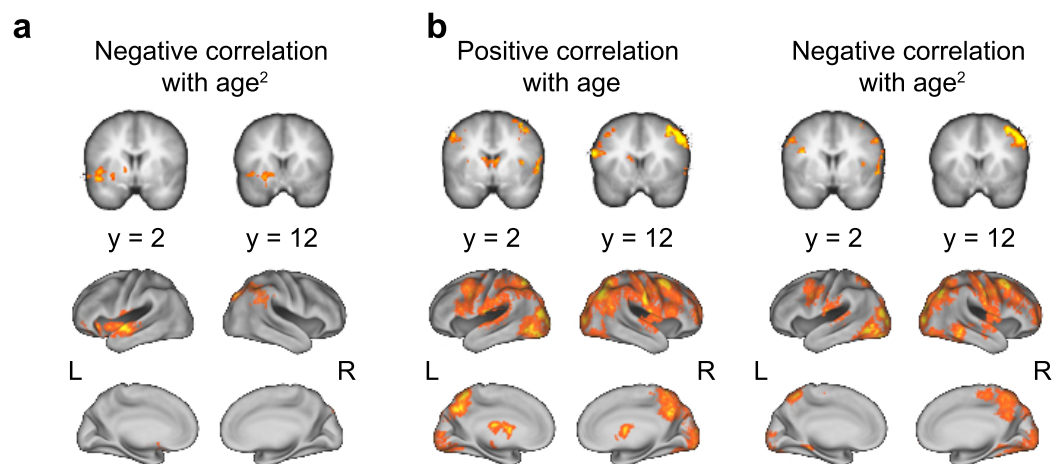


Figure 6: Regions that showed correlations with age in the secondary model when correcting at the whole-brain level at $z > 2.3$, $P < 0.05$. (a) Regions with a negative correlation between feedback prediction error and age². (b) Regions with relationships between stimulus decision value and age. See Supplementary Table 5 for coordinates.

Region of interest analysis as a function of age in the secondary model

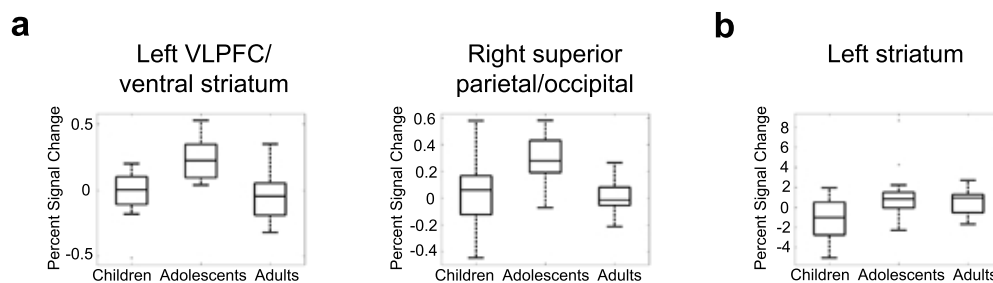


Figure 7: Boxplots for representative regions that showed relationships with age in whole-brain analyses corrected for multiple comparisons at a threshold of $z > 2.3$, $P < 0.05$. Note that these are non-independent analyses, and are shown only to demonstrate the functional form of the observed whole-brain results. (a) Negative correlation between feedback prediction error and age² in the left ventrolateral prefrontal cortex (VLPFC)/ventral striatum and right superior parietal/occipital cortex. (b) Positive correlation between stimulus decision value and age in the left striatum, although note that percent signal change asymptotes between adolescents and adults. All boxplot lines show the lower quartile, median, and upper quartile values. Whiskers show the extent of the remaining data. Outliers are marked by crosshairs past the extent of the whiskers. See Supplementary Fig. 6 for images and Supplementary Table 5 for coordinates.

Group comparisons for striatal activity related to positive – negative feedback

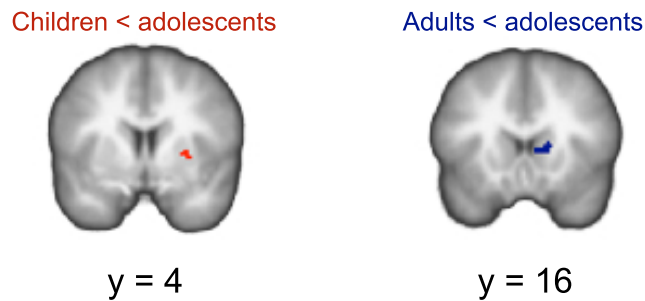


Figure 8: Striatal regions that showed a trend toward enhanced activity in adolescents as compared to children (red; corrected $P = 0.14$) and adults (blue; corrected $P = 0.06$) when viewing positive, as compared to negative, feedback. Results are cluster-mass corrected at $z > 2.3$.

Group comparisons for striatal activity related to positive prediction error

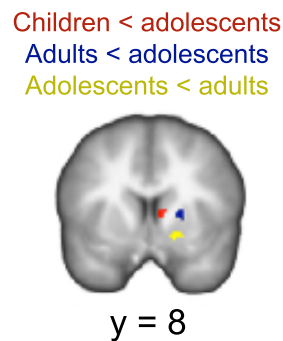


Figure 9: Regions within the striatum that showed differential activation related to positive prediction error across age groups. T-tests were conducted directly comparing children to adolescents and adults to adolescents. A region in the dorsal striatum was more active in adolescents than children (red; cluster-mass corrected $P = 0.08$). Another dorsal striatal region was more active in adolescents than adults at a more lenient threshold (blue; cluster-mass corrected $P = 0.48$, uncorrected $P < 0.05$). Lastly, a region in the ventral striatum was more active in adults than adolescents (yellow; cluster-mass corrected $P = 0.08$). There were no regions where children displayed more activity than adolescents.

Overall accuracy and response times by group

Behavioral Data	Accuracy(s.d.)			Response Times(s.d.)		
	Children	Adol.	Adults	Children	Adol.	Adults
Pred Overall	62.2(11.9)	69.4(5.7)	73.0(10.6)	911.9(218.0)	964.9(217.6)	1050.9(209.4)
Pred Block 1	54.3(15.6)	55.6(15.5)	57.7(17.9)	980.1(197.4)	1113.7(275.8)	1253.4(199.9)
Pred Block 6	68.3(18.7)	79.8(15.1)	84.1(16.0)	851.9(273.9)	868.6(171.9)	840.3(186.6)
Rand Overall	50.5(13.4)	50.5(12.6)	47.8(11.3)	929.6(263.6)	1017.5(254.6)	1128.9(208.2)
Rand Block 1	50.7(22.3)	58.9(14.2)	54.0(9.0)	943.8(231.9)	1141.4(298.2)	1310.6(265.8)
Rand Block 6	44.5(27.0)	48.8(19.1)	40.4(26.4)	884.6(420.2)	1008.8(315.5)	1035.8(235.0)
Large Rew Pred	63.4(15.5)	74.1(9.6)	78.6(13.0)	901.6(218.0)	908.9(206.6)	1016.9(205.7)
Small Rew Pred	61.1(10.7)	64.7(7.2)	67.3(10.1)	922.3(229.3)	1020.9(245.1)	1084.9(226.7)
Large Rew Rand	51.1(14.2)	52.0(17.6)	42.1(18.0)	945.4(268.8)	982.9(228.4)	1121.3(207.1)
Small Rew Rand	49.0(17.4)	49.1(19.2)	53.6(16.6)	913.8(276.9)	1052.2(308.3)	1136.5(225.0)

Table 1: Accuracy and response times of children, adolescents, and adults to stimuli throughout learning. Accuracy is the percent of optimal responses for predictable stimuli, and the percent of Northern responses for random stimuli. s.d. = standard deviation, Pred = predictable stimuli, Rand = random stimuli, Rew = rewards.

Regions showing greater activity for stimulus onset vs. baseline by group

Region	Coordinates (x,y,z in mm)	Max z	Extent (voxels)
Stimulus Onset			
<i>Children</i>			
L postcentral gyrus, L supramarginal gyrus, B superior parietal lobule, B fusiform gyrus, B precuneus, B lateral occipital cortex, B occipital pole, B cerebellum	36, -88, -2	5.39	11500
B anterior cingulate gyrus, B paracingulate gyrus, B pre-supplementary motor area, B supplementary motor area	-4, 14, 48	4.22	809
<i>Adolescents</i>			
R fusiform gyrus, R lateral occipital cortex, R occipital pole, R cerebellum	34, -52, -20	5.23	6711
L fusiform gyrus, L lateral occipital cortex, L occipital pole, L cerebellum	-24, -100, 6	5.41	4194
L precentral gyrus, L postcentral gyrus, L supramarginal gyrus, L superior parietal lobule, L precuneus, L superior lateral occipital cortex	-56, -26, 50	4.49	4089
R anterior cingulate gyrus, B paracingulate gyrus, B pre-supplementary motor area, B supplementary motor area	-4, 16, 42	4.61	1500
<i>Adults</i>			
B fusiform gyrus, B lateral occipital cortex, B occipital pole, B cerebellum	40, -90, 4	5.01	13846
L middle frontal gyrus, B anterior cingulate gyrus, B paracingulate gyrus, L superior frontal gyrus, B pre-supplementary motor area, B supplementary motor area, L precentral gyrus, L postcentral gyrus, L supramarginal gyrus, B angular gyrus, B precuneus, B superior lateral occipital cortex	0, 12, 54	4.15	9345
L frontal pole, L inferior frontal gyrus, L middle frontal gyrus, L frontal opercular cortex, L insula, L ventral striatum, L caudate, L putamen, L pallidum, L thalamus, L brainstem	-28, 4, 2	4.07	4755
R frontal pole, R inferior frontal gyrus, R middle frontal gyrus, R frontal opercular cortex, R insula, R ventral striatum, R caudate, R putamen, R pallidum, R thalamus, R amygdala	14, 8, 8	3.78	2755

Table 2: Clusters associated with stimulus onset. All clusters survived whole-brain correction at $z > 2.3$, $P < 0.05$ and are reported in MNI space (mm). B = bilateral, L = left, R = right.

Regions showing greater activity for positive vs. negative feedback by group

Region	Coordinates (x,y,z in mm)	Max z	Extent (voxels)
Positive – Negative Feedback			
<i>Children</i>			
L frontal pole, B medial prefrontal cortex, L orbitofrontal gyrus, B anterior cingulate gyrus, B paracingulate gyrus, B ventral striatum, B ventral caudate, B ventral putamen, B ventral pallidum, B amygdala, B hippocampus, B parahippocampal gyrus, B fusiform gyrus, B lateral occipital cortex, B occipital pole, B cerebellum	28, –52, –16	5.25	25602
<i>Adolescents</i>			
B frontal pole, B medial prefrontal cortex, B orbitofrontal gyrus, B inferior frontal gyrus, B anterior cingulate gyrus, B precentral gyrus, B postcentral gyrus, B ventral striatum, B ventral caudate, B ventral putamen, B amygdala, B hippocampus, B central opercular cortex, B posterior cingulate gyrus, B supramarginal gyrus, B angular gyrus, B superior temporal gyrus, B middle temporal gyrus, B inferior temporal gyrus, B fusiform gyrus, B precuneus, B lateral occipital cortex, B brainstem, B cerebellum	26, –78, –16	5.24	65240
<i>Adults</i>			
B fusiform gyrus, B lateral occipital cortex, R occipital pole, B cerebellum	6, –72, –8	3.68	4743
R postcentral gyrus, R supramarginal gyrus, R superior parietal lobule, R superior lateral occipital cortex	58, –34, 48	3.71	1531
L insula, L ventral striatum, L ventral caudate, L ventral putamen, L ventral pallidum, L amygdala	–18, 8, –10	4.18	883
R insula, R ventral striatum, R ventral caudate, R ventral putamen, R amygdala, R hippocampus	20, 10, –10	4.06	658
B anterior cingulate gyrus, B posterior cingulate gyrus	–6, –20, 26	3.21	594

Table 3: Clusters associated with increased activity when viewing positive as compared to negative feedback. All clusters survived whole-brain correction at $z > 2.3$, $P < 0.05$ and are reported in MNI space (mm). B = bilateral, L = left, R = right.

Regions showing associations with age

Region	Age Correlation	Coordinates (x,y,z in mm)	Max z	Extent (voxels)
a. Positive Prediction Error				
R striatum	Neg Age ²	14, 16, 4	3.72	343
R angular gyrus	Neg Age ²	60, -44, 28	3.80	286
b. Stimulus Decision Value				
B medial prefrontal cortex	Neg Age	0, 50, -8	4.11	680

Table 4: Clusters associated with age-related parametric changes in the response to (a) positive prediction error and (b) stimulus decision value. All clusters survived whole-brain correction at $z > 2.3$, $P < 0.05$ and are reported in MNI space (mm). B = bilateral, R = right.

Regions showing associations with age in the secondary model

Region	Age Correlation	Coordinates (x,y,z in mm)	Max z	Extent (voxels)
a. Feedback Prediction Error				
L frontal pole, L orbitofrontal gyrus, L insula, L ventral striatum, L ventral putamen, L ventral pallidum, L central opercular cortex, L planum polare, L superior temporal gyrus, L middle temporal gyrus	Neg Age ²	−56, −4, −8	3.91	1252
R supramarginal gyrus, R angular gyrus, R superior lateral occipital cortex	Neg Age ²	60, −46, 32	3.54	562
b. Stimulus Decision Value				
B precuneus, B superior lateral occipital cortex	Pos Age	48, −60, 38	4.35	5904
B lingual gyrus, B inferior lateral occipital cortex	Pos Age	−40, −78, 14	4.43	4349
R inferior frontal gyrus, R insula, R middle frontal gyrus, R precentral gyrus, R postcentral gyrus, R central opercular cortex, R planum polare, R supramarginal gyrus	Pos Age	40, 8, 56	4.63	3507
L postcentral gyrus, L central opercular cortex, L supramarginal gyrus, R parietal opercular cortex, L planum temporale, R superior temporal gyrus	Pos Age	−58, −16, 14	4.12	911
L inferior frontal gyrus, L middle frontal gyrus, L precentral gyrus	Pos Age	−60, 12, 20	3.90	878
R frontal pole	Pos Age	36, 64, 4	4.02	669
B caudate, L putamen, B thalamus	Pos Age	−24, −8, 18	3.61	659
R occipital fusiform gyrus, R lingual gyrus, L superior lateral occipital cortex, B lateral inferior occipital cortex, R cerebellum	Neg Age ²	−40, −78, 14	4.16	4397
R superior parietal lobule, R angular gyrus, B precuneus, B superior lateral occipital cortex	Neg Age ²	−4, −60, 60	4.29	3397
R inferior frontal gyrus, R insula, R middle frontal gyrus, R precentral gyrus, R postcentral gyrus, R central opercular cortex	Neg Age ²	38, 10, 54	4.55	2786
L middle frontal gyrus, L precentral gyrus	Neg Age ²	−60, 4, 34	3.65	548
R frontal pole	Neg Age ²	38, 64, 4	4.01	548
L postcentral gyrus, L superior temporal gyrus	Neg Age ²	−70, −22, 32	3.42	391

Table 5: Clusters associated with age-related parametric changes in the secondary model in the response to (a) feedback prediction error and (b) stimulus decision value. All clusters survived whole-brain correction at $z > 2.3$, $P < 0.05$ and are reported in MNI space (mm). B = bilateral, L = left, R = right.

SUPPLEMENTARY METHODS

Participants. We recruited 67 healthy developing right-handed participants between the ages of 8 and 30 from the community. They were recruited with online advertisements and by randomly calling families found through a commercially available list of households within a 25 mile radius of UCLA (Survey Sampling Inc., Fairfield, CT). We recruited participants as part of a larger study examining cortico-striatal functioning in typically developing children and siblings of probands with childhood onset schizophrenia; the analyses reported here only include typically developing individuals. We interviewed potential participants to determine whether they met inclusion/exclusion criteria for the study. For potential participants who were minors, we interviewed their parents to determine the minor's eligibility. Participants were eligible if they had no history of CNS disease, DSM-IV disorders, or learning disabilities and no treatment with anti-psychotic drugs or substance use in the past two years. Additionally, they could not have any metal in their bodies other than dental fillings, and could not be pregnant. Of these 67 potential participants, 61 were eligible to participate in the MRI study. Of these, 16 were excluded because of: technical issues (six; for example, button box recording responses did not work), excessive motion (six; more than one translational displacement of 3 mm or greater), or poor performance (four; some task conditions had no correct trials, which precluded data analysis). Of the remaining 45 participants included in this analysis, there were 18 children aged 8–12 (mean age 10.8, 9 females), 16 adolescents aged 14–19 (mean age 15.8, 6 females) and 11 adults aged 25–30 (mean age 26.5, 4 females). All participants (and their parents if they were under 18) provided written informed consent or assent (for minors) according to the procedures of the UCLA Institutional Review Board.

Experimental Design and Procedure. Participants performed a probabilistic learning task during fMRI acquisition. They viewed the task through LCD goggles and responded using an MR-compatible button box. Each trial lasted on average 5,000 ms (range 3,900–7,750 ms). On each trial, participants saw an abstract computer-generated stimulus (ArtMatic Pro, U&I Software LLC) that they were told represented the pattern on a college t-shirt. They were asked to classify the shirt as being worn by someone who went to Northern University (left button press) or Eastern University (right button press). The probabilistic nature of the task was emphasized by explaining that people occasionally wear t-shirts of schools other than that which they attended, therefore the feedback would not always be consistent for each stimulus. The stimuli were on the monitor for an average of 3,000 ms (jittered between 2,500 and 5,000 ms), during which time participants had to indicate which university the person wearing the shirt attended with a button press with their index or middle finger on their right hand. At the end of the stimulus presentation, feedback appeared on the screen for 1,250 ms. Feedback consisted of both information about the intended response (i.e., which university the person wearing the t-shirt actually went to) and, if the participant was correct, a reward of gold coins. The reward was either one coin or five coins, and after the task participants were paid 5 cents for every coin they amassed throughout the scan. A blank screen served as the intertrial interval, with an average of 750 ms (range 150–1500 ms; Fig. 1 in the main text). The trial order and length were optimized for separating the neural response to stimuli from the neural response to feedback using custom MATLAB code. There were 144 trials spread over two 306 second runs. Six stimuli were used in this task: four were predictable and associated 83% of the time with one of the two outcomes (two with Northern and two with Eastern) and two were random (associated 50% of the time with each university). Within each of the three stimulus types (predictable Northern, predictable Eastern, and random), one stimulus was associated with a large reward of five coins, and the other was

associated with a small reward of one coin.

MRI Data Acquisition. We collected imaging data with a 3T head-only Siemens Allegra MRI scanner at the Ahmanson-Lovelace Brain Mapping Center at the University of California, Los Angeles. For each functional run we collected 182 T2*-weighted echoplanar images (33 slices, slice thickness 4 mm, TR = 2000 ms, TE = 30 ms, flip angle = 90°, matrix 64 x 64, field of view 200 mm). We also acquired a T2-weighted matched-bandwidth high-resolution anatomical scan with the same slice prescription as the functional images. Lastly, we collected a T1-weighted magnetization-prepared rapid-acquisition gradient echo (MPRage; 160 sagittal slices, slice thickness 1 mm, TR = 2000 ms, TE = 2.1 ms, matrix 192 x 192, field of view 256) for anatomical registration.

Behavioral Data Analysis. We analyzed behavioral data for both response time and percent of optimal responses using R version 2.4.1 (The R Foundation for Statistical Computing, <http://www.R-project.org>). Accuracy was defined in terms of the proportion of optimal responses. Therefore, for predictable stimuli associated with Northern, a Northern response was considered optimal, even if feedback on any given trial was Eastern. For random stimuli, since there was no optimal response, a response of Northern was arbitrarily defined as optimal for calculations. For response time analyses, all trials were included since there were no clear-cut incorrect responses to exclude. We analyzed the data using a 2 (stimulus-type: predictable or random) x 2 (reward: large or small) x 6 (block: 1–6) repeated-measures ANOVA with age group (3: children, adolescents, adults) as the between-groups factor. To look at overall money earned, we conducted a one-way ANOVA with three levels (age group).

fMRI Data Analysis. We processed and analyzed imaging data using FSL (FMRIB's Software Library, www.fmrib.ox.ac.uk/fsl). For preprocessing we used FSL 3.3, including BET to extract the brain from the skull and MCFLIRT for motion correction. Following motion correction, we submitted the data to independent components analysis using MELODIC ICA, and we used the results from this analysis to identify and remove potentially artifactual components in the data. After manually identifying artifactual components from one run for each of six randomly-selected participants, we trained a classifier⁹ to automatically identify artifactual components in all the participants, and then used MELODIC to remove them from the data.

We conducted statistical analyses in FSL 4.1 using FEAT 5.98. The statistical model included events modeled at stimulus onset, participant response (choice), and feedback onset. In a separate model examining model-based responses, we derived parametric regressors for decision value, choice value, and prediction error from a simple network model that used one node per stimulus (six in this study) as an input layer and two nodes as an output layer, one node representing each outcome. The input and output layers of the model were fully connected, without the use of any intermediary hidden layers. Connection weights were based on individual participant learning and updated after each trial using the Rescorla-Wagner rule¹⁰. This model had a single free parameter: a learning rate (alpha), ranging from zero to one, that modulated how strongly individual trials altered the association between stimuli and responses. The decision value (DV) of the presented stimulus was updated after each trial by a multiplicative function of the participant's learning rate (alpha) and the prediction error (PE) on the trial:

$$DV_{(n+1)} = DV_{(n)} + (\alpha) * PE_{(n)}$$

We calculated decision value on each trial as the sum of the decision values relating to both possible outcomes for the given stimulus, even though only the decision value of the chosen outcome was updated with the above equation. The reason for using the sum of the decision values related to both possible outcomes as the trial decision value was because it is likely that the decision value for each response is encoded in similar areas. Choice value was the decision value for the chosen stimulus. We calculated prediction error as the difference between the actual observed outcome (1 or 0 for Northern and Eastern respectively) and the choice value on each trial.

We fit learning rate parameters independently for each participant. We ran a series of simulations where starting weights mapping stimuli to responses were initialized to random values between -1 and 1 , using grid search across learning rates from 0.02 to 0.6 . The random starting weights were biased in proportion with the participant's responses to each category in order to most optimally fit actual participant behavior. After these initial networks were generated, we ran the model using the same sequence of stimuli as observed by the participant. We averaged the simulations of each learning rate to find the learning rate parameter that best minimized the difference between model responses and participant responses. In these data it was the case that the global best fit was often a very small learning rate, but larger learning rate values still led to very close fits. The learning rate selected for the model was the largest rate that fit the participant's behavior within 10% of the absolute best fit, with a minimum value of $.02$ and a maximum value of 0.6 . This adjustment was due to the need for regressor variability when fitting the output of this model to the imaging data, as very small or very large learning rates produce less variable prediction error regressors. The minimum value chosen was less than the best learning rate and the maximum value was larger than the best learning rate for all participants so as not to miss the optimal learning rate for any participants. The estimated learning rate did not significantly differ across the three age groups ($F(2,42) = 1.34$, $P = 0.27$; alphas = 0.55 , 0.55 and 0.51 for children, adolescents and adults respectively (see Supplementary Fig. 2 for plots of how learning rate relates to both decision value and prediction error).

Once the participant's learning rate was determined, we ran the model again with all connection weights initialized to zero using both the selected learning rate and the actual input the participant received. From this model a quantification of decision value, choice value and prediction error for each stimulus/outcome pair could be determined on each trial.

The secondary model implemented to clarify the results of the above model included parametric regressors for decision value and prediction error calculated in a manner different from the above model (Lin, Rangel & Adolphs, unpublished). For each stimulus, the value of a response was 1 if it was the optimal response (i.e., Northern for a predictable stimulus associated with Northern or Eastern for a predictable stimulus associated with Eastern) and 0 if it was the alternate, non-optimal response. For random stimuli, as there was no optimal response, a value of 1 was arbitrarily assigned to Northern and 0 to Eastern. The decision value (DV) was defined as the mean of all chosen outcomes throughout the task thus far:

$$DV_{(n)} = \text{mean}(DV_{(1:n)})$$

We assigned a decision value of 0.5 to the first instance of each stimulus, as participants began with

no knowledge of optimal outcome. We calculated decision value using the actual value of the response associated with the first instance of each stimulus, not 0.5. We defined prediction error as the difference between the actual observed outcome (1 or 0 for optimal and non-optimal respectively) and the expected outcome (decision value) on each trial.

We expected participants to learn stimulus-response associations in the predictable, but not the random, stimuli. This was reflected in behavioral responses, so we focused the decision value analysis on predictable stimuli. As expected, the prediction error-related neural response to predictable and random stimuli was similar, given that an unexpected reward was responded to equivalently whether it was associated with a predictable or a random stimulus. Therefore, we included all stimuli in the prediction error analysis.

For the first level analysis, we spatially smoothed images using a Gaussian kernel of FWHM 5 mm. We carried out time-series statistical analysis using FILM (FMRIBs Improved Linear Model) with local autocorrelation correction after highpass temporal filtering (Gaussian-weighted least-squares straight line fitting, with $\sigma = 33.0$ s). We created regressors of interest by convolving a delta function representing each event of interest with a canonical (double-gamma) hemodynamic response function¹¹. We created parametric regressors by modulating the amplitude of a delta function using a demeaned version of the parameter of interest. In addition to regressors of interest, we included estimated motion parameters and their temporal derivatives (i.e., displacement) as nuisance regressors. We performed linear contrasts for comparisons of interest.

We applied a 3-step registration process using FSL's FLIRT module for linear registration. EPI functional images were first registered to an in-plane T2-weighted structural image (matched bandwidth; 7 DOF). The in-plane structural image was registered to the high-resolution structural image (MPRage; 7 DOF), and the high-resolution image was registered to standard MNI152 space using FLIRT linear registration with 12 degrees of freedom. These transformation matrices were combined to provide the transform from EPI to MNI space, and this transform was applied to the results from the first-level analysis.

We combined data across runs using a fixed effects model, and then used a mixed effects model at the group level with FSL's FLAME (Stage 1 only). The model included a regressor modeling mean activity and demeaned regressors for age, age², gender, and accuracy. Contrasts included positive and negative relationships with each of the regressors. We performed outlier deweighting using a mixture modeling approach¹². We thresholded results at a whole-brain level using cluster-based Gaussian random field theory, with a cluster-forming threshold of $z > 2.3$ and a whole-brain corrected cluster significance level of $P < 0.05$. We performed cortical surface renderings using CARET software (<http://brainmap.wustl.edu>). We mapped group statistical maps to the Population Average Landmark and Surface-based (PALS) atlas using a multifiducial mapping technique¹³. For the purposes of presentation, we overlaid data on the average atlas surface. For striatal region of interest (ROI) analyses we used a structural ROI including bilateral caudate, putamen and nucleus accumbens from the Harvard-Oxford Probabilistic Atlas (FSL; provided by the Harvard Center for Morphometric Analysis) thresholded at a 50% probability of being in each structure. For ROI analyses, we used FSL's Randomise tool to correct for multiple comparisons using a cluster-mass correction thresholded at $z > 2.3$ and $P < 0.05$.

SUPPLEMENTARY REFERENCES

1. Aron, A. R. *et al. J. Neurophysiol.* **92**, 1144–52 (2004).
2. Poldrack, R. A. *et al. Nature* **414**, 546–50 (2001).
3. Breiter, H. C., Aharon, I., Kahneman, D., Dale, A. & Shizgal, P. *Neuron* **30**, 619–39 (2001).
4. Knutson, B., Fong, G. W., Adams, C. M., Varner, J. L. & Hommer, D. *Neuroreport* **12**, 3683–7 (2001).
5. Liu, X. *et al. J. Neurosci.* **27**, 4587–97 (2007).
6. Nieuwenhuis, S. *et al. Neuroimage* **25**, 1302–9 (2005).
7. O’Doherty, J., Kringelbach, M. L., Rolls, E. T., Hornak, J. & Andrews, C. *Nat. Neurosci.* **4**, 95–102 (2001).
8. Rogers, R. D. *et al. Biol. Psychiatry* **55**, 594–602 (2004).
9. Tohka, J. *et al. Neuroimage* **39**, 1227–45 (2008).
10. Rescorla, R. A. & Wagner, A. R. In Black, A. & Prokasy, W. F. (eds.) *Classical Conditioning II: Current Research and Theory*, 64–99 (Appleton Century Crofts, New York, NY, 1972).
11. Woolrich, M. W., Ripley, B. D., Brady, M. & Smith, S. M. *Neuroimage* **14**, 1370–86 (2001).
12. Woolrich, M. W. *Neuroimage* **41**, 286–301 (2008).
13. Van Essen, D. C. *Neuroimage* **28**, 635–62 (2005).