Supplementary Materials

Motivational context modulates prediction error response in schizophrenia

Jenna Reinen, Jared X. Van Snellenberg, Guillermo Horga, Anissa Abi-Dargham, Nathaniel D. Daw, Daphna Shohamy

Reinforcement Learning Model Analyses

Subject Level Model. We used a Q learning model to estimate two parameters (learning rate α and beta β) for each subject. Estimation of these parameters on a trial-by-trial basis was used to generate trial-specific prediction error (δ) regressors for use with the functional imaging data. The value of the choice options were modeled as:

$$Q_{c(t)} = Q_{c(t)} + \alpha * \delta_t$$

$$\delta_t = fb(t) - Q_{c(t)}$$

Where learning rate α is a free learning rate parameter. On each trial (t), the model can learn from prediction error δ, scaled by the learning rate. Across conditions, positive feedback $fb(t)$ was coded as 1 and negative feedback was coded as -1, ensuring that any condition-wise effects on parameters would not be confounded by feedback coding. We assumed that subjects made their choice c (0 or 1) according to the learned values $Q_{c(t)}$ at each step using a softmax distribution with free inverse temperature parameter β[1]:

$$P(choose\ a|Q_1, Q_2) = \frac{\exp(\beta * Q_1)}{\exp(\beta * Q_1) + \exp(\beta * Q_2)}$$

Critically, the ability to update the cue value relies on information in the form of PE, the difference between value expected and outcome received. Prediction error was calculated as the difference between the expected Q value given choice at trial (t) and the actual feedback (fb) received.

For the parameter fits reported below, we use a change of variables in which α is transformed by an inverse Gaussian CDF (so that its value, otherwise bounded in [0,1] ranges between $-\infty$ and ∞ and group and condition effects can be summed along the unbounded real line), and β is reparameterized as $\beta \cdot \alpha$, which makes its magnitude more stable under variation in learning rate.[2]

Group Level Models**.** In order to assess effects on the model parameters of group and condition (and, in subsequent analyses, feedback valence), we used a multilevel mixed model, in which the individual subject and condition-specific parameters for the above subject-level model were nested under population-level distributions. We

then jointly estimated individual and population-level parameters, and effects of condition and disease on them.

At the population level, both β and α were drawn for each subject from a Gaussian distribution with free mean and standard deviation parameters representing the population tendency and across-subject variation in the parameter. We modeled separate group means for patients vs controls by adding an additional term to the population level mean. Thus the softmax temperature for subject s was drawn from:

$$\beta_s \sim N\left(\mu^{\beta} + k^{\beta * g} \cdot I(s, 'patient'), \sigma^{\beta}\right)$$

where I is a binary indicator of patient (1)/control (0) status and $k^{\beta * g}$ ('g' for group) is a free parameter coding the difference in mean softmax temperature for the patient group. Learning rate was coded analogously.

The model also included within-subject effects of condition on the temperature and learning rate, such that for condition c (1 gain, 0 loss), $\beta_s^c = \beta_s + \beta_{s,gain} \cdot I(c, 'gain')$. The per-subject condition effects $\beta_{s,gain}$ (representing the difference in softmax temperature in the gain condition) were again drawn from group-level distributions, with a mean and standard deviation and an additional parameter capturing any differences in this effect between patients and controls:

$$\beta_{s,gain} \sim N\left(\mu^{\beta * c} + k^{\beta * c * g} \cdot I(s, 'patient'), \sigma^{\beta * c}\right)$$

Again, analogous condition effects were included for the learning rate.

All together, the population-level model contained two means for the learning parameters ($\mu^{\alpha}, \mu^{\beta}$), two means for the condition effects, ($\mu^{\alpha * c}, \mu^{\beta * c}$), four effects of group or group x condition ($k^{\alpha * g}, k^{\alpha * g * c}, k^{\beta * g}, k^{\beta * g * c}$), and four standard deviations ($\sigma^{\alpha}, \sigma^{\alpha * c}, \sigma^{\beta}, \sigma^{\beta * c}$).

To complete the model, we specified hyperpriors on these parameters, which were taken to be uninformative within the range of parameters seen in the literature. For the softmax temperature, hyperpriors on the standard deviation terms were taken as half-$Cauchy$(0,2), and μs and ks as N(0,2). For the softmax temperature, we used N(0,1) for all parameters (the half of that distribution with positive support for variances). This was chosen so that the a priori distribution over the bounded subject level learning rate (when transformed by a Gaussian CDF with variance 5) was approximately uniform in 0,1.

We also considered two followup models, which augmented the aforementioned model by allowing for either the learning rate (model 2) or feedback scaling (model 3) to vary between trials with correct vs. incorrect feedback. For model 2, this involved adding an additional within-subject factor $\alpha_{s,inc}$ to the learning rate (prior to transformation with the Gaussian CDF) for negative feedback trials. This had its

own group level mean $\mu^{\alpha*inc}$, standard deviation $\sigma^{\alpha*inc}$, and interaction with group $k^{\alpha*g*inc}$. For model 3, we instead rescaled the feedback $fb(t)$ (1 or -1) by a within-subject factor $f_{s,cor}$ or $f_{s,inc}$, depending on the feedback. These were each drawn from group-level distributions with means (μ^{cor}, μ^{inc}), standard deviations (σ^{cor}, σ^{inc}), condition means (μ^{cor*c}, μ^{inc*c}) and standard deviations (σ^{cor*c}, σ^{inc*c}), and group or group x condition effects (k^{cor*g}, k^{inc*g}, $k^{cor*g*c}$, $k^{inc*g*c}$). This model omits the softmax temperature β (equivalently, takes it to be everywhere 1) and the associated group-level parameters, because the rescaling of feedback terms has equivalent effect by rescaling the learned Q. The model also omitted condition and group x condition effects on the learning for simplicity.

Parameter Estimation and Significance Testing**.** We estimated the model parameters by using Markov Chain Monte Carlo (MCMC) inference, implemented in the Stan programing language[3, 4]. This is a procedure which provides samples from the posterior distribution over all model parameters conditional on the observed choices. We ran four chains of 1000 samples each, discarding the first 100 samples per chain for burnin. This approach is similar to those used by Otto et al[2] and Sharp et al[5]. We verified convergence and mixing of the model by visual inspection of traceplots and by examining the diagnostic statistic $\hat{r}$ [6], for which values close to 1 are consistent with convergence. This value was less than 1.1 for all parameters in all models. We also augmented Model 1 to additionally generate samples of the per-subject and trial prediction errors δ, which are a function of the learning rate, according to its posterior distribution. The posterior median value of each of these samples was calculated and used in our imaging analysis as the trial- and subject-unique parametric repressor. We report model likelihoods for each group in Table S2.

We used the distribution of the samples to generate symmetric 95% credible intervals (the range between the 2.5th and 97.5th percentiles) for each parameter of interest. We treated a parameter as significant if zero lay outside this range. We also report "P" as one minus the size of the largest symmetric credible interval that excludes zero[7], which is roughly analogous to a two-tailed P value.

Results**.** In **Table S1**, we report the median, 95% credible interval, and significance of the group level parameter, estimated using MCMC for each of the three models Here, values of k in bold indicate we can, with more than 95% confidence, exclude zero for the effect on learning. This shows a significant effect of the softmax temperature by group, which (in model 3, which divides it into a separate scaling term for correct vs. incorrect feedback), is significant only for the correct trials. The implications for these findings are discussed within the main text.

In addition to the model described in the manuscript, the fits of two followup models, we examined whether the groupwise effect on learning might reflect a more specific effect on positive (correct) or negative (incorrect) feedback trials. In a model (model 2) where learning rate was allowed to vary as a function of feedback type, we found no differences between learning rates by trial type, nor any group

effect on this difference (P=0.21, P=0.23); the group effect on softmax temperature remained. However, in a model where the feedback itself was instead scaled separately on each trial type (thereby splitting up the softmax temperature, which scales value, into separate components associated with positive and negative feedback, model 3), the effect of group was significant only for positive feedback trials (P=.005), though note that the difference in effects between positive and negative feedback trials was not, itself, significant (P=.31). Altogether these results suggest that patients' choices were less driven by feedback; that the effect was not detectably related to learning rate; and that the effect was particularly pronounced for positive feedback trials.

We also assessed the fit of the model to each subject's behavior by computing the log-likelihood of each subject's choices using the model as fit to the remaining subjects (leave one out). We augmented the MCMC models to compute samples from the likelihood of each subject's choices, according to the distribution of model parameters fit to the remaining subjects in a leave-one-out procedure. We then took the log average of the samples to compute expected per-subject predictive log likelihoods, which we compared between patients and controls using two-sample t-tests, and between models using paired-sample t-tests over subjects. While there was a significant group difference in fit ($t = -2.37$; $p = 0.02$) both groups fit better than chance (patients: $t = 4.10$ $p < 0.001$; controls: $t = 8.71$, $p < 0.001$). Further, both groups showed a spread across likelihoods. For this analysis, chance was calculated to be -83 (below chance being below -83). One patient was at -83, and one patient was slightly below chance, demonstrating that, while arguably two subjects were not fit well by the data, the majority of patient subjects were above chance as defined by the model.

Across subjects, the three models were not significantly different from one another in terms of fit, reflecting the fact that they represented alternative parameterizations for the same underlying effect (Supplementary Table S2).

MODEL 1	**Lower CI**	**Posterior Median**	**Upper CI**	**Bayesian P**
μ - beta	0.58	0.74	0.91	<0.001
μ - beta*cond	-0.23	-0.04	0.18	0.65
μ - alpha	-1.38	-0.62	0.21	0.93
μ - alpha * cond	-1.28	-0.51	0.3	0.9
k - beta * group	-0.68	-0.44	-0.21	**<0.001**
k – beta * group * cond	-0.18	0.09	0.37	0.75
k - alpha * group	-1.67	-0.49	0.68	0.214
k - alpha * group * cond	-1.29	-0.04	1.27	0.48
σ - beta	0.19	0.27	0.39	
σ - beta * cond	0.16	0.28	0.44	
σ - alpha	1.25	1.76	2.45	
σ - alpha* cond	0.13	0.95	1.8	
MODEL 2	**Lower CI**	**Posterior Median**	**Upper CI**	**Bayesian P**
μ - beta	0.58	0.74	0.91	<0.001
μ - beta*cond	-0.2	0	0.23	0.493
μ - alpha	-1.16	-0.3	0.66	0.726
μ - alpha * cond	-1.21	-0.3	0.6	0.753

μ - alpha * neg fb	-2.07	-0.98	0.1	0.9653
k - beta * group	-0.66	-0.42	-0.19	**<0.001**
k – beta * group * cond	-0.21	0.08	0.35	0.707
k - alpha * group	-1.43	-0.15	1.15	0.405
k - alpha * cond * group	-1.35	0.04	1.38	0.523
k - alpha*group * neg fb	-2.2	-0.62	0.99	0.228
σ - beta	0.2	0.29	0.41	
σ - beta * cond	0.16	0.29	0.46	
σ - alpha	0.86	1.6	2.38	
σ - alpha* cond	0.23	1.28	2.23	
σ – alpha * neg fb	0.29	1.62	3	
MODEL 3	**Lower CI**	**Posterior Median**	**Upper CI**	**Bayesian P**
μ - pos fb	0.84	1.15	1.5	<0.001
μ - pos fb * cond	-0.21	0.11	0.43	0.235
μ - neg fb	-0.03	0.29	0.56	0.032
μ - neg fb * cond	-0.59	-0.2	0.22	0.855
μ - alpha	-0.52	-0.25	0.01	0.9695
k - pos fb * group	-1.09	-0.61	-0.17	**0.005**
k - pos fb * cond * group	-0.34	0.13	0.65	0.7
k- neg fb * group	-0.62	-0.22	0.19	0.133
k - neg fb * cond * group	-0.46	0.07	0.62	0.61
σ - pos fb	0.25	0.45	0.69	
σ - pos fb * cond	0.11	0.39	0.69	
σ - neg fb	0.08	0.25	0.48	
σ - neg fb * cond	0.07	0.28	0.59	
σ - alpha	0.43	0.65	0.94	

Supplementary Data Table 1. 95% credible interval and posterior median are given for each group level parameter. Values were taken from the distribution of samples generated from Bayesian statistical inference using an MCMC procedure. Significant values indicate a nonzero effect with more than 95% confidence.

MODEL	**Mean**	**+/- 1 SEM**
Model 1, Patients	-64.18	4.59
Model 2, Patients	-64	4.64
Model 3, Patients	-64.51	4.71
Model 1, Controls	-50.12	3.77
Model 2, Controls	-50.06	3.79
Model 3, Controls	-50.68	3.83

Supplementary Data Table 2. Model likelihoods for each group.

Finally, to check the success of our parameter estimation, we ran a data simulation. We augmented the MCMC model to generate samples from the per-trial, per-subject simulated probabilities of optimal choices (the so-called posterior predictive density over choices), in order to compare simulated model behavior to observed choices. Our results matched the actual data relatively well (see Supplementary Figure 1). Similar to analyses in the actual data, an ANOVA revealed an effect of group ($F=5.36$, $p=0.02$ but not of condition ($F=9$, $p=0.95$) and no interaction ($F=0.75$,$p=0.38$).

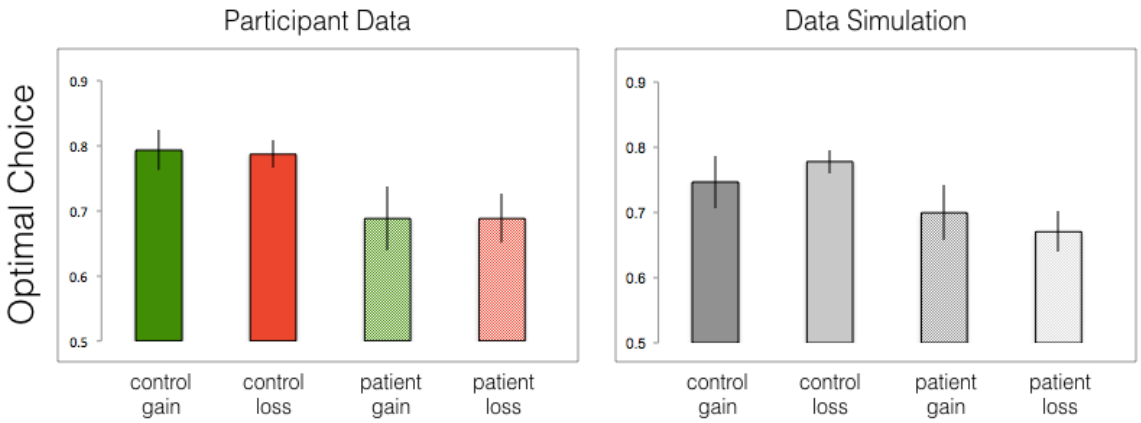

Supplementary Data Figure 1. Participant data versus data generated from data simulation.

Functional MRI Preprocessing Steps

Functional images were preprocessed with SPM8 (Wellcome Department of Imaging Neuroscience, London, UK) and in-house Matlab code. Data was converted to a 32-bit floating point precision Analyze format to reduce the impact of rounding errors, and an in-brain mask was used to detect artifactual volumes by comparing global signal across a sliding window; volumes which differed more than 8 mean deviations from the adjacent volume's mean signal value were subsequently not included in analysis. Functional images were then slice-time corrected and realigned using INRIAlgin[8] to correct for motion. T1 and EPI images were realigned to match templates from the International Consortium for Brain Imaging (ICBM), and a six-parameter affine coregistration technique, along with manual visual inspection, was implemented to complete and check reorientation and coregisration. T1 images were then segmented into their tissue components (gray matter, white matter, and cerebrospinal fluid) and the normalization parameters generated in the segmentation algorithm were applied to the coregistered T1 and EPI images. This process was checked visually. Images were then smoothed with an 8mm Gaussian kernel, which was scaled by the first-level hemodynamic response function (HRF) across runs. For more information on these preprocessing steps, see Van Snellenberg et al, 2015 [9].

Data quality was assessed by examining global temporal signal-to-noise, as well as visual inspection of template-normalized images (to ensure that segmentation and normalizations algorithms functioned correctly), and were deemed to be of sufficient quality in all cases. To assess movement for participants, we calculated mean absolute displacement for each subject based on raw measurements. A direct comparison of controls versus patients revealed that patients demonstrated greater movement than controls in the gain ($t = -2.03$, $p = 0.04$) but not the loss condition ($t = -1.45$, $p = 0.15$). However, we account for motion by using motion as a regressor of no interest in our GLM.

Exploratory Analyses

We assessed relationships between learning measures and disease symptoms as an exploratory analysis, given the limited sample size. Behaviorally, we examined whether there existed a relationship between negative symptoms with learning performance. To determine whether working memory capacity is related to reward responses and learning, we examined correlations between working memory performance on the self-ordered working memory task (SOT)[10] and learning performance. Because prior findings have reported that patients with negative symptoms tend to be the most impaired in terms of learning from positive outcomes [11-13], we limited our analyses to the gain condition. For the functional imaging data, we examined whether PE responses during feedback in mPFC and striatum were associated with these measures. To do so, we extracted the beta values from the contrast of feedback PE in the gain > loss conditions in these ROIs in patient subjects. We examined the Pearson correlation between these beta values with performance and symptom measures, and then used a two-sample t-test to determine whether beta values in each of these regions differed based on the patients' medication history, i.e. whether subjects were drug-naïve or drug-free at the time of testing (9 of 16 were medication free). Of note, these analyses were performed with a group of 16 participants, and as such, any inferences should be taken with this sample size in mind.

We first examined behavioral and neural relationships to learning performance. We did not find a relationship between PANSS negative symptoms in beta from Model 1 (r=0.14, p=0.61) or in Model 2(r=0.20, p=0.46).

Similarly, we did not find any significant relationships with optimal choice learning performance and working memory performance (gain condiiton: r=0.32, p=0.22). We also tested the relationship between behavioral performance and PANSS negative symptoms, and did not find any relationship between them (optimal choice: r=0.27, p=0.31, β: r=0.14, p=0.61). Likewise, there was no correlation between PANSS negative symptoms and the gain > loss feedback prediction error beta values extracted from nucleus accumbens (r=-0.10, p=0.71) or mPFC (r=-0.11, p=0.68), or between optimal choice performance in the gain condition and in patients betas within mPFC (r=-0.34, p=0.20) or in bilateral accumbens (r=-0.18,p=0.48).

To assess the effect of working memory (SOT) on prediction error response in striatum, we used the extracted beta values described above, and then assessed this measure in an analysis of covariance (ANCOVA) with a factor of group to determine if working memory performance measured by the capacity estimation of the self-ordered working memory test (SOT)[10] was related to the accumbens beta (gain>loss) response of interest. We did not find any significant effects of working memory from this analysis (F(group) = 0.5, p = 0.82, F(WM) = 0, p = 0.97, F(group * WM) = 0.04, p = 0.84).

Finally, for group differences between subjects who were medication naïve and medication free, and found no differences in the G-L beta in accumbens ($t_{(37)}$ = -1.05, p = 0.14) or in mPFC ($t_{(37)}$ = -0.67, p =0.51).

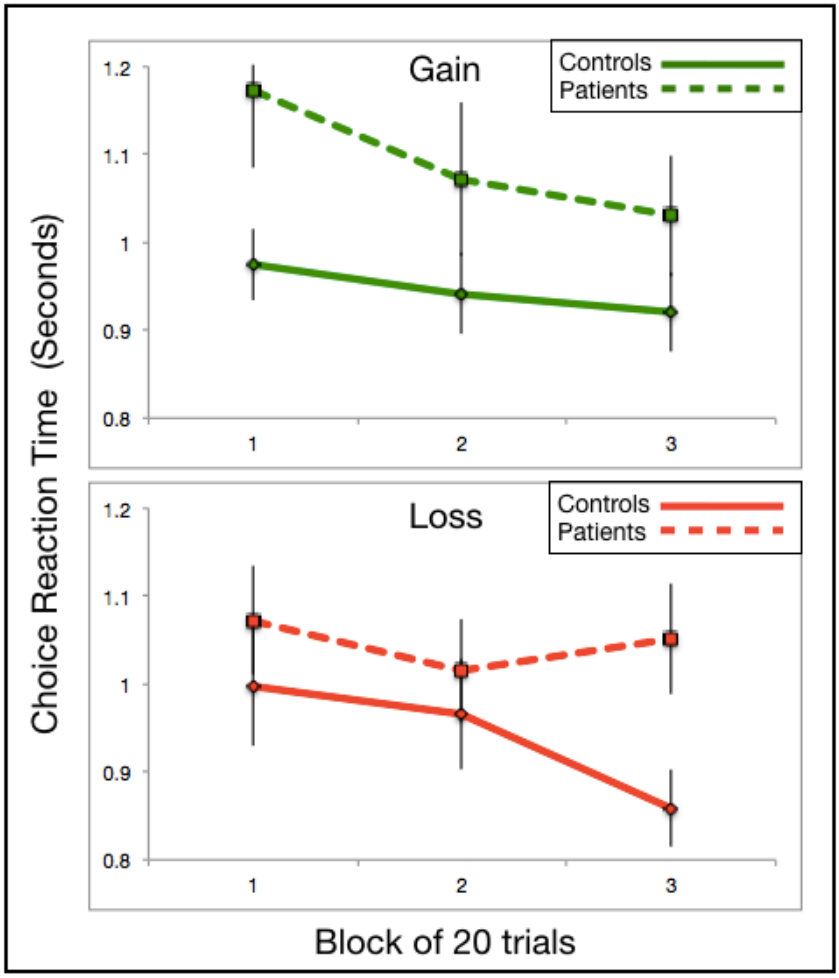

Supplementary Data Figure 2. Reaction time by block of 20 trials. Group average reaction times were calculated for each group of 20 trials for patients (SP) and controls (SC). There was no effect of block or condition, but there was an effect of group, as patients' reaction times were slower than controls.

Analyses Excluding Patient Subjects With Model Fits at Chance

Additional analyses sought to examine the data without two patients who performed at or worse than chance. We find that after removing the two subjects, the pattern of positive and negative results were largely similar. Behaviorally, the analysis of optimal choices by group and condition was as before: there still existed an overall effect of group but not condition or interaction for the optimal choice ($F(\text{group})=4.56$, $p = 0.03$; $F(\text{condition}) =0.07$, $p=0.79$; $F(\text{interaction})=0$, $p=0.95$). Effects on reaction time using this subgroup failed to reach significance ($F(\text{group})=1.94$, $p=0.16$; $F(\text{condition})=0.01$, $p=0.93$; $F(\text{interaction})=0$; $p=0.99$). In the model fits to choices, as before, we still found an effect of beta on group ($P<0.001$), but not for beta by group by condition ($P=0.70$), group by learning rate (0.15), or learning rate by condition (0.39). Corrected (FWE at $p < 0.05$) neural results also remained very similar (**Supplementary Data Figure 4**) for the key

group x condition interaction effects in the medial temporal lobe (-45, -24, -6, k = 139, t-max = 4.98). For the group x condition interaction, the large reported cluster on the right was present at uncorrected thresholds in the subgroup, but it just failed to survive whole brain correction, whereas a region in left putamen did survive whole brain correction (-18, -3, 0, k = 12, t-max = 2.86). As before, we also found group differences in the gain but not loss condition in prefrontal cortex (gain condition: 6, 57, -3, k = 57, t-max = 2.80; -24, 54, -6, k = 39, t-max = 2.72).

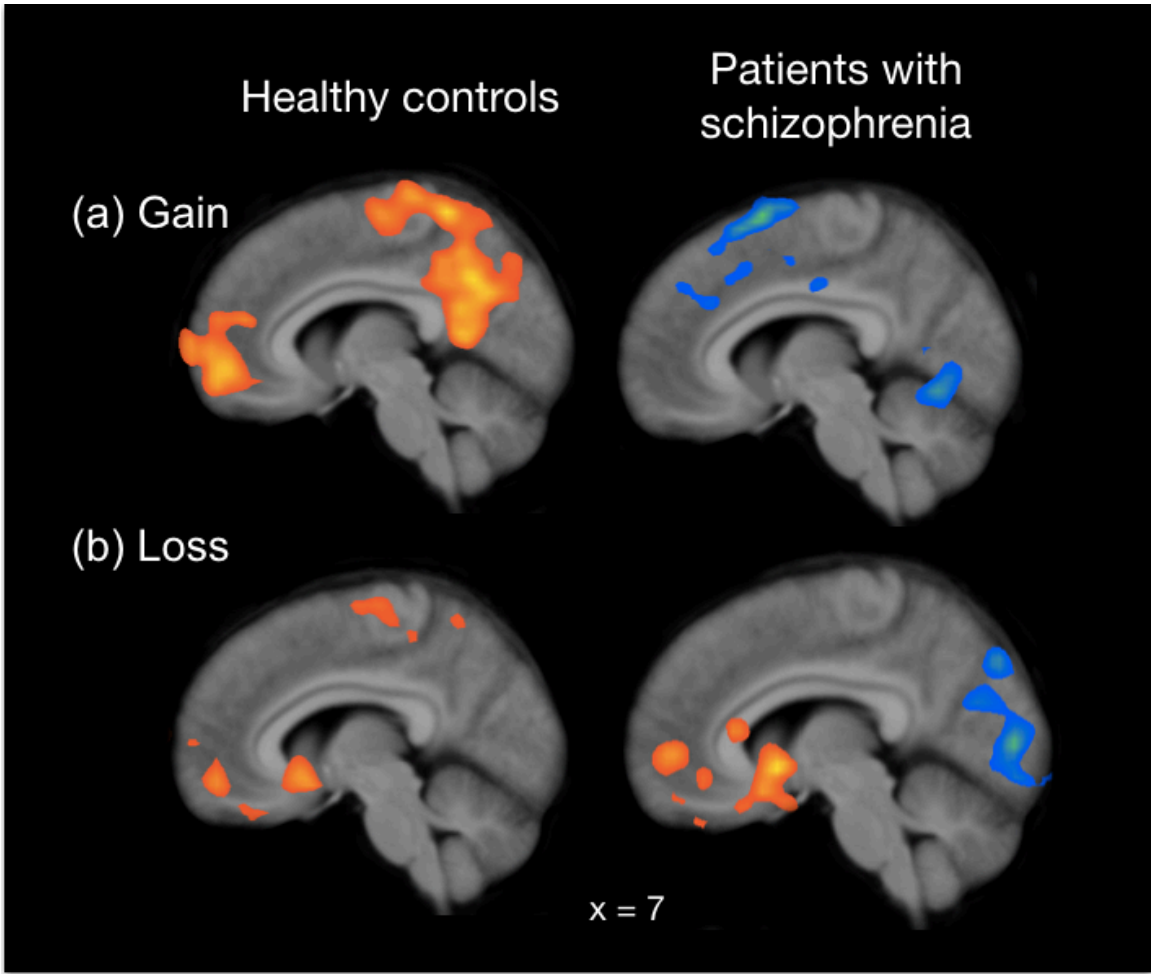

Supplementary Data Figure 3. Whole-brain fMRI results for feedback prediction error. Whole-brain results were FWE corrected at $p < 0.05$ during for the feedback prediction error learning signal when subjects learned from gains (a, above) and from avoiding losses (b, below).

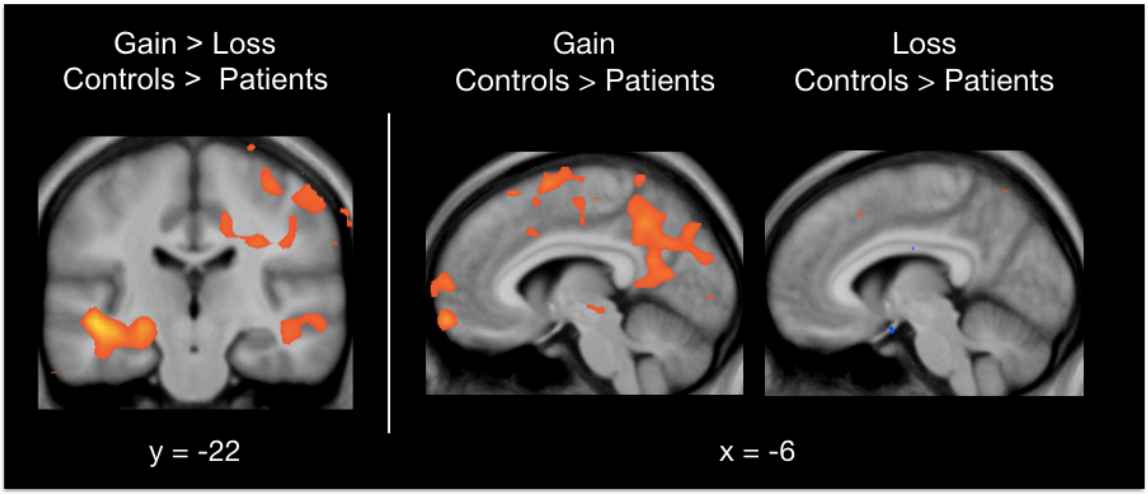

Supplementary Data Figure 4. Whole-brain fMRI results for feedback prediction error without 2 patient subjects. Results are whole brain corrected (FWE $p<0.05$), and are consistent with findings within the entire sample.

Region	Coordinates (Peak) X	Y	Z	K	T-Max
Superior Temporal Gyrus	-42	**-21**	**-9**	**829**	**4.93**
Superior Temporal Gyrus	-42	-21	-9	121	4.93
Sub-Gyral	-24	-18	-12	66	4.12
Inferior Temporal Gyrus	-54	-18	-36	38	3.79
Inferior Temporal Gyrus	-51	-6	-42	42	3.75
Middle Temporal Gyrus	-45	12	-36	32	3.68
Fusiform Gyrus	-39	-63	-12	54	3.33
Middle Temporal Gyrus	-54	0	-21	53	3.32
Parahippocampal Gyrus	-42	-33	-6	62	3.31
Fusiform Gyrus	-36	-45	-15	62	3.31
Inferior Temporal Gyrus	-30	-6	-45	21	3.26
Inferior Temporal Gyrus	-39	-3	-45	24	3.25
Inferior Temporal Gyrus	-54	-6	-30	36	3.23
Fusiform Gyrus	-36	-3	-21	94	3.22
Culmen	-15	-54	-18	29	2.90
Sub-Gyral	-45	-48	0	12	2.70
Uncus	-30	-18	-33	17	2.62
Uncus	-27	-3	-30	33	2.59
Precuneus	**12**	**-51**	**39**	**472**	**4.39**
Precuneus	12	-51	39	216	4.39
Precuneus	-15	-66	30	60	3.70
Precuneus	6	-51	54	65	3.36
Precuneus	-12	-45	30	34	3.30
Precuneus	-9	-54	39	68	2.98
Posterior Cingulate	9	-36	21	23	2.44
Superior Temporal Gyrus	**27**	**12**	**-36**	**571**	**4.21**
Superior Temporal Gyrus	27	12	-36	53	4.21
Putamen	18	3	0	42	4.07
Putamen	27	3	3	43	3.88
Uncus	18	6	-24	58	3.79

Caudate/Accumbens	3	9	-3	71	3.54
Precentral Gyrus	69	-9	27	58	3.33
Parahippocampal Gyrus/Amygdala	30	0	-18	35	3.30
Precentral Gyrus	60	-3	27	56	3.06
Precentral Gyrus	57	-6	12	50	3.06
BA 11/Orbitofrontal Cortexc	-9	24	-24	26	2.84
Lentiform Nucleus	18	-9	-6	13	2.61
Insula	42	0	3	18	2.48
Superior Temporal Gyrus	30	21	-39	18	2.44
Precentral Gyrus	**30**	**-15**	**30**	**508**	**3.51**
Precentral Gyrus	24	-9	69	63	3.45
Precentral Gyrus	36	-18	57	59	3.42
Middle Frontal Gyrus	30	-6	39	24	3.37
Precentral Gyrus	33	-15	72	45	3.31
Postcentral Gyrus	42	-27	36	58	3.22
Postcentral Gyrus	33	-30	33	45	3.20
Postcentral Gyrus	60	-18	51	26	3.06
Insula	45	-24	24	20	2.88
Postcentral Gyrus	48	-18	42	29	2.63
Postcentral Gyrus	27	-36	57	32	2.42

Supplementary Data Table 3. Clusters and subclusters in the gain > loss condition for control > patient subjects in response to feedback prediction error**.** Results are FWE corrected at $p < 0.05$, and are in MNI space. Local maxima (subclusters > 10 voxels) are listed beneath each cluster (bold).

References

1. Daw ND. Trial-by-trial data analysis using computational models. In: Delgado MR, Phelps EA, Robbins TW, eds. *Decision Making, Affect, and Learning* Oxford University Press; 2011.
2. Otto AR, Raio CM, Chiang A, Phelps EA, Daw ND. Working-memory capacity protects model-based learning from stress. *Proc Natl Acad Sci U S A* Dec 24 2013;110(52):20941-20946.
3. *Stan: A C++ Library for Probability and Sampling* [computer program]. Version Version 2.7.0; 2015.
4. Team SD. *Stan Modeling Language Users Guide and Reference Manual, Version 2.7.0*; 2015.
5. Sharp ME, Foerde K, Daw ND, Shohamy D. Dopamine selectively remediates 'model-based' reward learning: a computational approach. *Brain* Feb 2016;139(Pt 2):355-364.
6. Gelman A, Rubin DB. Inference from iterative simulation using multiple sequences. *Stat Sci* 1992(7):457-472.
7. Sharp ME, Foerde K, Daw ND, Shohamy D. Impaired goal-directed learning in Parkinson's disease. *Brain* 2015;In press
8. Freire L RA, Mangin JF. What is the best similarity measure for motion correction in fMRI time series? *IEEE Trans Med Imaging* 2002;21(5):470-484.
9. Van Snellenberg JX, Slifstein M, Read C, et al. Dynamic shifts in brain network activation during supracapacity working memory task performance. *Human brain mapping* Apr 2015;36(4):1245-1264.

10. Van Snellenberg JX, Conway AR, Spicer J, Read C, Smith EE. Capacity estimates in working memory: Reliability and interrelationships among tasks. *Cognitive, affective & behavioral neuroscience* Mar 2014;14(1):106-116.
11. Gold JM, Waltz JA, Matveeva TM, Kasanova Z, Strauss GP, Herbener ES, Collins AG, Frank MJ. Negative symptoms and the failure to represent the expected reward value of actions: behavioral and computational modeling evidence. *Arch Gen Psychiatry* Feb 2012;69(2):129-138.
12. Strauss GP, Frank MJ, Waltz JA, Kasanova Z, Herbener ES, Gold JM. Deficits in positive reinforcement learning and uncertainty-driven exploration are associated with distinct aspects of negative symptoms in schizophrenia. *Biological psychiatry* Mar 1 2011;69(5):424-431.
13. Waltz JA, Frank MJ, Wiecki TV, Gold JM. Altered probabilistic learning and response biases in schizophrenia: behavioral evidence and neurocomputational modeling. *Neuropsychology* Jan 2010;25(1):86-97.