

Journal section: Behavioral/Systems/Cognitive

Title: Falsifying the Reward Prediction Error Hypothesis with an Axiomatic Model

Abbreviated title: Testing the Reward Prediction Error Model Class

Authors: Robb B. Rutledge¹, Mark Dean², Andrew Caplin², and Paul W. Glimcher^{1,2}

Author addresses: ¹Center for Neural Science ²Department of Economics, New York University, New York, NY 10003 USA

Corresponding author address: Robb B. Rutledge, Center for Neural Science, New York University, 4 Washington Place, Room 809, New York, NY 10003, Phone: +1.212.998.3904, Fax: +1.212.995.4011, email: robb@cns.nyu.edu

Number of figures: 8

Supplemental material: consists of text, 5 figures, and 1 table

Number of pages: 52

Number of words: Abstract: 159
Introduction: 533
Discussion: 1782

Keywords: reward prediction error, reinforcement learning, dopamine, axiomatic model, neuroeconomics

Acknowledgments: This work was supported by a National Defense Science and Engineering Graduate Fellowship (R.B.R.) and National Institutes of Health Grants F31-AG031656 (R.B.R.) and RO1-NS054775 (P.W.G.). We thank Dan Burghart, Eric DeWitt, and Stephanie Lazzaro for helpful comments.

Note: Some of the data presented here are reported as a portion of another article currently in press at an economics journal (Caplin, Dean, Glimcher & Rutledge. Measuring beliefs and rewards: a neuroeconomic approach. Forthcoming at the *Quarterly Journal of Economics*). That article was provided to the editors. The cover letter submitted to the editors details the degree of overlap between these two manuscripts. To summarize briefly: Figures 1, 2, 3a, and 3c, which explain the task, axiomatic model, present an example of a traditional regression-based analysis, and show the anatomical definition and trial averages for the nucleus accumbens are similar to figures in the *QJE* paper. Finally, figure 4a, although a unique analysis, shows a result similar to a figure panel in the *QJE* paper. Of 80 statistical tests in Table 1 and the 12,000 statistical tests summarized in figure 5, analysis related to 10 of those tests are included in the *QJE* paper. All salience-related analyses including figure 7, the 224 statistical tests in Table 2, and the 33,600 statistical tests summarized in figure 8, and all axiom-related analyses for the other 10 brain areas besides the nucleus accumbens are unique to this manuscript. We also respectfully point out to our reviewers that there is very little overlap between the readers of the *Quarterly Journal of Economics* and the *Journal of Neuroscience*. While we share the conviction of many that double publication must be avoided at all costs, we urge the reviewers to consider the differences between the neurobiological analyses and conclusions and the economic theories developed as two separate goals of this joint interdisciplinary research program.

Abstract

Neuroimaging studies typically identify neural activity correlated with the predictions of highly parameterized models, like the many reward prediction error (RPE) models used to study reinforcement learning. Identified brain areas might encode RPEs or alternatively simply have activity correlated with RPE model predictions. Here we use an alternate axiomatic approach rooted in economic theory to formally test the entire class of RPE models on neural data. We show that measurements of neural activity from the striatum, medial prefrontal cortex, amygdala, and posterior cingulate cortex satisfy necessary and sufficient conditions for the entire class of RPE models. However, activity measured from the anterior insula falsifies the axiomatic model and therefore no RPE model can account for this activity. Further analysis suggests the anterior insula might instead encode something related to the salience of an outcome. As cognitive neuroscience matures and models proliferate, formal approaches that assess entire classes of models rather than specific model exemplars may take on increased significance.

Introduction

Our understanding of the natural world progresses through the development of explanatory models designed to capture compact descriptions of physical events. Within neuroscience, models tend to develop through a process of competitive evolution in which highly specified models are tested against each other. Other disciplines, including physics and economics, often employ an alternative approach, dividing the space of all possible models into subdomains and then attempting to falsify the hypothesis that one or more members of an entire class of models can account for a set of empirical observations. These model classes are typically defined by sets of testable rules called *axioms*. Popper (1959) argued that the most powerful test of any theory derives from formal efforts aimed at falsification. In this tradition, the axiomatic approach explicitly attempts to falsify entire model classes.

Dopamine neurons are thought to encode a reward prediction error (RPE) signal, the difference between experienced and predicted rewards. Numerous studies have fit specific parameterized RPE models to measurements of dopamine neuron activity (Schultz et al., 1997; Hollerman and Schultz, 1998; Nakahara et al., 2004; Bayer and Glimcher, 2005; Joshua et al., 2008; Matsumoto and Hikosaka, 2009) and fMRI measurements of neural activity in dopamine target areas (McClure et al., 2003; O'Doherty et al., 2003, 2004; Seymour et al., 2004; Abler et al., 2006; Li et al., 2006; Pessiglione et al., 2006; Behrens et al., 2008; D'Ardenne et al., 2008; Hare et al., 2008). Model competitions have shown that parameterized temporal-difference approaches (Sutton and Barto, 1990) better account for electrophysiological data (e.g., Schultz et al., 1997) than RPE models related to the approach of Rescorla and Wagner (1972).

Unfortunately, comparing correlation coefficients for different RPE models cannot tell us whether key features of dopamine-related activity are fundamentally incompatible with specific critical features of the entire class of all possible RPE models. The regression approach cannot, in principle, falsify the hypothesis that dopamine neurons encode some kind of RPE signal.

Caplin and Dean recently examined the necessary and sufficient properties of a RPE signal (Caplin and Dean, 2008a, 2008b), finding that any such signal must possess three critical features. They showed that if any one of these features is absent, the observed signal cannot represent a RPE regardless of whether it is correlated with parameterized RPE models. If all of these features are present then the measured signal meets criteria of both necessity and sufficiency for representing a RPE. By empirically testing these formal mathematical axioms, it is possible to test the entire class of RPE models for a neural signal measured from any brain area.

To axiomatically test the hypothesis that specific neural signals can encode RPEs, we used fMRI to measure blood-oxygen-level dependent (BOLD) activity as subjects played monetary lotteries for real money. We asked whether BOLD responses in specific candidate RPE brain areas satisfied the necessary and sufficient criteria for encoding a RPE signal. Any signal falsifying one or more axioms cannot *in principle* encode any type of RPE. Such a signal cannot be accounted for by any model in the entire RPE model class. We also tested whether any candidate RPE areas might alternatively encode the absolute value of the RPE signal, a quantity related to saliency.

Materials and Methods

Subjects. Fourteen paid volunteers participated in the experiment (9 women, 5 men, all right-handed, mean age = 26.0 years). All subjects participated in two scanning sessions. Two subjects were excluded from further analysis due to excessive head motion during the scanning sessions. Participants gave informed consent in accordance with the procedures of the University Committee on Activities involving Human Subjects of New York University.

Experimental task. Prior to scanning, subjects were endowed with \$100 in cash. Subjects also received a show-up fee of \$35 at the end of each scanning session, regardless of task earnings. On each trial, subjects chose between two monetary lotteries where the probability of each prize was represented by the area of that prize's slice (Fig. 1A). To test the axiomatic model, it was necessary to collect data with two prizes available (+\$5, -\$5) at a variety of probabilities (0% to 100% in 25% increments). Thus the observation set consisted of five lotteries and eight possible trial types (lottery-prize pairs). A lottery from the observation set appeared in every trial. To ensure that subjects usually chose from the observation set, the decoy lottery always had a lower mathematical expected value (ranging from \$1.25 to \$5 lower). The decoy set also included additional prizes (+\$0, -\$10). Subjects were given a choice between options to ensure that they were actively engaged in the task. After a 12.5 s fixation period, options were presented for 5 s. The fixation cross was extinguished, indicating that the subject had 1.25 s to make their selection by button press. After a 7.5 s delay period, the prize was revealed for 3.75 s as a change in the color of that prize's slice in the chosen lottery. If a subject failed to make a button press in the required time window, they lost \$10. Out

of 3,024 total trials completed, subjects missed 21 trials and chose the decoy lottery in 28 trials, completing 2,975 trials to the observation set.

Imaging. Imaging data were collected with a Siemens Allegra 3-Tesla head-only scanner equipped with a head coil from Nova Medical. T2*-weighted images were collected using an EPI sequence. We collected 23 slices oriented parallel to the AC-PC plane (TR = 1.25 s, TE = 30 ms, ascending interleaved order, 3 x 3 x 3 mm, 64 x 64 matrix in a 192-mm FOV). This volume provided coverage of the subcortical, frontal, and midbrain regions of interest while omitting part of the parietal lobe and the crown of the skull in all subjects. Each scan consisted of 396 images. The first four images were discarded to avoid T1 saturation effects. There were 16 choice trials during each scan. Each trial lasted 30 s (Fig. 1A). Each subject completed 13-16 scans over two sessions, with most subjects ($n = 9$) completing eight scans in each session. The data set consisted of 74,844 volumes, with an average of 130 min of functional data per subject. We also collected high-resolution T1-weighted anatomical images using a MP-RAGE pulse sequence (144 sagittal slices, TR = 2.5 s, TE = 3.93 ms, TI = 900 ms, flip angle = 8°, 1 x 1 x 1 mm, 256 x 256 matrix in a 256-mm FOV) for coregistration of functional data.

Data analysis. Functional imaging data were analyzed using BrainVoyager QX (Brain Innovation, Maastricht, the Netherlands), with additional analyses performed in MATLAB (MathWorks, Natick, MA) and STATA (StataCorp, College Station, TX). We sinc-interpolated functional data in time to adjust for staggered slice acquisition. We corrected for any head movement by realigning all images to the first volume of the session using six-parameter rigid-body transformations. We detrended and high-pass filtered (cutoff of 3 cycles per scan) to remove low-frequency drift in the signal. We then

co-registered images to each subject's high-resolution anatomical scan, rotated into the AC-PC plane, and normalized into Talairach space using piecewise affine Talairach grid scaling using trilinear interpolation. Data were spatially and temporally unsmoothed, except for the group random-effects analysis.

To demonstrate the standard regression approach, we performed group random-effects analysis using the summary statistics approach. For this analysis we spatially smoothed all data with an 8-mm FWHM Gaussian kernel. The regression model consisted of a single regressor of interest with the “predicted RPE” on each trial during the outcome period. This was defined for these purposes as the difference between the reward received in dollars and the expected value of the lottery. Three additional regressors modeled the options, button, and outcome onset for all trials. All four regressors were convolved with the canonical two-gamma ($\tau_1 = 6$ s, $\tau_2 = 15$ s, ratio of peak to undershoot = 6) hemodynamic impulse response function. A statistical map was then generated for the regressor of interest using one-sample *t*-tests. This map is shown for demonstration purposes without any minimum cluster threshold or corrections for multiple comparisons (Fig. 1*B*).

For further analysis, we independently defined anatomical regions of interest (ROIs) in individual subjects for 11 brain regions: the nucleus accumbens, anterior insula, caudate, putamen, medial prefrontal cortex, amygdala, posterior cingulate cortex, thalamus, ventral tegmental area, substantia nigra, and habenula. These regions were chosen because they have been found to have activity consistent with specific RPE models in previous neuroimaging and neurophysiological studies. Criteria for these structural definitions, primarily using those established by the Center for Morphometric

Analysis (Rademacher et al., 1992; Caviness et al., 1996), are described in the Supplemental Data and distributions for these definitions across subjects are shown (Fig. 3 and supplemental Fig. 1 and 3).

Our ROIs were largely located in subcortical and midbrain areas. The amygdala and posterior cingulate cortex ROIs were located near the boundaries of our acquisition volume, making these ROIs particularly susceptible to artifacts from the motion correction algorithm. To limit these artifacts, we excluded from our ROIs any voxel from a given scan for which the standard deviation of percent signal change exceeded 2%, a degree of variance incompatible with a continuous BOLD signal (see supplemental Table 1). In practice, this excluded < 5% on average for all structures except the amygdala (12%) and posterior cingulate cortex (23%). We limited the effects of motion on BOLD activity in individual voxels using a regression model that included the six motion predictor regressors and their temporal derivatives. We then averaged data across each anatomical ROI to produce a mean time course for each ROI which was converted to percent signal change using two baseline TRs as indicated in the text and figures.

We made no assumptions about the shape of the hemodynamic response functions in our anatomical ROIs, but removed correlations between timepoints at the subject and trial type level using an AR4 autoregressive model while maintaining consistent timepoint averages. We then averaged activity within the 5-TR window, weighting each timepoint equally. We computed parameter estimates by ordinary least squares for each of the eight trial types for the 2,975 trials in the observation set controlling for subject-level differences in activity. We evaluated the following axiomatic RPE model and RPE absolute value model, testing for differences between parameter estimates using Wald

tests of linear restriction. We tested the robustness of our results by evaluating a wide range of baselines and starting times for the 5-TR analysis window.

Axiomatic RPE model. To determine whether the BOLD signal measured in the striatum and other possible RPE areas meets the criteria of necessity and sufficiency for encoding a RPE signal, we formally tested the RPE hypothesis using the axiomatic model developed by Caplin and Dean (2008a). This approach makes no specific assumptions about the precise form of subjective variables like “reward” and “expectation” that are not part of the RPE hypothesis, but that the traditional regression approach requires. Using this model, we can thus explicitly test whether a given neural signal falsifies or satisfies the three conditions of necessity and sufficiency for the entire class of RPE models. For example, all RPE models assume that a RPE signal responds similarly to any fully anticipated outcome, whether it be winning or losing \$5 or winning an apple or an orange, and the model’s third axiom formally captures that intuition. Surprisingly, this assumption has never been tested on dopamine neurons for prizes with relatively similar sensory properties, like apple juice and orange juice. If any neural signal does not satisfy this or either of the other two axioms then it cannot, in principle, represent a RPE signal. For the two-prize case we tested, the three axioms are necessary and sufficient criteria for the RPE model class (Caplin et al., in press).

We tested our measures of neural activity against the three axioms: consistent prize ordering (axiom 1), consistent lottery ordering (axiom 2), and no surprise equivalence (axiom 3). The three axioms are as follows, where $\delta(z, p)$ is neural activity associated with receiving prize z (e.g., winning \$5) from lottery p (e.g., 50% probability of winning \$5). $\delta(z)$ is the one-prize “lottery” where prize z has 100% probability:

- 1) $\delta(z, p) > \delta(z', p) \Rightarrow \delta(z, p') > \delta(z', p')$
- 2) $\delta(z, p) > \delta(z, p') \Rightarrow \delta(z', p) > \delta(z', p')$
- 3) $\delta(z') = \delta(z)$

Axiom 1: consistent prize ordering. Consider two different lotteries p and p' (they have 25% and 75% probabilities of winning \$5, respectively, in Figure 2A). Caplin and Dean (2008a) demonstrated that if a RPE signal responds with higher activity to winning than losing \$5 from lottery p , then it must be the case, according to the RPE hypothesis, that winning \$5 has a higher experienced reward than losing \$5. Therefore, the signal must also respond with higher activity to winning than losing \$5 from lottery p' . Figure 2A shows a hypothetical result that would falsify this first criterion. Hypothetical neural activity (for example, BOLD activity from some brain area) is plotted against the probability of winning \$5; each point represents activity associated with receiving a particular prize (+\$5 in red, -\$5 in blue) from one of the five lotteries in the observation set. Open circles represent unobservable outcomes; for example, observing the activity associated with losing \$5 when the probability of winning \$5 is 100% is impossible. Higher activity for winning than losing \$5 from lottery p implies that winning \$5 has the higher experienced reward (Fig. 2A). Higher activity for losing than winning \$5 from lottery p' implies the opposite, and this contradiction violates the first axiom. Any crossing of the red and blue lines contradicts consistent prize ordering and proves that the activity under study cannot, in principle, encode any form of RPE signal. This is true for any two prizes, for example, comparing apples and oranges; for a RPE signal, if the activity is higher for apples than oranges for one lottery, it must be higher for all lotteries.

Axiom 2: consistent lottery ordering. Consider again lotteries p and p' (they again have 25% and 75% probabilities of winning \$5, respectively, in Figure 2B). A RPE signal

that responds with lower activity to losing \$5 from lottery p' than from lottery p implies that, according to the RPE hypothesis, p' has higher predicted reward. Therefore, it must also respond with lower activity to winning \$5 from lottery p' than from lottery p . Figure 2B shows a violation of this axiom. Lower activity for losing \$5 from lottery p' than lottery p implies that lottery p' (the lottery with a 75% chance of winning \$5) has the higher predicted reward. Higher activity for winning \$5 from lottery p' than from lottery p implies the opposite. The blue line's downward slope implies that the lotteries with a higher probability of winning \$5 have higher predicted reward. The red line's upward slope implies the opposite. For any two lotteries, any difference in signs of slopes between red and blue lines contradicts consistent lottery ordering and proves that the activity under study cannot, in principle, encode any form of RPE signal. The activity for a prize received from two different lotteries can be identical, implying that two lotteries have equal predicted reward, as long as the activity for receiving the other prize from the same two lotteries is also identical.

Axiom 3: no surprise equivalence. The final criterion of necessity and sufficiency identified by Caplin and Dean (2008a) was that RPE signals must respond identically to all fully predicted outcomes, conditions under which the reward prediction error is zero. If there is no reward prediction error, the signal must always generate the same response regardless of the prediction. Consider the two one-prize “lotteries” shown as the filled endpoints of the red and blue lines in Figure 2C. If, as shown in the plot, the signal responds with less activity to losing than winning \$5 when both outcomes are fully anticipated, this violates the third axiom and proves that the activity under study cannot, in principle, encode any form of RPE signal.

These three representational constraints that must be obeyed by any member of the class of RPE models (Caplin and Dean, 2008a, 2008b; Caplin et al., in press), whether a Rescorla-Wagner model, a temporal-difference model, a RPE model with a high or low learning rate, or a RPE model with any arbitrary utility function. If an observed neural signal fails to meet any of these criteria, then the proposition that it can encode a RPE signal can be considered formally falsified. In contrast, a neural signal that demonstrates all three properties is one that, in the two-prize case, meets the sufficient criteria for encoding a RPE signal (as proven for the two-prize case by Caplin and Dean, 2008). A pattern of activity satisfying all three axioms is shown in Figure 2D.

We do not make any assumptions about the magnitude of experienced and predicted rewards for prizes or lotteries, nor about the hemodynamic response function of subjects or brain areas. However, since our analysis is performed at the group level, we do assume that subjects have the same ordering over prizes and lotteries. For example, we assume that all subjects *either* prefer winning to losing \$5 or alternatively prefer losing to winning \$5. Although the theory itself does not require this, if we additionally assume that subjects prefer winning to losing \$5 and also prefer lotteries with a higher probability of winning \$5, then we can predict that the axioms will be satisfied *specifically* in the way indicated in the left-most column in Table 1. They could be satisfied in many other ways, including if all the signs in the left-most column were reversed (Caplin and Dean, 2008). We looked for signals that satisfied the axioms in this manner by counting the number of tests with the predicted sign at $p < 0.05$ for a wide range of baselines and analysis windows. Baselines were selected around the end of the fixation period and the end of the delay period. BOLD activity in all areas was observed to be relatively similar

across trial types during these periods. A range of analysis windows were tested starting before the outcome period and lasting into the next trial. For the habenula, we also tested the axiomatic RPE model with the signs of all statistical tests reversed, as electrophysiological data has suggested this area may encode a sign-reversed RPE signal (Matsumoto and Hikosaka, 2007).

RPE absolute value model. To test whether a signal can represent the magnitude of the RPE signal for the two-prize case we tested, we must make two assumptions about how the RPE is constructed. First, we assume that the RPE is the mathematical difference between the experienced and predicted reward. Second, we assume (for the two-prize case) that the predicted reward is equal to $p_z u_z + (1-p_z) u_{z'}$ where p_z , u_z , and $u_{z'}$ are the probability and utility of prize z and the utility of prize z' respectively. Thus the RPE absolute value when prize z is received is $(1-p_z) \text{abs}(u_z - u_{z'})$. When prize z' is received, it is $(1-p_{z'}) \text{abs}(u_{z'} - u_z)$. Since the second term is always the same for the two-prize case, the RPE absolute value should be a decreasing function of probability. We test whether activity decreases with prize probability with the following condition:

- 1) The activity associated with receiving prize z from lottery p is higher than for receiving prize z' from lottery p' if and only if the probability of receiving prize z from lottery p is higher than the probability of receiving prize z' from lottery p' .

In this way we examine the possibility that how surprising an outcome is whether for good or bad, a property related to salience, can be encoded by the BOLD response in a particular brain area. For the habenula we also tested for a sign-reversed RPE absolute value signal.

Results

Traditional regression-based analysis

A number of previous studies have examined the RPE hypothesis by selecting a fully parameterized member of the RPE model class and correlating some element of the model with measured BOLD activity. We first completed a standard random-effects regression analysis of this type (Fig. 1B) to allow comparison with the results of our axiomatic RPE model analysis. To accomplish this we had to make several assumptions about concepts like “reward” and “expectation”, variables which cannot be measured directly. We therefore assumed, as have previous studies (e.g., D’Ardenne et al., 2008), that the utility function for gains and losses was a linear function of monetary reward with no change in slope at the origin (Pascal’s utility function), and that the predicted reward was equal to the utilities of the prizes weighted by their objective probabilities (expected utility theory’s independence axiom). These assumptions imply that the RPE signal would be proportional to the difference in dollars between the outcome received and the lottery’s expected value. We also assumed, as have previous studies (e.g., Li et al., 2006), that the BOLD response in all areas would follow the canonical two-gamma hemodynamic impulse response function which has been well validated in sensory and motor cortex (e.g., Vazquez and Noll, 1998; Friston et al., 1999). We found that BOLD activity in the striatum (including parts of the nucleus accumbens, putamen, and caudate) was significantly correlated ($p < 0.001$, uncorrected) with the predicted RPE specified in this way. This result is consistent with numerous previous studies (e.g., McClure et al., 2003; O’Doherty et al., 2003, 2004; Pessiglione et al., 2006) that have shown that activity in this area is well correlated with specific RPE models. At a more liberal threshold ($p <$

0.01), BOLD activity in the medial prefrontal cortex was also correlated with the predicted RPE, but not activity in other candidate RPE areas including the anterior insula, amygdala, and posterior cingulate cortex. While these data clearly indicate that BOLD activity in the striatum is correlated with the predictions of this particular RPE model, they cannot tell us whether the data is actually compatible with the RPE hypothesis. Is the observed correlation limited by a fundamental and insurmountable mismatch between critical properties of the signal and the model? To answer that question we turned next to a test of the necessary and sufficient signal properties required for a RPE representation.

Neuroimaging test of the RPE hypothesis by the axiomatic method

Neuroimaging studies have identified activity in numerous brain areas that is correlated with the predictions of particular RPE models. To test the hypothesis that BOLD activity in these brain areas can actually encode a RPE signal, we first anatomically defined ROIs and then computed estimates of the average BOLD activity for each of the eight trial types from the observation set of lotteries. This allowed us to produce plots of the kind shown in Figure 2 for each brain area. We then performed statistical tests on these data in an effort to falsify one or more of the axioms.

We first extracted BOLD responses from the nucleus accumbens and the anterior insula in all subjects (Fig. 3*A-B*), both regions identified as possible RPE areas in previous studies (e.g., Pessiglione et al., 2006; Voon et al., 2010). We then plotted the average BOLD responses for the eight trial types (Fig. 3*C-D*) converted to percent signal change relative to a baseline selected as the last two TRs of the fixation period (TR 9-10). The outcome of each trial was presented on the screen for 3 TRs (TR 22-24). Due to the

lag in the hemodynamic response (approximately 5 s or 4 TRs), we specified our initial analysis window as TR 26-30 (later analyses relax this assumption). For each brain area, we then estimated parameters for each trial type, averaging activity across the analysis window, weighting timepoints equally. Our methodology assumes that all subjects have the same ordering of prizes and lotteries. For example, we assume that all subjects either prefer winning to losing \$5 or alternatively all prefer losing to winning \$5. We make no assumptions about the shape of the hemodynamic response function in different brain areas. The resulting parameter estimates are plotted for the nucleus accumbens (Fig. 4A) and anterior insula (Fig. 4B). For each area, we then performed 10 Wald tests of linear restriction (Wald, 1943) on the relations between these parameter estimates which instantiate the three critical axiomatic criteria. Test results are shown in Table 1.

For BOLD activity in the nucleus accumbens (Fig. 4A), a subregion of the ventral striatum, axiom 1 is satisfied, with higher activity for winning than losing \$5 for the three two-prize lotteries (all $p < 0.001$). Axiom 2 is satisfied with all lines significantly downward sloping (all $p < 0.05$). Finally, axiom 3 is satisfied, with activity not significantly different for the two fully anticipated outcomes ($p = 0.29$). This signal thus satisfies all three necessary and sufficient conditions of the axiomatic RPE model and can unambiguously encode a RPE signal.

Perhaps surprisingly, the data for the anterior insula indicate a very different conclusion (Fig. 4B). Axiom 1 is falsified at $p < 0.05$; the activity is higher for losing than winning \$5 from the 75% lottery ($p < 0.001$), but this is not true for the other lotteries at $p < 0.05$. Axiom 2 is also falsified at $p < 0.05$ in two different ways: activity is higher for losing \$5 from the 50% than the 25% lottery ($p = 0.032$) but this is not true for

winning \$5, and activity is lower for winning \$5 from the 75% than the 25% lottery ($p < 0.001$), but this is not true for losing \$5. Finally, axiom 3 is also falsified; the activity is significantly higher for losing \$5 than winning \$5 for the fully anticipated outcomes ($p < 0.001$). Therefore, this signal falsifies all three necessary and sufficient conditions of the axiomatic RPE model, and cannot possibly encode any type of RPE signal. BOLD activity in the anterior insula, despite the fact that it is correlated with the predictions of specific RPE models in some studies, cannot in principle encode a RPE signal under the conditions we have examined.

We also tested several other areas that previous studies suggest might encode RPE signals. Anatomical definitions and BOLD time series for six additional areas are shown in supplemental Figure 1 with tests presented in Table 1. BOLD activity in the caudate also satisfies all three axioms at $p < 0.05$ and can encode a RPE signal. Activity in the putamen, medial prefrontal cortex, and amygdala, but not the posterior cingulate cortex, satisfies the first axiom at $p < 0.05$. However, activity in all four areas falsifies the second axiom, so these signals cannot, in principle, represent a RPE if the representation is constrained at the time of our analysis window relative to this specific baseline. The signal in the thalamus also falsifies all three axioms and cannot encode a RPE signal for this specific baseline and analysis window.

Although testing the axioms requires no assumptions about the precise ordering of prizes or lotteries, we predicted that subjects would both prefer winning to losing \$5 and would prefer lotteries with a higher probability of winning \$5 and that BOLD activity would be related to this preference. This led us to predict that the axioms would be satisfied in the specific way specified in the left-most column in Table 1. For any lottery,

winning \$5 should lead to higher activity than losing \$5. Winning either prize should lead to lower activity from lotteries with a higher probability of winning \$5. Although, for example, the medial prefrontal cortex falsifies the second axiom, most of the tests for this signal had the predicted sign. Because measurements of BOLD activity are noisy, whether or not a signal satisfies the axioms might depend on the baseline and analysis window used to estimate the responses. To test this possibility, and to examine the robustness of our findings, we analyzed signals for a wide range of baselines and analysis windows.

Assessing the robustness of axiomatic RPE model tests

In the preceding section, to estimate neural signals to test the axiomatic RPE model, we averaged the signal across a 5-TR analysis window (TR 26-30) beginning around the expected peak of the hemodynamic response. We also converted the raw signal to percent signal change using the last two TRs of the fixation period as a baseline. This standard practice in fMRI time series analysis adjusts for magnetic field drift that detrending and high-pass filtering fail to correct. To assess the robustness of our results, we counted the number of tests which were significant at $p < 0.05$ with the predicted sign for a range of baselines and analysis windows. We plot the results of this analysis in Figure 5, with results for 11 possible baselines (including no baseline) plotted against the starting time of the 5-TR analysis window. Color indicates the number of significant tests with the predicted sign for that particular baseline and analysis window, with more significant tests indicated in red colors and fewer in blue colors.

BOLD activity in the nucleus accumbens, caudate, amygdala and posterior cingulate cortex (Fig. 5) had the predicted sign for all 10 tests of the axiomatic model. Swaths of red in Figure 5 indicate that most tests had the predicted sign for a range of baselines and analysis windows, suggesting that the RPE model is robustly appropriate for these areas. Signals measured from each of these areas satisfy the axioms in exactly the way predicted, and these signals thus can encode a RPE (supplemental Fig. 2). Although BOLD activity in the putamen and medial prefrontal cortex did not have the predicted sign for all 10 tests for any baseline or analysis window (Fig. 5D-E), there are signals for both areas which satisfy all three axioms at $p < 0.05$ (supplemental Fig. 2). For example, for a baseline TR 8-9 and analysis window TR 28-32, the medial prefrontal cortex satisfies all three axioms; all tests have the predicted sign except tests 2.1 and 2.2. The signs of these two tests are both equal (rather than minus) and therefore satisfy the second axiom. The putamen signal for a baseline TR 22-23 and analysis window TR 25-29, for example, also satisfies all three axioms. For all these areas, the majority of tests have the predicted sign for a wide range of analysis windows and baselines.

In contrast, the anterior insula does not appear to satisfy the criteria for a RPE for any baseline or analysis window. There is only a single baseline and analysis window over the entire range tested for which this area (Fig. 5B) even has the predicted sign for the majority (six) of the tests. There exists no baseline and no analysis window within the range of TR 22-36 for which all three axioms are satisfied for either the anterior insula or the thalamus; the signal from both areas cannot possibly encode a RPE representation under the conditions we examined and this result is robust to choice of baseline and analysis window.

Because our traditional random-effects regression analysis revealed correlations with our particular RPE model only in the striatum (Fig. 1B) and, at a more liberal threshold ($p < 0.01$, uncorrected), the medial prefrontal cortex, we were surprised to see several other brain areas from which signals satisfied the axiomatic RPE model, some of which (amygdala and posterior cingulate cortex) are rarely identified in neuroimaging studies of the RPE hypothesis. Plotting the average BOLD response to positive and negative outcomes from the three two-prize lotteries reveals that the hemodynamic responses in the amygdala and posterior cingulate cortex, and also the medial prefrontal cortex, bear little similarity to the canonical hemodynamic response function (Fig. 6). For example, all three signals terminate at a higher level than they started. This may suggest that prior regression-based analyses have failed to identify several of these RPE signals due to incorrect assumptions about hemodynamics. In fact, we note that even in the nucleus accumbens the hemodynamic prediction appears to fit the data poorly, with the signal rising initially for all outcomes and then dipping well below the starting level. As shown here, our analysis methods circumvent these issues.

Although imaging the dopaminergic midbrain structures is notoriously difficult and few studies have reported success at identifying possible RPE signals in the midbrain (although see D'Ardenne et al., 2008), we tested whether signals extracted from the ventral tegmental area and substantia nigra might satisfy the axiomatic RPE model. We found no evidence of RPE signals in *BOLD responses* in either area (supplemental Fig. 4), although we cannot address here whether spiking patterns are consistent with the RPE theory. We also tested whether BOLD responses in the habenula might encode a RPE signal or alternatively a sign-reversed RPE signal, as a recent electrophysiological study

has suggested is carried by spiking activity (Matsumoto and Hikosaka, 2007). We found no evidence for either a RPE or a sign-reversed RPE BOLD signal (supplemental Fig. 4). Supplemental Figure 3 displays ROI definitions and trial averages for all three areas.

Understanding the anterior insula: RPE absolute value signals

Given the previous reports indicating that BOLD activity in the anterior insula is often correlated with the predictions of specific RPE models and our finding that anterior insula activity cannot serve as a RPE signal, we examined whether the signal in the anterior insula might encode some other reward-related information. One possibility is that the signal encodes something about how surprising or salient an outcome is to a subject. Although there is little formal agreement regarding the definition of the term “salience”, one natural assumption would be that an outcome is more salient if it is less likely. In our experimental setting, a greater response to an outcome with lower probability is equivalent to encoding the absolute value of the RPE signal, if we assume that subjects form their expectations by linearly combining the utilities of prizes weighted by their probabilities. A RPE absolute value (“salience”) model has a testable restriction that it places on our two-prize data set. Activity associated with receiving a prize z must be higher than activity for z' if and only if the probability of receiving prize z is less than the probability of receiving prize z' . Testing this restriction requires evaluating the 28 pairwise comparisons between all pairs of outcomes.

We replot the parameter estimates against the probability of the prize received for the nucleus accumbens and anterior insula for baseline TR 22-23 and analysis window TR 24-28 (Fig. 7). For this baseline and analysis window, we found that the anterior

insula signal was largely a decreasing function of prize probability as would be predicted for a salience signal (Fig. 7B). This was not the case for the nucleus accumbens (Fig. 7A). We evaluated the 28 tests of the RPE absolute value model for eight brain areas (Table 2). We found that in the anterior insula, 27 of 28 tests had the predicted sign at $p < 0.05$. In the nucleus accumbens only 15 of 28 tests had the predicted sign at $p < 0.05$.

We conducted our tests of the RPE absolute value model on neural signals estimated with a range of baseline and analysis windows, as we did for the axiomatic RPE model. For each baseline and analysis window, we counted how many of the 28 statistical tests were significant at $p < 0.05$ with the predicted sign. In Figure 8, we plot the results of this analysis, using the conventions in Figure 5. While most tests have the predicted sign for a range of baseline and analysis windows for the anterior insula (Fig. 8B), this is not the case for the nucleus accumbens (Fig. 8A) or amygdala (Fig. 8F). Some evidence for a RPE absolute value signal was present in other areas, including the thalamus and caudate in particular (Fig. 8) and the substantia nigra (supplemental Fig. 5).

Discussion

Neuroimaging studies have identified numerous brain areas where BOLD activity is correlated with the predictions of highly specified RPE models, most frequently including the ventral striatum. Here we used an axiomatic model to show that BOLD activity in the nucleus accumbens, a subregion of the ventral striatum, satisfies necessary and sufficient conditions for the RPE model class. This signal can represent RPEs in tasks like ours, as previous studies have suggested but never formally tested. This is also true for signals measured from the other subregions of the striatum, the caudate and putamen, as well as

the amygdala, medial prefrontal cortex, and posterior cingulate cortex. For each area, there must exist some RPE model that accounts for the BOLD responses measured in our experiment. This axiomatic approach required none of the auxiliary assumptions about unobservable variables like reward and expectation necessary with the traditional regression approach and rather than simply looking for any correlation it specifically tests the properties critical to the RPE model class. The importance of this distinction is highlighted by the fact that traditional regression-based studies have also found activity in the anterior insula correlated with the predictions of specific RPE models. This observed correlation could arise either because the BOLD activity encodes some type of RPE signal or alternatively because the BOLD activity is correlated with features of some RPE models. We show here that the latter is the case. The signal in the anterior insula falsifies the axiomatic model and cannot, in principle, encode a RPE signal under these conditions. This activity may instead encode the absolute value of the RPE signal, a signal that is correlated with the predictions of some RPE models, and this quantity may be referred to as salience.

Reward prediction error models and the anterior insula

Perhaps the most surprising result presented here is that the signal measured in the anterior insula falsifies the axioms of the RPE model. There is no way of defining or parameterizing a RPE model to account for the BOLD signal measured in this area in our task. This is a critical logical feature of the axiomatic approach, allowing us to unambiguously contradict the hypothesis presented in several previous studies that BOLD activity in the anterior insula might encode some kind of RPE signal (Seymour et

al., 2004; Pessiglione et al., 2006; Wittman et al., 2008; Voon et al., 2010). This does not mean that the findings reported in those papers are in error. There can be no doubt that the activity in the anterior insula is correlated with the predictions of some RPE models. However, our tests suggest that the limits of those observed correlations arise from properties of the signal that are *fundamentally* incompatible with any RPE representation.

Reward prediction error absolute value models

Many studies identifying correlations with RPE models involved painful stimuli (Seymour et al., 2004) or financial losses (Pessiglione et al., 2006; Voon et al., 2010), which may be particularly salient outcomes. Neuroimaging studies have also found evidence for a role for the anterior insula in representing uncertainty (Huettel et al., 2005; Grinband et al., 2006), prediction errors related to the variance in rewards (Preuschoff et al., 2006, 2008), and in processing salient stimuli (Jensen et al., 2007; Seeley et al., 2007). Ullsperger and von Cramon (2003) identified the anterior insula as having greater activity for negative than positive feedback in a task in which negative feedback is the less frequent (and thus more salient) class of feedback.

To explore this possibility, we tested BOLD activity in the anterior insula with a RPE absolute value model. Anterior insula activity measured in our task almost completely satisfied this model. BOLD activity in the anterior insula is largely a decreasing function of prize probability and might encode the absolute value of RPE, consistent with some notions of salience.

This possibility is of particular importance because it has been argued that dopamine neurons, and BOLD activity in dopamine target areas, may actually encode

salience either in addition to or instead of RPEs (Berridge and Robinson, 1998; Redgrave et al., 1999; Horvitz, 2000; Zink et al., 2003, 2004). A recent electrophysiology study has identified an anatomically distinct subpopulation of neurons in dorsolateral substantia nigra that increases their activity in response to unexpected appetitive and aversive events (Matsumoto and Hikosaka, 2009) and it has been suggested that these are dopamine neurons although this has not been verified pharmacologically or histologically. Another study has identified a subset of dopamine neurons in anesthetized rats in the ventral tegmental area that responds to aversive events (Brischoux et al., 2009), although the relationship of these findings to the hypothesis that all dopamine neurons encode RPEs remains unclear. However, although we found that BOLD activity in the anterior insula may encode the absolute value of the RPE, we did not find signals in any other area that satisfied the constraints of the RPE absolute value model.

Relating BOLD activity to dopamine

Electrophysiological results suggest that midbrain dopamine neurons encode a RPE signal (Schultz et al., 1997; Hollerman and Schultz, 1998; Nakahara et al., 2004; Bayer and Glimcher, 2005; Joshua et al., 2008; Matsumoto and Hikosaka, 2009; Zaghoul et al., 2009) and the regions in which we identified BOLD activity that could encode RPEs are all regions to which dopamine neurons are known to project. Although dopaminergic drugs influence learning rates associated with RPE signals (Rutledge et al., 2009; Voon et al., 2010) and also modulate the magnitude of BOLD activity for putative RPE signals in the striatum (Pessiglione et al., 2006; Voon et al., 2010), it is important to note that we

cannot conclude that the RPE signals we measured using fMRI are due to dopaminergic activity.

Although we did not find that BOLD activity in the midbrain dopamine structures can encode a RPE signal, imaging these structures is notoriously difficult. Whether this reflects the widely acknowledged discrepancy between BOLD activity and spiking activity or the limitations of our imaging protocol is unclear. D'Ardenne and colleagues (2008) found evidence for positive (but not negative) RPE signals in the ventral tegmental area using high-resolution imaging and midbrain-specific alignment algorithms. The habenula is another difficult-to-image structure which might encode a sign-reversed RPE signal (Matsumoto and Hikosaka, 2007), although we were unable to find evidence for this here using a standard imaging protocol.

Medial prefrontal cortex, amygdala, and posterior cingulate cortex

Our finding of RPE signals in all three regions of the striatum (nucleus accumbens, caudate, and putamen) is not surprising. However, there are far fewer reports consistent with RPE signals in the medial prefrontal cortex (Behrens et al., 2008), amygdala (Yacubian et al., 2006), and posterior cingulate cortex (de Bruijn et al., 2009), although electrophysiological studies have found activity consistent with RPE signals in all three areas (McCoy et al., 2003; Belova et al., 2007; Matsumoto et al., 2007). Our traditional random-effects correlation analysis using a typical RPE model convolved with the canonical two-gamma hemodynamic response function (HRF) revealed correlations in the striatum and, at a very liberal threshold, the medial prefrontal cortex, but not the amygdala or posterior cingulate cortex. Inspection of the BOLD time series in Figure 6

suggests one possible explanation. The apparent HRFs for the amygdala and posterior cingulate cortex appear to bear little similarity to the canonical two-gamma HRF used most commonly in standard regression analyses. This is also true for the nucleus accumbens and medial prefrontal cortex, although BOLD responses in these areas are likely strong enough to still produce significant correlations. This finding may suggest that future regression-based studies of reward areas should either use HRFs demonstrated to be appropriate for the regions under study or should use deconvolution or autoregressive methods that are less susceptible to the differences in the HRFs between brain areas.

Another problem with the standard regression approach is apparent in the BOLD time series for all six areas that can encode RPEs. Previous studies have always assumed that the responses to outcomes received from one-prize “lotteries” (like a tone followed by a juice reward) are intermediate between responses to positive and negative outcomes, but inspection of the BOLD time series reveals that this is not the case for our data (Fig. 3C and supplemental Fig. 1). Although an advantage of our axiomatic methodology is that we make no assumptions about how these responses relate to responses for two-prize lotteries, these qualitatively different signals identify another failing of the standard regression approach.

The axiomatic approach

The axiomatic methodology is of particular interest because it adds an additional tool to neuroscientific methodologies. Where we can falsify all the axioms, as we did for the anterior insula, we can reject the entire class of RPE models and look instead for alternate

hypotheses that might account for data. Where all three axioms are satisfied, additional axioms can be specified to refine our model. One direction for future research would be to establish whether the quantity of dopamine released in these areas, which can be measured with electrochemical methods (Phillips et al., 2003; Day et al., 2007), satisfies the axiomatic model. Such data would test the linkage between the RPE representation and dopamine most directly. Future research could also further investigate the anterior insula signal we identified by axiomatizing one of the many salience hypotheses and designing an experiment specifically to test the conditions of necessity and sufficiency for that specific axiomatic model.

Conclusion

This study introduces axiomatic modeling to neuroscience and shows the value of that approach. We formally tested the RPE hypothesis, showing both that signals from dopamine target areas satisfy the axioms of a RPE representation and that the signal from the anterior insula falsifies the axioms and cannot possibly encode a RPE signal under the conditions we examined. In contrast, the standard regression approach that dominates fMRI today relies on highly parameterized models with specific assumptions about reward, beliefs, and learning when it examines a theory like the RPE hypothesis. Such an analysis yields a correlation coefficient but no direct test of the actual hypothesis under scrutiny. The axiomatic approach provides a powerful alternative in the Popperian tradition of testing a hypothesis by attempting to falsify it. By breaking hypotheses down into their basic assumptions, not only can entire classes of models be tested, but these assumptions identify the possible ways in which the model can be proven false and

suggest how these assumptions can be tested experimentally. This approach also points the way for further model development. When the data falsifies a specific axiom, new theoretical approaches are suggested. This is not the case when low correlations are observed in traditional region-based analyses. In this sense, the axiomatic approach offers novel benefits that complement existing approaches to the analysis of brain function.

References

Abler B, Walter H, Erk S, Kammerer H, Spitzer M (2006) Prediction error as a linear function of reward probability is coded in human nucleus accumbens. *Neuroimage* 31:790-795.

Bayer HM, Glimcher PW (2005) Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* 47:129-141.

Behrens TEJ, Hunt LT, Woolrich MW, Rushworth MFS (2008) Associative learning of social value. *Nature* 456:245-250.

Belova MA, Paton JJ, Morrison SE, Salzman CD (2007) Expectation modulates neural responses to pleasant and aversive stimuli in primate amygdala. *Neuron* 55:970-984.

Berridge KC, Robinson TE (1998) What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? *Brain Res Rev* 28:309-369.

Brischoux F, Chakraborty S, Brierley DI, Ungless MA (2009) Phasic excitation of dopamine neurons in ventral VTA by noxious stimuli. *Proc Natl Acad Sci USA* 106:4894-4899.

de Bruijn ERA, de Lange FP, von Cramon DY, Ullsperger M (2009) When errors are rewarding. *J Neurosci* 29:12183-12186.

Caplin A, Dean M (2008a) Dopamine, reward prediction error, and economics. *Q J Econ* 123:663-701.

Caplin A, Dean M (2008b) Axiomatic methods, dopamine and reward prediction error. *Curr Opin Neurobiol* 18:197-202.

Caplin A, Dean M, Glimcher PW, Rutledge RB (in press) Measuring beliefs and rewards: a neuroeconomic approach. *Q J Econ*.

Caviness VSJ, Meyer J, Makris N, Kennedy DN (1996) MRI-based topographic parcellation of the human neocortex: an anatomically specified method with estimate of reliability. *J Cogn Neurosci* 8:566–587.

D'Ardenne K, McClure SM, Nystrom LE, Cohen JD (2008) BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science* 319:1264-1267.

Day JJ, Roitman MF, Wightman RM, Carelli RM (2007) Associate learning mediates dynamic shifts in dopamine signaling in the nucleus accumbens. *Nat Neurosci* 10:1020-1028.

Friston KJ, Josephs O, Rees G, Turner R (1998) Nonlinear event-related responses in fMRI. *Magn Reson Med* 39:41-52.

Grinband J, Hirsch J, Ferrera VP (2006) A neural representation of categorization uncertainty in the human brain. *Neuron* 49:1-7.

Hare TA, O'Doherty J, Camerer CF, Schultz W, Rangel A (2008) Dissociating the role of orbitofrontal cortex and striatum in the computation of goal values and prediction errors. *J Neurosci* 28:5623-5630.

Hollerman JR, Schultz W (1998) Dopamine neurons report an error in the temporal prediction of reward during learning. *Nat Neurosci* 1:304-309.

Horvitz JC (2000) Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neurosci* 96:651-656.

Huettel SA, Song AW, McCarthy G (2005) Decisions under uncertainty: probabilistic context influences activation of prefrontal and parietal cortices. *J Neurosci*. 25:3304-3311.

Jensen J, Smith AJ, Willeit M, Crawley AP, Mikulis DJ, Vitcu I, Kapur S (2007) Separate brain regions code for salience vs. valence during reward prediction in humans. *Hum Brain Mapp* 28:294-302.

Joshua M, Adler A, Mitelman R, Vaadia E, Bergman H (2008) Midbrain dopaminergic neurons and striatal cholinergic interneurons encode the difference between reward and aversive events at different epochs of probabilistic classical conditioning trials. *J Neurosci* 28:11673-11684.

Li J, McClure SM, King-Casas B, Montague PR (2006) Policy adjustment in a dynamic economic game. *PloS ONE* 1:103-113.

Matsumoto M, Hikosaka O (2007) Lateral habenula as a source of negative reward signals in dopamine neurons. *Nature* 447:1111-1115.

Matsumoto M, Matsumoto K, Abe H, Tanaka K (2007) Medial prefrontal cell activity signaling prediction errors of action values. *Nat Neurosci* 10:647-656.

Matsumoto M, Hikosaka O (2009) Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature* 459:837-841.

McClure SM, Berns GS, Montague PR (2003) Temporal predication errors in a passive learning task activate human striatum. *Neuron* 38:339-346.

McCoy AN, Crowley JC, Haghighian G, Dean HL, Platt ML (2003) Saccade reward signals in posterior cingulate cortex. *Neuron* 40:1031-1040.

Nakahara H, Itoh H, Kawagoe R, Takikawa Y, Hikosaka O (2004) Dopamine neurons can represent context-dependent prediction error. *Neuron* 41:269-280.

O'Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ (2003) Temporal difference models and reward-related learning in the human brain. *Neuron* 38:329-337

O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ (2004) Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* 304:452-454.

Pessiglione M, Seymour B, Flandin G, Dolan RJ, Frith CD (2006) Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature* 442:1042-1045.

Phillips PEM, Stuber GD, Heien MLAV, Wightman RM, Carelli RM (2003) Subsecond dopamine release promotes cocaine seeking. *Nature* 422:614-618.

Popper K (1959) *The logic of scientific discovery*. New York: Basic Books.

Preuschoff K, Bossaerts P, Quartz SR (2006) Neural differentiation of expected reward and risk in human subcortical structures. *Neuron* 51:381-390.

Preuschoff K, Quartz SR, Bossaerts P (2008) Human insula activation reflects risk prediction errors as well as risk. *J Neurosci* 28:2745-2752.

Rademacher J, Galaburda AM, Kennedy DN, Filipek PA, Caviness VSJ (1992) Human cerebral cortex: localization, parcellation, and morphometry with magnetic resonance imaging. *J Cogn Neurosci* 4:352-374.

Redgrave P, Prescott TJ, Gurney K (1999) Is the short-latency dopamine response too short to signal reward error? *Trends Neurosci* 22:146-151.

Rescorla RA, Wagner AR (1972) A theory of pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In: *Classical conditioning II: current research and theory* (Black AH, Prokasy WF, eds), pp 64-99. New York: Appleton.

Rutledge RB, Lazzaro SC, Lau B, Myers CE, Gluck MA, Glimcher PW. Dopaminergic drugs modulate learning rates and perseveration in Parkinson's patients in a dynamic foraging task. *J Neurosci* 29:15104-15114 (2009).

Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. *Science* 275:1593-1599.

Seeley WW, Menon V, Schatzberg AF, Keller J, Glover GH, Kenna H, Reiss AL, Greicius MD (2007) Dissociable intrinsic connectivity networks for salience processing and executive control. *J Neurosci* 27:2349-2356.

Seymour B, O'Doherty JP, Dayan P, Koltzenburg M, Jones AK, Dolan RJ, Friston KJ, Frackowiak RS (2004) Temporal difference models describe higher-order learning in humans. *Nature* 429:664-667.

Sutton RS, Barto AG (1990) Time-derivative models of Pavlovian reinforcement. In: *Learning and computational neuroscience: foundations of adaptive networks* (Gabriel M, Moore J, eds), pp 497-537. Cambridge, MA: MIT.

Ullsperger M, von Cramon DY (2003) Error monitoring using external feedback: specific roles of the habenular complex, the reward system, and the cingulate motor area revealed by functional magnetic resonance imaging. *J Neurosci* 23:4308-4314.

Vazquez AL, Noll DC (1998) Nonlinear aspects of the BOLD response in functional MRI. *Neuroimage* 7:108-118.

Voon V, Pessiglione M, Brezing C, Gallea C, Fernandez HH, Dolan RJ, Hallett M (2010) Mechanisms underlying dopamine-mediated reward bias in compulsive behaviors. *Neuron* 65:135-142.

Wald A (1943) Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Trans Am Math Soc* 54:426-482.

Wittman BC, Daw ND, Seymour B, Dolan RJ (2008) Striatal activity underlies novelty-based choice in humans. *Neuron* 58:967-973.

Yacubian J, Gläscher J, Schroeder K, Sommer T, Braus DF, Büchel C (2006) Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. *J Neurosci* 26:9530-9537.

Zaghloul KA, Blanco JA, Weidemann CT, McGill K, Jaggi JL, Baltuch GH, Kahana MJ (2009) Human substantia nigra neurons encode unexpected financial rewards. *Science* 323:1496-1499.

Zink CF, Pagnoni G, Martin-Skurski ME, Mukeshwar D, Berns GS (2003) Human striatal response to salient nonrewarding stimuli. *J Neurosci* 23:8092-8097.

Zink CF, Pagnoni G, Martin-Skurski ME, Chappelow JC, Berns GS (2004) Human striatal responses to monetary reward depend on saliency. *Neuron* 42:509-517.

Figure Legends and Tables

Figure 1. Experimental task and group reward prediction error (RPE) analysis. **A**, Experimental task design with timing indicated. On each trial, subjects were presented with two options, lotteries with the probability of each prize indicated by the area of the prize's slice. After 5 s, the fixation cross was extinguished and the subject had 1.25 s to indicate their decision by pressing a button. After a delay period, the prize was revealed by a change in the color of the associated slice, here winning \$5 from a lottery with a 50% chance of doing so. **B**, Areas in which neural activity was correlated with predicted RPE in a random-effects group analysis. At a threshold of $p < 0.001$ (uncorrected), areas of correlation were found in the bilateral nucleus accumbens (coronal and axial images at $y = +5$ and $z = -4$, respectively), left putamen (coronal image), and right caudate. Predicted RPE was defined as the mathematical difference in dollars between the prize received and the lottery's expected value. The color scale indicates the t-value of the contrast testing for a significant effect of predicted RPE during the outcome period. Data are overlaid on the mean normalized image and shown in radiological convention, with the right hemisphere on the left.

Figure 2. The axiomatic RPE model. Hypothetical neural activity is shown for two prizes (winning \$5 in red and losing \$5 in blue) received from five lotteries with probabilities of winning from 0% to 100%. Only two prizes are possible so, for example, the lottery with a 50% probability of winning \$5 also has a 50% probability of losing \$5. **A**, Example of a

violation of axiom 1. **B**, Example of a violation of axiom 2. **C**, Example of a violation of axiom 3. **D**, A pattern of activity with no axiomatic violations.

Figure 3. BOLD responses in the nucleus accumbens and anterior insula. **A**, **B**, ROIs were defined in individual subjects by anatomical criteria for the nucleus accumbens (coronal image) and anterior insula (axial image). The color scale indicates the number of subjects containing a particular voxel in the individual ROI definitions. Data are overlayed on the mean normalized image and shown in radiological convention, with the right hemisphere on the left. **C**, **D**, Data were averaged across all voxels in the individual anatomical ROIs and replotted as trial averages. Trial averages are color-coded by predicted RPE for each of the eight trial types. The outcome period (TR 22-24) is indicated. The window (TR 26-30) for which the axioms were tested is shown in gray. The largest standard error is shown on the right. Anatomical ROIs and trial averages for additional areas are shown in supplemental Figures 1 and 3.

Figure 4. Testing the axiomatic RPE model. **A**, **B**, Parameter estimates and 95% confidence intervals are plotted for each trial type for the two prizes (winning \$5 in red and losing \$5 in blue) against the probability of winning \$5. The data from the nucleus accumbens satisfies all three axioms at $p < 0.05$. The data from the anterior insula falsifies all three axioms at $p < 0.05$. Test results are shown in Table 1. Results for additional areas are shown in supplemental Figure 2.

Figure 5. Assessing the robustness of axiomatic RPE model analyses. **A-H**, Heatmaps show results of the axiomatic analysis for a variety of baselines and starting times for the 5-TR analysis window. Testing the axiomatic model across areas requires 10 statistical tests. The first TR of the baseline is indicated for each 2-TR baseline. The color scale indicates the number of tests with the predicted result for a RPE signal at $p < 0.05$. The baseline and analysis windows used for the analyses in Figure 4 and Table 1 is indicated by rectangles. All ROIs are defined by anatomical criteria in individuals. The neural activity in the nucleus accumbens, caudate, putamen, medial prefrontal cortex, amygdala, and posterior cingulate cortex has the predicted result for the majority of tests for a variety of baseline and analysis windows. The neural activity in the anterior insula and thalamus does not have the predicted result for a RPE signal regardless of the choice of baseline and analysis window. Nb, no baseline. Dopaminergic midbrain and habenula results are shown in supplemental Figure 4.

Figure 6. BOLD responses to positive and negative outcomes. **A-H**, BOLD responses for positive (red) and negative (blue) outcomes are plotted against model fits with the canonical two-gamma hemodynamic response function. Results are for the three two-prize lotteries. Error bars reflect \pm SEM across subjects. Dotted lines represent best fits for a regression model with regressors modeled for options, choice, and outcome onset, convolved with the canonical two-gamma hemodynamic impulse response function.

Figure 7. Testing the RPE absolute value model. **A, B**, Parameter estimates and 95% confidence intervals are plotted for the two prizes (winning \$5 in red and losing \$5 in

blue) against the probability of receiving that prize for the anatomical ROIs shown in Figures 3A and 3B. The baseline is TR 22-23 and the analysis window TR 24-28. The neural activity in the anterior insula is a largely decreasing function of prize probability, consistent with encoding the absolute value of the RPE signal, a quantity related to some notions of salience. The neural activity in the nucleus accumbens does not appear to be a decreasing function of prize probability.

Figure 8. Assessing the robustness of RPE absolute value model analyses. *A-H*, Heatmaps show results of the analysis for a variety of baselines and starting times for the 5-TR analysis window. Testing the RPE absolute value model requires 28 statistical tests. The first TR of the baseline is indicated for each 2-TR baseline. The color scale indicates the number of tests with the predicted result for a RPE signal at $p < 0.05$. The baseline and analysis window used for the analyses in Figure 7 and Table 2 is indicated by rectangles. All ROIs are defined by anatomical criteria in individuals. The neural activity in the anterior insula has the predicted result for a RPE absolute value signal for most tests for a variety of baseline and analysis windows. The nucleus accumbens and amygdala do not have the predicted result for a RPE absolute value signal regardless of the choice of baseline and analysis window. Nb, no baseline. Dopaminergic midbrain and habenula results are shown in supplemental Figure 5.

Axiom		NAcc	AI	Caud	Put	MPFC	Am	PCC	Thal
1.1	+	+	=	+	+	+	+	+	+
1.2	+	+	=	+	+	+	+	=	=
1.3	+	+	-	+	+	+	+	+	=
2.1	-	-	=	-	-	-	=	=	-
2.2	-	-	+	-	-	+	-	+	=
2.3	-	-	-	-	-	-	-	-	-
2.4	-	-	-	-	=	-	=	-	-
2.5	-	-	-	-	-	-	-	-	-
2.6	-	-	=	-	-	-	-	-	-
3	=	=	-	=	-	=	=	=	-

Table 1. Axiomatic RPE model statistical tests. Testing the three axioms of the axiomatic

RPE model on our data requires 10 statistical tests. Wald tests of linear restriction were performed on parameter estimates computed with a baseline of TR 9-10 and an analysis window of TR 26-30 (parameter estimates for the nucleus accumbens and anterior insula are shown in Figure 4) with the sign of all significant tests indicated ($p < 0.05$). We predicted that RPE signals would satisfy the axioms in the way indicated by the signs in the left-most ‘predicted sign’ column. At $p < 0.05$, the nucleus accumbens and caudate each satisfy all three axioms. The anterior insula and thalamus falsify all three axioms. The amygdala and medial prefrontal cortex each satisfy two axioms and the putamen and posterior cingulate cortex each satisfy one axiom. Axiomatic statistical test 1.1, {+\$5, 25% probability of winning \$5} - {- \$5, 25%}; 1.2, {+\$5, 50%} - {- \$5, 50%}; 1.3, {+\$5, 75%} - {- \$5, 75%}; 2.1, {+\$5, 50%} - {+\$5, 25%}; 2.2, {- \$5, 50%} - {- \$5, 25%}; 2.3, {+\$5, 75%} - {+\$5, 50%}; 2.4, {- \$5, 75%} - {- \$5, 50%}; 2.5, {+\$5, 75%} - {+\$5, 25%}; 2.6, {- \$5, 75%} - {- \$5, 25%}; 3, {+\$5, 100%} - {- \$5, 0%}. NAcc, nucleus accumbens; AI, anterior insula; Caud, caudate; Put, putamen; MPFC, medial prefrontal cortex; Am, amygdala; PCC, posterior cingulate cortex; Thal, thalamus.

Condition		NAcc	AI	Caud	Put	MPFC	Am	PCC	Thal
1.1	+	=	+	+	+	+	=	+	+
1.2	+	+	+	+	=	+	=	+	+
1.3	+	+	+	+	+	+	+	+	+
1.4	+	+	+	+	+	+	+	+	+
1.5	+	+	+	+	+	+	+	+	+
1.6	+	+	+	+	+	+	+	+	+
1.7	+	-	+	=	-	-	-	-	=
1.8	+	-	+	+	-	-	-	-	=
1.9	+	+	+	+	+	+	-	-	+
1.10	+	=	+	+	+	=	=	-	=
1.11	+	=	+	+	=	=	-	-	+
1.12	+	=	+	+	+	+	=	=	+
1.13	+	+	+	+	-	=	=	=	=
1.14	+	+	+	+	+	+	+	+	=
1.15	+	+	+	+	+	+	+	=	+
1.16	+	+	+	+	+	+	+	+	+
1.17	+	-	+	-	-	-	-	=	=
1.18	+	+	+	+	+	+	-	+	=
1.19	+	-	+	-	-	=	-	=	+
1.20	+	=	+	+	+	+	=	+	+
1.21	+	+	+	+	+	+	+	+	+
1.22	+	+	+	+	+	+	+	+	+
1.23	+	+	+	+	+	+	=	+	+
1.24	+	+	+	+	+	+	+	+	+
1.25	=	+	=	+	+	+	+	+	+
1.26	=	+	=	+	+	+	+	=	=
1.27	=	+	=	+	+	+	+	=	=
1.28	=	-	-	-	=	-	=	+	=

Table 2. RPE absolute value model statistical tests. Testing the RPE absolute value model requires 28 tests. Wald tests of linear restriction were performed on parameter estimates with an analysis window of TR 24-28 and a baseline of TR 22-23 (parameter estimates for the nucleus accumbens and anterior insula are in Figure 7) with the sign of all significant tests indicated ($p < 0.05$). The left-most column indicates the predicted signs for a RPE absolute value signal. Tests 1.1-1.24 compare outcomes to other outcomes with lower probability. Tests 1.25-1.28 compare outcomes to other outcomes with the same prize probability. Tests are listed in Supplemental Data. ROI abbreviations are as in Table 1.

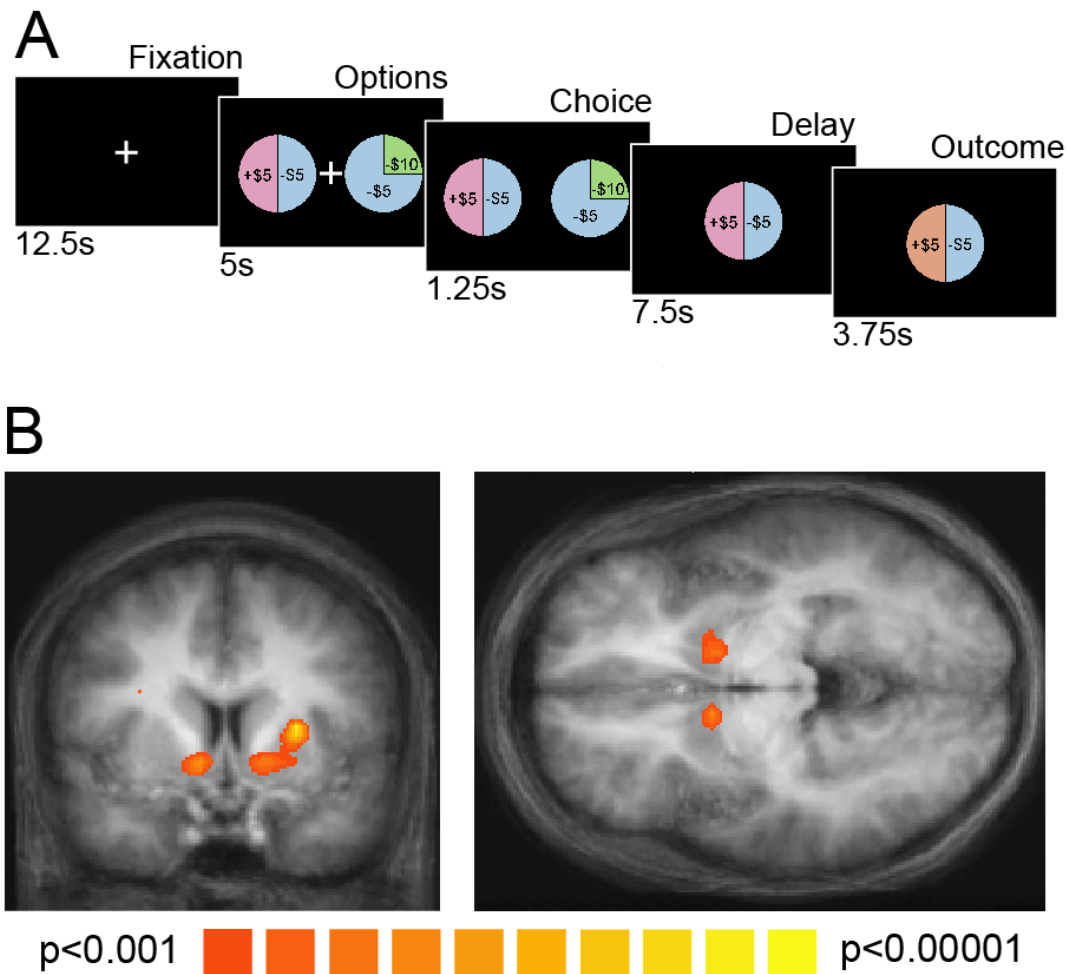


Figure 1. Experimental task and group reward prediction error (RPE) analysis. **A**, Experimental task design with timing indicated. On each trial, subjects were presented with two options, lotteries with the probability of each prize indicated by the area of the prize's slice. After 5 s, the fixation cross was extinguished and the subject had 1.25 s to indicate their decision by pressing a button. After a delay period, the prize was revealed by a change in the color of the associated slice, here winning \$5 from a lottery with a 50% chance of doing so. **B**, Areas in which neural activity was correlated with predicted RPE in a random-effects group analysis. At a threshold of $p < 0.001$ (uncorrected), areas of correlation were found in the bilateral nucleus accumbens (coronal and axial images at $y = +5$ and $z = -4$, respectively), left putamen (coronal image), and right caudate. Predicted RPE was defined as the mathematical difference in dollars between the prize received and the lottery's expected value. The color scale indicates the t-value of the contrast testing for a significant effect of predicted RPE during the outcome period. Data are overlaid on the mean normalized image and shown in radiological convention, with the right hemisphere on the left.

Axiomatic RPE model

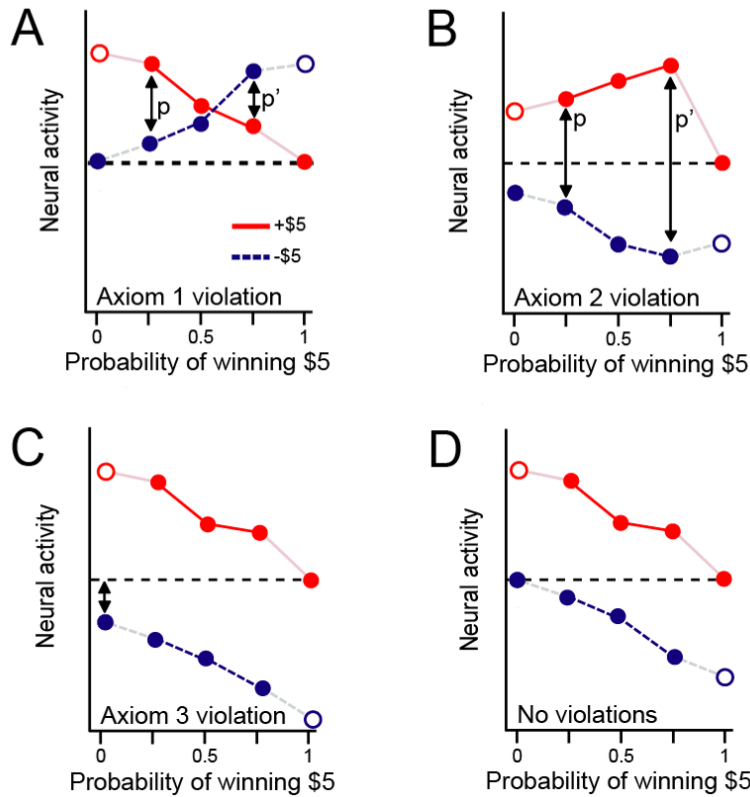


Figure 2. The axiomatic RPE model. Hypothetical neural activity is shown for two prizes (winning \$5 in red and losing \$5 in blue) received from five lotteries with probabilities of winning from 0% to 100%. Only two prizes are possible so, for example, the lottery with a 50% probability of winning \$5 also has a 50% probability of losing \$5. **A**, Example of a violation of axiom 1. **B**, Example of a violation of axiom 2. **C**, Example of a violation of axiom 3. **D**, A pattern of activity with no axiomatic violations.

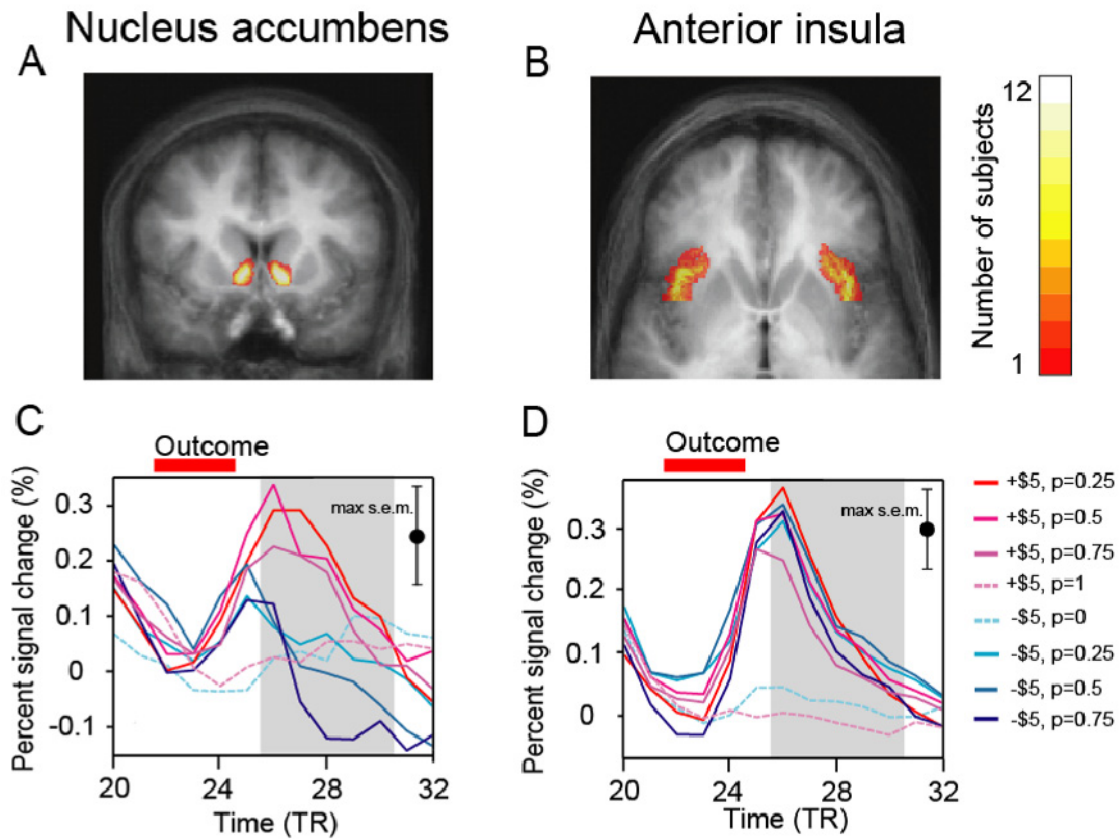


Figure 3. BOLD responses in the nucleus accumbens and anterior insula. **A**, **B**, ROIs were defined in individual subjects by anatomical criteria for the nucleus accumbens (coronal image) and anterior insula (axial image). The color scale indicates the number of subjects containing a particular voxel in the individual ROI definitions. Data are overlaid on the mean normalized image and shown in radiological convention, with the right hemisphere on the left. **C**, **D**, Data were averaged across all voxels in the individual anatomical ROIs and replotted as trial averages. Trial averages are color-coded by predicted RPE for each of the eight trial types. The outcome period (TR 22-24) is indicated. The window (TR 26-30) for which the axioms were tested is shown in gray. The largest standard error is shown on the right. Anatomical ROIs and trial averages for additional areas are shown in supplemental Figures 1 and 3.

Testing the axiomatic RPE model

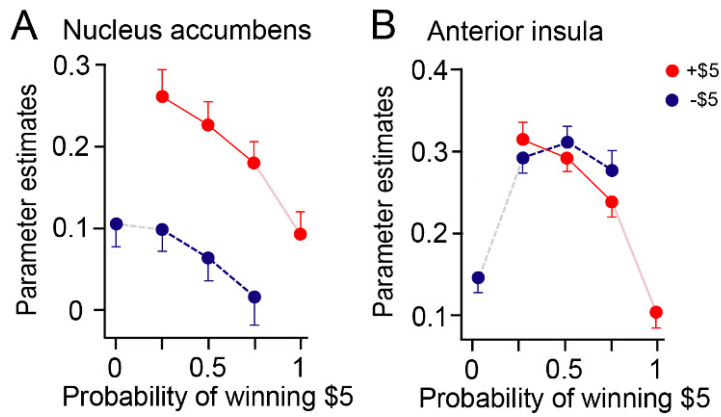


Figure 4. Testing the axiomatic RPE model. *A, B*, Parameter estimates and 95% confidence intervals are plotted for each trial type for the two prizes (winning \$5 in red and losing \$5 in blue) against the probability of winning \$5. The data from the nucleus accumbens satisfies all three axioms at $p < 0.05$. The data from the anterior insula falsifies all three axioms at $p < 0.05$. Test results are shown in Table 1. Results for additional areas are shown in supplemental Figure 2.

Testing the axiomatic RPE model

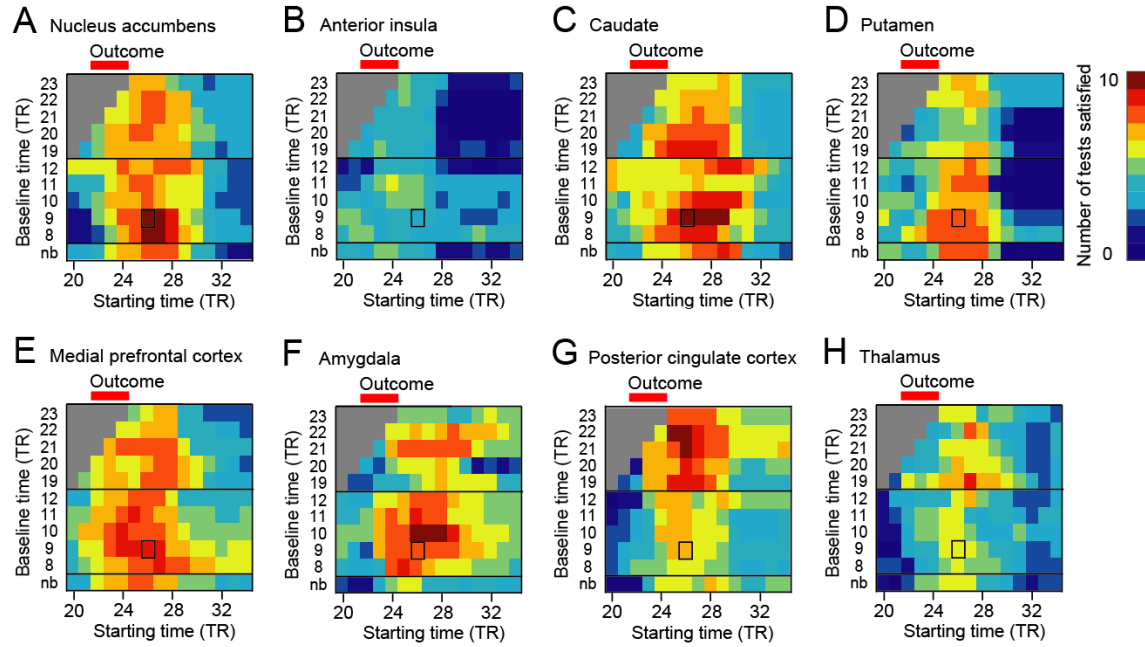


Figure 5. Assessing the robustness of axiomatic RPE model analyses. **A-H**, Heatmaps show results of the axiomatic analysis for a variety of baselines and starting times for the 5-TR analysis window. Testing the axiomatic model across areas requires 10 statistical tests. The first TR of the baseline is indicated for each 2-TR baseline. The color scale indicates the number of tests with the predicted result for a RPE signal at $p < 0.05$. The baseline and analysis windows used for the analyses in Figure 4 and Table 1 is indicated by rectangles. All ROIs are defined by anatomical criteria in individuals. The neural activity in the nucleus accumbens, caudate, putamen, medial prefrontal cortex, amygdala, and posterior cingulate cortex has the predicted result for the majority of tests for a variety of baseline and analysis windows. The neural activity in the anterior insula and thalamus does not have the predicted result for a RPE signal regardless of the choice of baseline and analysis window. Nb, no baseline. Dopaminergic midbrain and habenula results are shown in supplemental Figure 4.

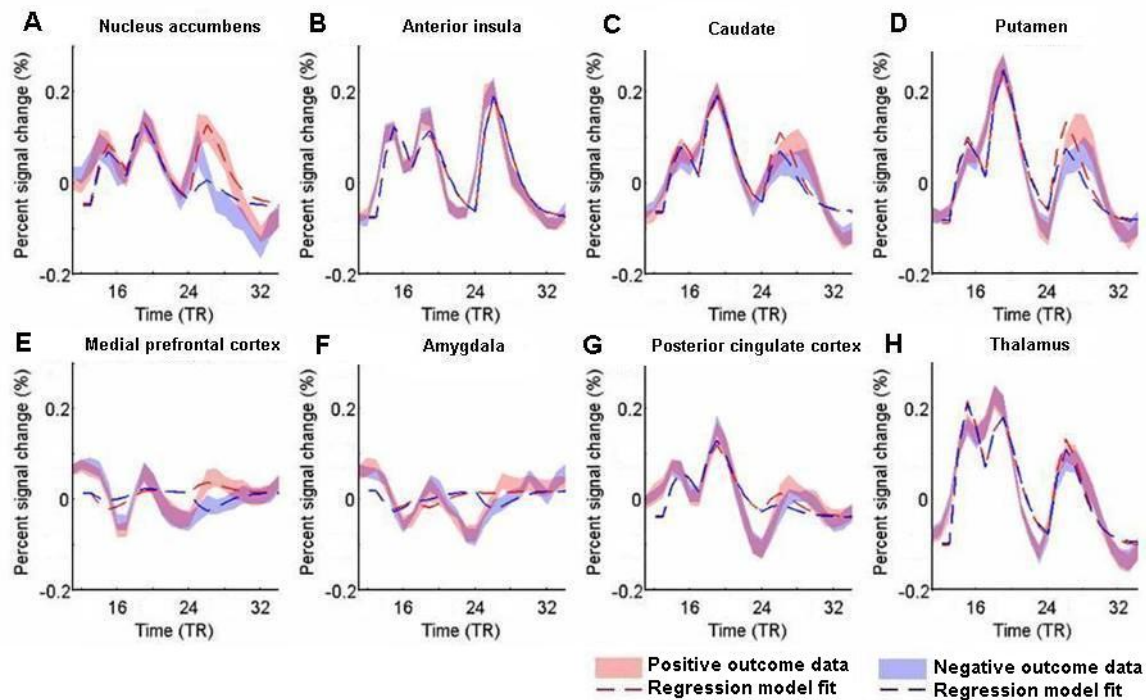


Figure 6. BOLD responses to positive and negative outcomes. *A-H*, BOLD responses for positive (red) and negative (blue) outcomes are plotted against model fits with the canonical two-gamma hemodynamic response function. Results are for the three two-prize lotteries. Error bars reflect \pm SEM across subjects. Dotted lines represent best fits for a regression model with regressors modeled for options, choice, and outcome onset, convolved with the canonical two-gamma hemodynamic impulse response function.

Testing the RPE absolute value model

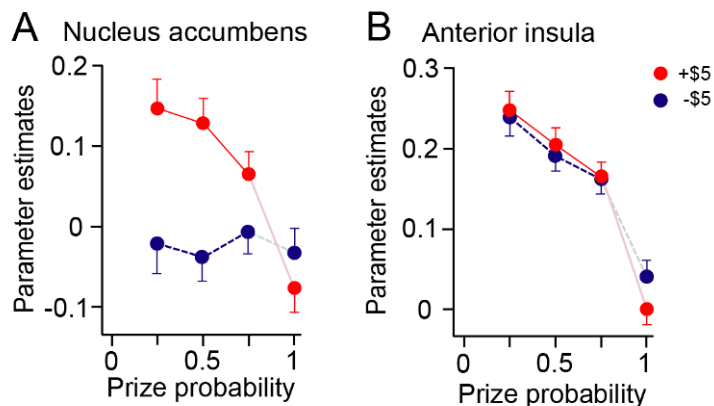


Figure 7. Testing the RPE absolute value model. **A, B,** Parameter estimates and 95% confidence intervals are plotted for the two prizes (winning \$5 in red and losing \$5 in blue) against the probability of receiving that prize for the anatomical ROIs shown in Figures 3A and 3B. The baseline is TR 22-23 and the analysis window TR 24-28. The neural activity in the anterior insula is a largely decreasing function of prize probability, consistent with encoding the absolute value of the RPE signal, a quantity related to some notions of salience. The neural activity in the nucleus accumbens does not appear to be a decreasing function of prize probability.

Testing the RPE absolute value model

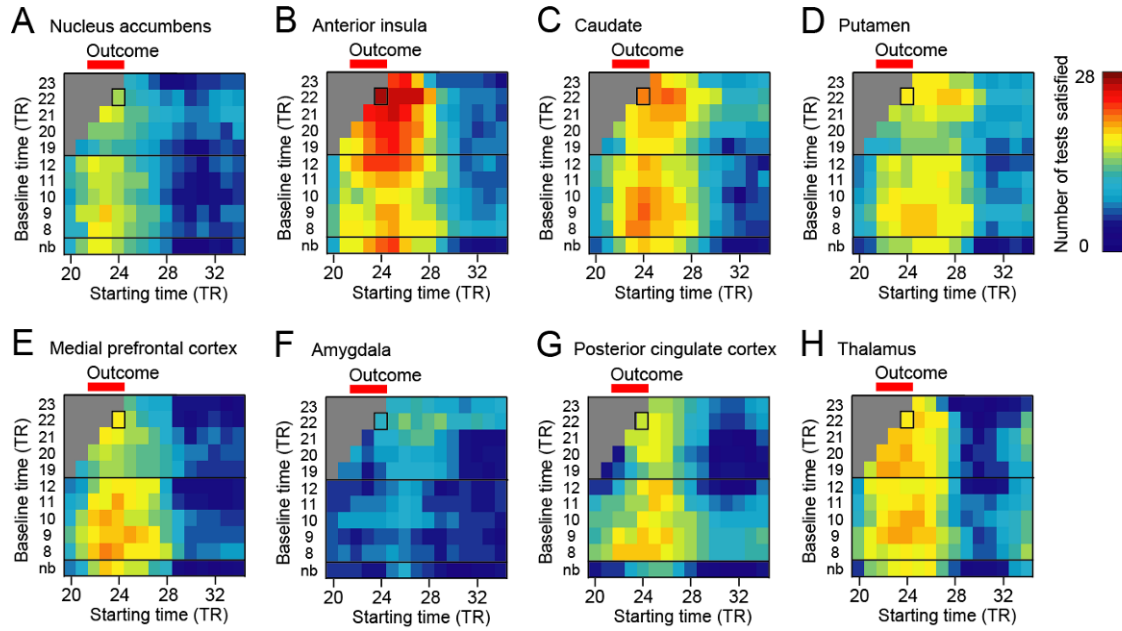


Figure 8. Assessing the robustness of RPE absolute value model analyses. *A-H*, Heatmaps show results of the analysis for a variety of baselines and starting times for the 5-TR analysis window. Testing the RPE absolute value model requires 28 statistical tests. The first TR of the baseline is indicated for each 2-TR baseline. The color scale indicates the number of tests with the predicted result for a RPE signal at $p < 0.05$. The baseline and analysis window used for the analyses in Figure 7 and Table 2 is indicated by rectangles. All ROIs are defined by anatomical criteria in individuals. The neural activity in the anterior insula has the predicted result for a RPE absolute value signal for most tests for a variety of baseline and analysis windows. The nucleus accumbens and amygdala do not have the predicted result for a RPE absolute value signal regardless of the choice of baseline and analysis window. Nb, no baseline. Dopaminergic midbrain and habenula results are shown in supplemental Figure 5.

SUPPLEMENTAL DATA

Falsifying the Reward Prediction Error Hypothesis with an Axiomatic Model

Robb B. Rutledge, Mark Dean, Andrew Caplin & Paul W. Glimcher

Absolute value of RPE model statistical tests

Condition 1.1, {+\$5, 25% probability of winning \$5} - {+\$5, 50%}; 1.2, {+\$5, 25%} - {+\$5, 75%}; 1.3, {+\$5, 25%} - {+\$5, 100%}; 1.4, {+\$5, 25%} - {- \$5, 50%}; 1.5, {+\$5, 25%} - {- \$5, 25%}; 1.6, {+\$5, 25%} - {- \$5, 0%}; 1.7, {- \$5, 75%} - {+\$5, 50%}; 1.8, {- \$5, 75%} - {+\$5, 75%}; 1.9, {- \$5, 75%} - {+\$5, 100%}; 1.10, {- \$5, 75%} - {- \$5, 50%}; 1.11, {- \$5, 75%} - {- \$5, 25%}; 1.12, {- \$5, 75%} - {- \$5, 0%}; 1.13, {+\$5, 50%} - {+\$5, 75%}; 1.14, {+\$5, 50%} - {+\$5, 100%}; 1.15, {+\$5, 50%} - {- \$5, 25%}; 1.16, {+\$5, 50%} - {- \$5, 0%}; 1.17, {- \$5, 50%} - {+\$5, 75%}; 1.18, {- \$5, 50%} - {+\$5, 100%}; 1.19, {- \$5, 50%} - {- \$5, 25%}; 1.20, {- \$5, 50%} - {- \$5, 0%}; 1.21, {+\$5, 75%} - {+\$5, 100%}; 1.22, {+\$5, 75%} - {- \$5, 0%}; 1.23, {- \$5, 25%} - {+\$5, 100%}; 1.24, {- \$5, 25%} - {- \$5, 0%}; 1.25, {+\$5, 25%} - {- \$5, 75%}; 1.26, {+\$5, 50%} - {- \$5, 50%}; 1.27, {+\$5, 75%} - {- \$5, 25%}; 1.28, {+\$5, 100%} - {- \$5, 0%}.

Anatomical ROI definitions

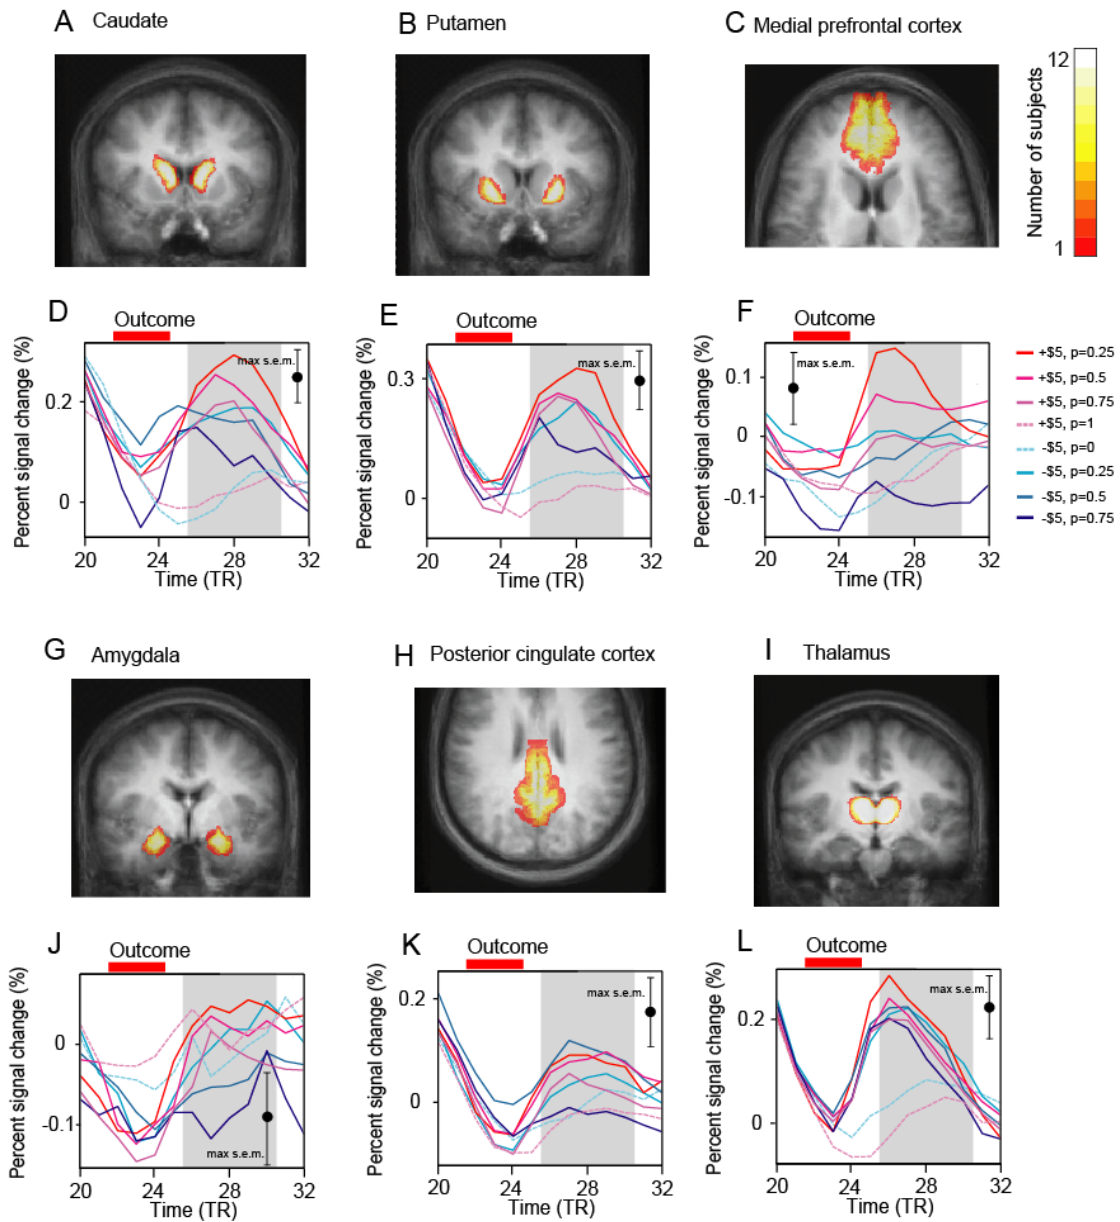
We defined each of the 11 anatomical ROIs on the T1-weighted MP-RAGE anatomical images for each subject in Talairach space. The nucleus accumbens, caudate, putamen, amygdala, insula, and thalamus were defined according to the criteria established by the Center for Morphometric Analysis (Caviness et al., 1996; Rademacher et al., 1992; instruction manuals at <http://www.cma.mgh.harvard.edu/manuals/>). We defined the anterior insula by the portion of the insula at least 5 mm anterior to the anterior commissure. We defined the posterior cingulate cortex and medial prefrontal cortex by selecting all gray matter within Brodmann areas 23 and 31 and Brodmann areas 10 (medial portion) and 32 (ventral portion), respectively. We identified the substantia nigra using the cerebral peduncles and the red nucleus as boundaries. D'Ardenne and colleagues (2008) identified the ventral tegmental area as a region approximately 60 mm³ in volume anterior to the red nucleus and bounded laterally by the substantia nigra, we defined the ventral tegmental area by those boundaries as a region exactly 56 mm³ in volume. We defined the habenula bilaterally as a region of 60-80 mm³ superior and anterior to the habenular commissure in the dorsal medial thalamus.

REFERENCES

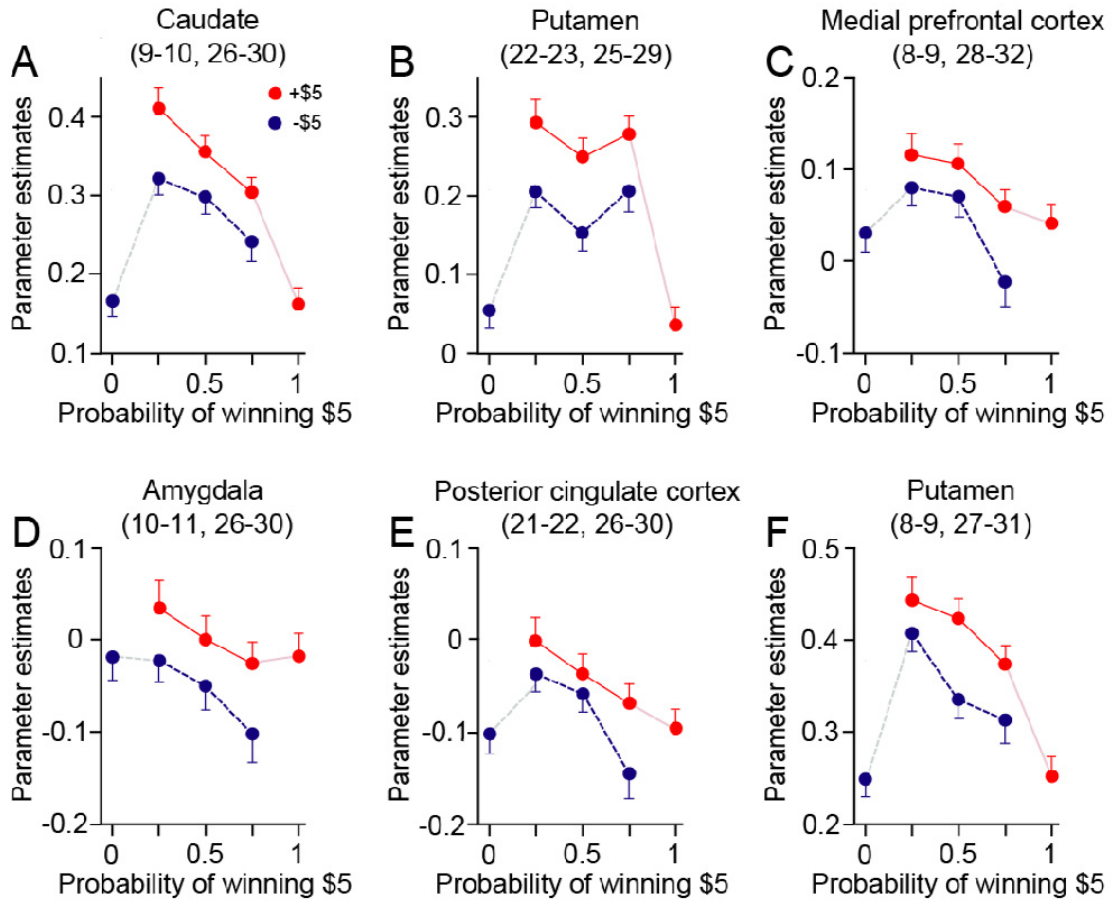
Caviness VSJ, Meyer J, Makris N, Kennedy DN (1996) MRI-based topographic parcellation of the human neocortex: an anatomically specified method with estimate of reliability. *J Cogn Neurosci* 8:566–587.

D'Ardenne K, McClure SM, Nystrom LE, Cohen JD (2008) BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science* 319:1264-1267.

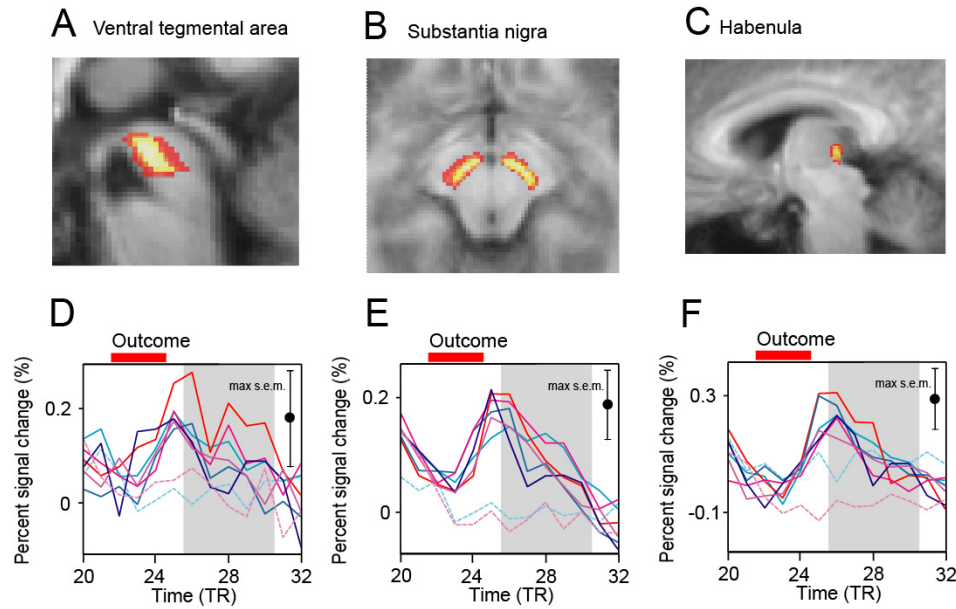
Rademacher J, Galaburda AM, Kennedy DN, Filipek PA, Caviness VSJ (1992) Human cerebral cortex: localization, parcellation, and morphometry with magnetic resonance imaging. *J Cogn Neurosci* 4:352-374.



Supplemental Figure 1. Anatomical ROI definitions and trial averages for six brain regions. The color scale indicates the number of subjects containing a particular voxel in the individual ROI definitions. Data are overlaid on the mean normalized image and shown in radiological convention, with the right hemisphere on the left. Average BOLD time courses for each ROI are shown for each of the eight trial types. Trial averages are color-coded by predicted RPE for each of the eight trial types. The outcome period (TR 22-24) is indicated. The analysis window (TR 26-30) used in Table 1 is shown in gray. The largest SEM for each ROI is shown on the right.

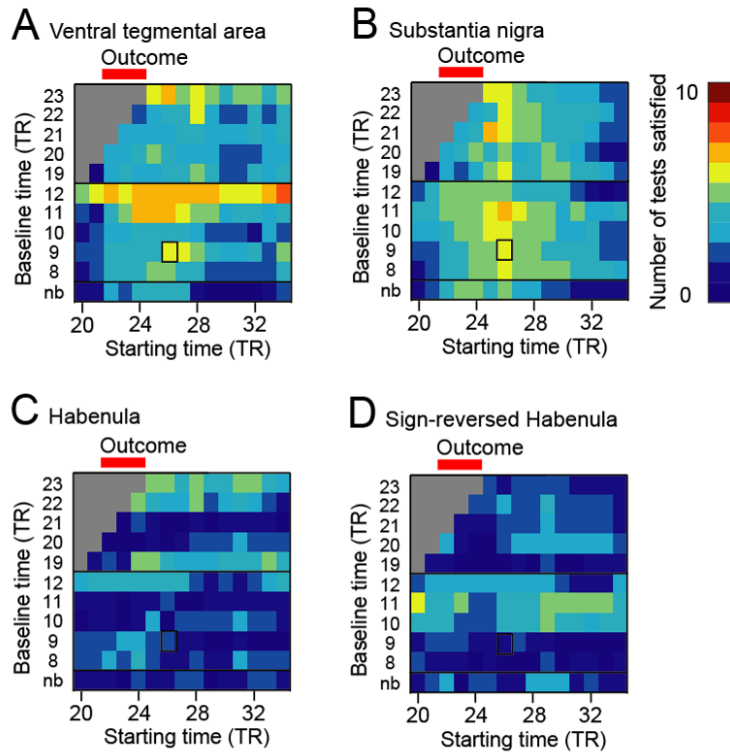


Supplemental Figure 2. Testing the axiomatic model. Parameter estimates and 95% confidence intervals are plotted for each trial type for the two prizes (winning \$5 in red and losing \$5 in blue) against the probability of winning \$5. Baseline TRs and analysis window TRs are indicated in parentheses. **A-E**, Signals for which the three axioms are satisfied at $p < 0.05$. **F**, In the putamen, the axioms are satisfied with the predicted signs at $p < 0.10$.



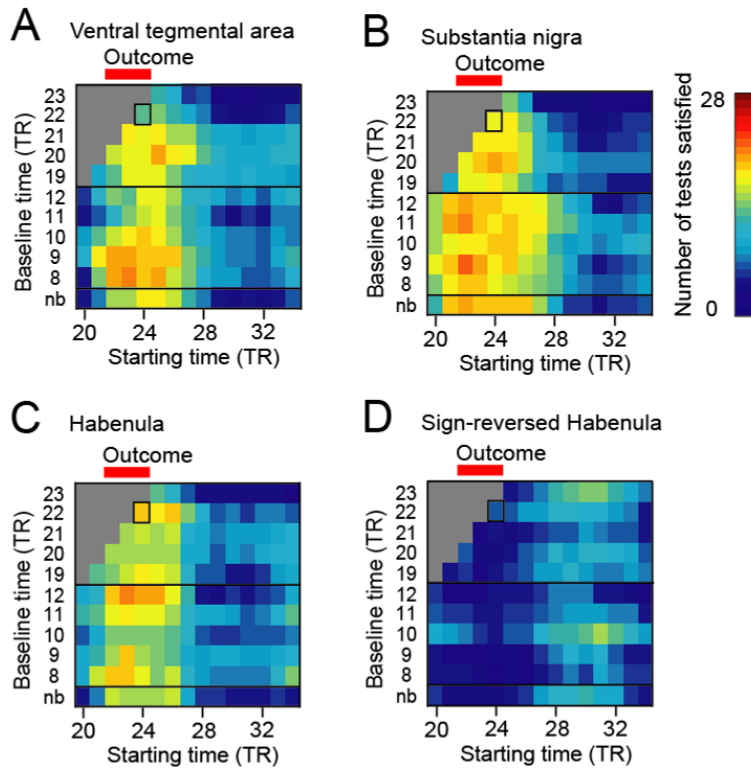
Supplemental Figure 3. Anatomical ROI definitions for three additional brain regions for the dopaminergic midbrain and habenula. ROIs are defined in individual subjects in sagittal (**A**, **C**) and axial (**B**) images. The color scale indicates the number of subjects containing a particular voxel in the individual ROI definitions. Data are overlaid on the mean normalized image and shown in radiological convention, with the right hemisphere on the left. Average BOLD time courses for each ROI are shown for each of the eight trial types. Trial averages are color-coded by predicted RPE for each of the eight trial types. The outcome period (TR 22-24) is indicated. The analysis window (TR 26-30) used in Table 1 is shown in gray. The largest SEM for each ROI is shown on the right.

Testing the axiomatic RPE model



Supplemental Figure 4. Tests of the axiomatic RPE model on the dopaminergic midbrain and habenula. Heatmaps show results of the axiomatic analysis for a variety of baselines and starting times for the 5-TR analysis window. The first TR of the baseline is indicated for each 2-TR baseline. The color scale indicates the number of statistical tests with the predicted result for a RPE at $p < 0.05$. The baseline and analysis windows used for the analyses in Table 1 are indicated. The heatmap for the sign-reversed habenula is shown in **D**.

Testing the RPE absolute value model



Supplemental Figure 5. Tests of the RPE absolute value model on the dopaminergic midbrain and habenula. Heatmaps show results of the axiomatic analysis for a variety of baselines and starting times for the 5-TR analysis window. The first TR of the baseline is indicated for each 2-TR baseline. The color scale indicates the number of statistical tests with the predicted result for a RPE absolute value signal at $p < 0.05$. The baseline and analysis windows used for the analyses in Table 2 are indicated. The heatmap for the sign-reversed habenula is shown in **D**.

ROI	Number of voxels	% in analysis
Nucleus accumbens	945 (157)	97.0 (7.7)
Anterior insula	2944 (465)	98.6 (1.6)
Caudate	4344 (729)	99.7 (0.7)
Putamen	5189 (913)	99.8 (0.5)
Medial prefrontal cortex	13357 (1183)	96.8 (1.1)
Amygdala	1337 (276)	87.9 (16.4)
Posterior cingulate cortex	20003 (2118)	77.2 (34.3)
Thalamus	10433 (1010)	98.9 (1.2)
Ventral tegmental area	56 (0)	98.6 (4.1)
Substantia nigra	353 (40)	98.8 (1.8)
Habenula	68 (4)	96.7 (8.0)

Supplemental Table 1. Anatomical ROI sizes. Mean (standard deviation) of the number of 1 x 1 x 1 mm voxels in Talairach space in individual anatomical ROI definitions. Ventral tegmental area has no variation in ROI size because it was defined to be exactly 56 voxels. Percent (standard deviation) of voxels included in further analysis is indicated for each ROI. Voxels were excluded from further analysis if the scan-level standard deviation of the signal exceeded 2%. This often occurred in voxels near the boundaries of the scanning volume, due to artifacts from the motion correction algorithm. The amygdala and posterior cingulate cortex are the ROIs most affected by these artifacts, and a significant percentage of the voxels in those ROIs were excluded from further analysis.