

Dissociating the Role of the Orbitofrontal Cortex and the Striatum in the Computation of Goal Values and Prediction Errors

Todd A. Hare,¹ John O'Doherty,¹ Colin F. Camerer,¹ Wolfram Schultz,² and Antonio Rangel¹

¹Division of Humanities and Social Sciences, California Institute of Technology, Pasadena, California 91125, and ²Department of Physiology, Development, and Neuroscience, University of Cambridge, Cambridge CB2 3DY, United Kingdom

To make sound economic decisions, the brain needs to compute several different value-related signals. These include goal values that measure the predicted reward that results from the outcome generated by each of the actions under consideration, decision values that measure the net value of taking the different actions, and prediction errors that measure deviations from individuals' previous reward expectations. We used functional magnetic resonance imaging and a novel decision-making paradigm to dissociate the neural basis of these three computations. Our results show that they are supported by different neural substrates: goal values are correlated with activity in the medial orbitofrontal cortex, decision values are correlated with activity in the central orbitofrontal cortex, and prediction errors are correlated with activity in the ventral striatum.

Key words: decision making; neuroeconomics; fMRI; goal values; decision values; prediction errors

Introduction

The brain must perform multiple value computations to make sound choices. Among them are the computation of goal values (GVs), decision values (DVs), and prediction errors (PEs). Goal values measure the predicted amount of reward associated with the outcome generated by each of the actions under consideration. Decision values measure the net value of taking the different actions; i.e., the benefits minus the costs. Prediction errors measure deviations from individuals' previous reward expectations; they are positive every time something better than expected happens, and negative when the opposite occurs (Schultz et al., 1997; Sutton and Barto, 1998). These three computations play different roles in decision making. Goal values and decision values are used to guide decisions to those actions with the largest net benefit at the time of decision making. Prediction errors are used to learn the value of states of the world and, thus, are critical for learning how to make better decisions in the future. An example of the calculation of these values is given in the Materials and Methods.

A large number of psychiatric diseases involve deficits in decision-making mechanisms (Dom et al., 2005; Paulus, 2007). An improved understanding of the neural basis of goal values, decision values, and prediction errors would advance our understanding of the impact that different types of neuropathologies and brain lesions might have on decision making. However, these

dissociable effects are predicated on the idea that separate neural systems perform different value computations in decision making.

The extent to which goal value, decision value, and prediction error computations have a different neural basis is an important open question in behavioral neuroscience and neuroeconomics. Prediction error signals have been shown to be expressed by mid-brain dopaminergic neurons that project widely to the striatum and prefrontal cortex in nonhuman primates (Schultz et al., 1997; Hollerman and Schultz, 1998; Bayer and Glimcher, 2005; Bayer et al., 2007). Recent human neuroimaging studies have separately suggested that blood oxygen level-dependent (BOLD) signal in the ventral striatum (VS) correlates with both prediction errors (McClure et al., 2003; O'Doherty et al., 2003; Abler et al., 2006; Li et al., 2006; Rodriguez et al., 2006; Tobler et al., 2006; Bray and O'Doherty, 2007; Murray et al., 2008; Rolls et al., 2008) and goal values (Aharon et al., 2001; Knutson et al., 2001, 2007; Hariri et al., 2006; Yacubian et al., 2006, 2007; Kable and Glimcher, 2007; Schaefer and Rotte, 2007). However, other studies designed to identify regions of the brain expressing goal value computations have shown activity reflecting goal values in orbitofrontal and medial prefrontal cortex, but not in the ventral striatum (Plassmann et al., 2007; Rolls et al., 2008).

Thus, the extent to which these computations have a common or separate neural substrate remains unanswered. We address this question using functional magnetic resonance imaging (fMRI) and a novel decision-making paradigm that allowed us to dissociate the neural substrates related to the goal value, decision value, and prediction error computations. The crucial aspect of this paradigm is that it decouples the parameters for the three value signals, thus permitting us to answer the question of what computation was performed in each brain region.

Received Jan. 8, 2008; accepted April 21, 2008.

This work was supported by the Moore Foundation.

Correspondence should be addressed to Antonio Rangel, Division of Humanities and Social Sciences, California Institute of Technology, HSS 228-77, Pasadena, CA 91125-7700. E-mail: rangel@hss.caltech.edu.

DOI:10.1523/JNEUROSCI.1309-08.2008

Copyright © 2008 Society for Neuroscience 0270-6474/08/285623-08\$15.00/0

Materials and Methods

Subjects

Sixteen subjects participated in the experiment (nine males; mean age, 24.1 years; age range, 19–38 years). Three additional subjects participated in the experiment but were excluded from the analysis because their low bids for the food items or their behavioral response curves indicated that they did not desire the items. Another additional subject was excluded because of difficulties in aligning his anatomy to the standard MNI space. All subjects were right handed, healthy, had normal or corrected-to-normal vision, had no history of psychiatric diagnoses, neurological or metabolic illnesses, and were not taking medications that interfere with the performance of fMRI. All subjects had no history of eating disorders and were screened for liking and at least occasionally eating the types of foods that we used. Subjects were told that the goal of the experiment was to study food preferences and gave written consent before participating. The review board of the California Institute of Technology (Pasadena, CA) approved the study.

Stimuli

Subjects made decisions on 50 different sweet and salty junk foods (e.g., chips and candy bars). The foods were presented to the subjects using high-resolution color pictures (72 dpi). The stimulus presentation and response recording was controlled by E-prime (Psychology Software Tools). The visual stimuli were presented using video goggles.

Value computations

In the current study, all value computations were expressed in US dollars. Goal values refer to the predicted reward associated with each of the food items. Before the fMRI task, the goal values for each subject and food item were measured using a Becker–DeGroot–Marschak (BDM) auction procedure that has been shown to elicit an individual's willingness to pay for a consumer good (Becker et al., 1964; Plassmann et al., 2007). Willingness to pay is the maximum amount that an individual will spend to obtain the item offered for sale. Note that an individual's goal value for a particular item does not remain constant and depends on the current state. For example, when hungry, a person will pay more for a snack than she will pay for the same snack when satiated. The decision value is equal to the net value of getting the food item. Knowing the goal value (i.e., willingness to pay) for each food item allows us to compute the decision value by subtracting the price at which the food item is offered from the goal value. For example, if a candy bar with a goal value of \$2 is offered at the cost of \$1, then the decision value would be $\$2 - \$1 = \$1$. However if the cost of the candy bar were \$3, then the decision value would be $\$2 - \$3 = -\$1$. As stated above, prediction error measures the difference between the current state of the world and the predicted value. If in the current trial a candy bar valued by the subject at \$2 is offered at a cost to the subject of \$1 and in addition the subject receives a random prize of \$2, then the total value of the current trial is the goal value (\$2) minus the cost (\$1) plus the prize (\$2) = \$3. If the predicted value for the trial, based on outcomes of previous trials was \$1, then there would be a positive prediction error of \$2.

Task

Subjects were instructed not to eat for 4 h before the experiment, which increased the value that they placed on the foods. They were also informed that they would have to remain in the lab for 30 min at the conclusion of the experiment, and that the only thing that they will be able to eat is whatever food they earned or purchased during the task. In addition to a \$60 participation fee, each subject received five \$1 bills in "spending money" to purchase food from us. Whatever money they did not spend was theirs to keep. The design of the behavioral task is a critical component of the current study. In most experimental paradigms examining reward-based decision making, goal values, decision values, and prediction, errors are highly correlated. This correlation makes it extremely difficult to distinguish areas that compute one value signal from areas that compute another, and thus might have lead to the misattribution of function. The current paradigm reduced the correlation between the goal value, decision value, and prediction error parameters by simultaneously presenting three sources of information that combined to form the overall value measure for each decision-making trial. The subject's

task was to decide whether or not to purchase a familiar snack food item. Immediately before entering the MRI scanner, subjects placed bids in a computerized BDM auction (Becker et al., 1964) for the right to eat a snack at the end of the experiment in 50 different bidding trials. The rules of the BDM auction create a situation in which the optimal strategy for the subject is to bid exactly what she is willing to pay for that food item (Becker et al., 1964; Plassmann et al., 2007). This optimal strategy was explained in detail to the subjects before the experiment began. This procedure yielded a precise estimate of the goal value for each food item on an individual subject basis. Food items were presented in random order and subjects were given four seconds to place a bid for the current food item by clicking the mouse on a continuous scale from 0 to 5 dollars. Each trial ended as soon as the subject placed a bid and the subsequent trial began immediately. Two subjects failed to place a bid for one or two items within four seconds during the prescan auction. Trials containing the missed food items were assigned as errors during fMRI analysis for those subjects.

Once inside the MRI scanner, subjects completed 300 trials of a forced-choice (yes/no) task using the same 50 food items. Figure 1 describes the time structure of the fMRI experiment. On each trial subjects were shown a screen containing a high-resolution image of a food item with the price (x) above it and a gain/loss (z) below it. Both x and z were drawn at random from the distribution $(-3:3)$, where negative numbers represented a subtraction from the subjects "spending money," and positive numbers an addition to it. The subjects responded with a button press to indicate whether or not they would pay the indicated price for the food item shown. Buttons assigned to the yes and no responses were counter-balanced across subjects. The amount represented by z was added or subtracted regardless of the subject's decision. The presence of the positive/negative price and random gain/loss amount in this paradigm minimizes the relationship between goal values, decision values, and prediction errors allowing the three signals to be dissociated. For example, on one trial a highly valued food item (goal value) might be offered in conjunction with a negative price and large loss resulting in a negative prediction error, whereas in a different trial a low-value food item might be offered along with a positive price and a gain, thus resulting in a positive prediction error. Similarly, the fact that price could be either positive or negative allowed the dissociation of goal and decision values. For example, a low-value food item offered at a very positive price would result in a high decision value, but a low goal value. On 61% of the trials, the food images, x , and z were drawn at random. On the remaining trials, these variables were manipulated in the manner described below to increase the orthogonality of the three value measures (GV, DV, and PE). For 13% of the trials, x and z were set to 0 so that only the goal value was present; on another 26% of the trials, the food image was replaced by a yellow square. On half of the yellow square trials, z was set to 0, leaving only the decision value, and on the other half x was set to 0, leaving only the prediction error signal. All trials were modeled by the parametric regressors for goal value, decision value, and prediction error. As a result of this design the average correlation between the regressors for goal value and prediction error was $r = 0.306$. Had our design included only the food items and not x or z , the average correlation between goal value and prediction error would have been $r = 0.777$. Similarly, the correlation between decision value and prediction error was $r = 0.573$ in the current design, whereas if we had not included z , the correlation would have been $r = 0.853$. Thus, the current design greatly reduced the linear relationship between goal and decision values, and prediction errors.

At the end of the experiment, one of the 50 prescan or 300 scanner trials was randomly selected and the outcome of that trial, and only that trial was implemented. As a result, subjects did not have to worry about spreading their \$5 dollar budget over the different items and they could treat each trial as if it were the only decision that counted.

fMRI data acquisition

The functional imaging was conducted using a Siemens 3.0 Tesla Trio MRI scanner to acquire gradient echo T2*-weighted echoplanar (EPI) images with BOLD contrast. To optimize functional sensitivity in the orbitofrontal cortex (OFC), we used a tilted acquisition in an oblique orientation of 30° to the anterior commissure–posterior commissure line

(Deichmann et al., 2003). In addition, we used an eight-channel phased array coil that yields a 40% signal increase in signal in the medial OFC over a standard head coil. Each volume comprised 44 axial slices. A total of 1064 volumes were collected over four sessions during the experiment in an interleaved ascending manner. The imaging parameters were as follows: echo time, 30 ms; field of view, 192 mm; in-plane resolution and slice thickness, 3 mm; repetition time, 2.75 s. Whole-brain high-resolution T1-weighted structural scans ($1 \times 1 \times 1$ mm) were acquired from the 16 subjects and coregistered with their mean EPI images and averaged together to permit anatomical localization of the functional activations at the group level.

fMRI data preprocessing

Image analysis was performed using SPM5 (Wellcome Department of Imaging Neuroscience, Institute of Neurology, London, UK). Images were corrected for slice acquisition time within each volume, motion corrected with realignment to the last volume, spatially normalized to the standard Montreal Neurological Institute EPI template, and spatially smoothed using a Gaussian kernel with a full width at half maximum of 8 mm. Intensity normalization and high-pass temporal filtering (using a filter width of 128 s) were also applied to the data.

fMRI data analysis

Primary model. The data analysis proceeded in three steps. First, we estimated a general linear model (GLM) with AR(1) and 2 regressors including response trials (R1), and missed trials (R2). To take advantage of the parametric nature of our design, the general linear model also included the following parametric modulators: response trials modulated by GV (M1), response trials modulated by DV (M2), and response trials modulated by PE (M3). The parametric modulators are defined as follows. Goal value equals the amount bid for the item sold in that trial during the prescan auction and, thus, is a measure of the subject's willingness to pay for the item being shown. Decision value is equal to the amount bid plus the value of x . (Recall that negative values of x represent prices, whereas positive values of x represent compensations to the subject.) Prediction error was calculated using the following equations: $B(1) = 0$; and for all $t > 1$: $PE(t) = V(t) - B(t)$; $B(t + 1) = B(t) + \lambda \times PE(t)$, where PE is the prediction error, t the trial number, $\lambda = 0.1$, and B is the expected value. The value of the current trial [$V(t)$] depended on the subject's decision. If the subject purchased the food item, then $V(t) = \text{bid}(t) + x(t) + z(t)$. However, if the subject choose not to purchase the food, then $V(t) = z(t)$. The data were modeled separately using λ values ranging from 0.1 to 0.8. These models were evaluated using the average β values from 8 mm spheres centered in the bilateral ventral striatum and encompassing the voxels reported to reflect prediction error signals for primary and monetary rewards in previous studies (Knutson et al., 2005; O'Doherty et al., 2006). $\lambda = 0.1$ was found to provide the best fit for the current data. All of the results reported below are based on this value.

Second, we calculated first-level single subject contrasts for (1) response trials modulated by goal value (M1), (2) response trials modulated by decision value (M2), and (3) response trials modulated by prediction error (M3).

Third, we calculated second-level group contrasts using a one-sample t test for each parametric regressor (GV, DV, and PE). Anatomical localizations were performed by overlaying the t maps on a normalized structural image averaged across subjects, and with reference to an anatomical atlas (Duvernoy, 1999).

Post hoc analyses. Post hoc examinations of how activity in regions of interest (ROIs) identified by the second-level group analysis scaled with the values of GV, DV, PE, x , and z were conducted by running 10 additional general linear models. In the first, trials were split according to the quartile values of GV for each subject. This resulted in a general linear model with five regressors including responses with first quartile GV [low level of goal value] (R1), responses with second quartile GV [medium-low level of goal value] (R2), responses with third quartile GV [medium-high level of goal value] (R3), responses with fourth quartile goal [[high level of goal value] (R4), and missed response trials (R5). The general linear model also included parametric modulators for regressors R1–R4, each modulated by DV and PE. The other additional

models were conducted in a similar manner, splitting the trials based on the quartile values for DV, PE, x , and z . In the quartile models for DV and x , regressors were modulated by GV and PE. In the quartile models for PE and z , regressors were modulated by GV and DV. The β weights resulting from these *post hoc* analyses were used to create the bar graphs shown in Figures 2 and 7. The time-course graphs in Figure 3 and supplemental Figure 1 (available at www.jneurosci.org as supplemental material) were created by running these same five quartile models using finite impulse response (FIR) basis functions rather than the canonical hemodynamic response function. Responses were modeled for five 2 s intervals after stimulus onset covering a total of 10 s.

To calculate the GLM in our primary analysis, the second parametric regressor (DV) was orthogonalized with respect to the first parametric regressor (GV), and the third parametric regressor (PE) was orthogonalized with respect to the first and second by the SPM5 software. As a consequence, any shared variance between the three parametric regressors was assigned to the GV regressor and any shared variance between the regressors for PE and GV or DV was assigned to first two regressors. We specified the model in this order to give GV the maximal explanatory power in all brain regions including the ventral striatum. For completeness, we also ran two additional GLMs specifying the regressors for DV or PE first. Both of these models also failed to show that activity in the ventral striatum was correlated with GV.

In Figures 2, 3, and 7, functionally defined ROIs in the medial OFC (mOFC), central OFC (cOFC), and VS were specified, respectively, using the group contrasts for GV, DV, and PE from our primary model. Voxels were included in the ROIs for DV and PE if they exceeded a threshold of $p < 0.01$ uncorrected. We used a liberal threshold to create larger ROIs that would allow for intersubject variability in the peak voxel. The ROI for GV was created using a threshold of $p < 0.0005$ uncorrected because more liberal thresholds included a large number of voxels beyond the desired anatomical location. In the mOFC ROI, the peak voxel for each subject was identified based on the GV contrast from the primary model in which GV was the first parametric modulator. In the cOFC ROI, the peak voxel for each subject was identified based on the contrast for DV from the version of the GLM in which DV was the first regressor. In the VS ROI, the peak voxel for each subject was identified based on the contrast for PE from the version of the GLM in which PE was the first regressor. This method resulted in peak voxels that were identified by regressors retaining their original shape rather than being orthogonalized with respect to another regressor.

To directly compare the fits for the GV, DV, and PE regressors, we recalculated our primary model after first z scoring the values of GV, DV, and PE within subjects to put them on the same scale. The average β for each parametric regressor across all voxels in functionally defined ROIs in the mOFC, cOFC, and VS was computed and compared within and across ROIs using Wilcoxon signed rank tests. Functional ROIs were defined by contrasts from our primary model for GV in the mOFC, DV in the cOFC, and PE in the VS. ROIs were created at a threshold of $p < 0.0005$ uncorrected in each region.

Results

Behavioral results

Figure 1 shows that subjects made purchases based on decision value computations (recall that the decision value is equal to the benefit minus the cost). As expected, subjects were more likely to purchase food items when the decision value was high and less likely to purchase when the decision value was low.

Neuroimaging results

Our analysis of the fMRI data capitalized on the fact that the nature of our design yielded three separate parameters for goal value, decision value, and prediction error on every trial. We analyzed the fMRI data using a GLM including parametric regressors based on these three values to identify regions of the brain that reflected the computation of the three value signals. The contrast for goal values showed activity in several areas in-

cluding the mOFC, the medial prefrontal cortex, and the amygdala (for a complete list of regions and coordinates, see Tables 1–3). The contrast for decision values showed activity in the cOFC. However, activity in the ventral striatum did not correlate with either goal or decision value even at liberal thresholds ($p < 0.01$, uncorrected); in fact, only the contrast for prediction error showed activity in the ventral striatum. Figure 2 shows that activity in medial OFC, central OFC, and ventral striatum best reflected goal value, decision value, and prediction error, respectively.

The findings so far support the conclusion that the BOLD signal in the mOFC is correlated with goal values, but the signal in ventral striatum is not. Because the latter part of the conclusion is based on a failure to reject the null hypothesis, we performed additional analyses to look more closely at this conclusion. In particular, we conducted *post hoc* analyses on the regions of interest in medial OFC, central OFC, and the ventral striatum identified by the primary analysis. To directly compare the fit of the regressors for GV, DV, and PE in each region, we ran the GLM again after first z scoring the three value computations within subjects to put them on the same scale. Figure 3 shows the mean and 95% confidence interval of the β (fitted parameter) for the parametric regressors for GV, DV, and PE in medial OFC, central OFC, and ventral striatum. These graphs again show that in the medial OFC there is a significant effect only for GV, in the central OFC there is an effect only for DV, and in the ventral striatum there is an effect only for PE. Furthermore, in ventral striatum, the effect of PE is significantly greater than GV, despite the fact that GV was entered into the model as the first parametric regressor and PE entered as the last resulting in any shared variance between GV and PE being attributed to GV. Although the effects of PE and DV were not significantly different in ventral striatum in this model, when the order of regressors was changed to put PE first giving it the maximum explanatory power, paired t tests showed that the effect of PE was significantly greater than DV ($t_{(15)} = 2.53$; Bonferroni corrected, $p < 0.05$) (Fig. 4). Finally, Figure 5 shows that the effect of GV was greater in medial OFC than central OFC and ventral striatum. The effect of PE was greater in the ventral striatum than the central OFC (Fig. 5). However, although there was no significant effect of PE in the medial OFC, the effect size did not differ between medial OFC and ventral striatum. Similarly, although central OFC alone showed a significant effect of DV, there was no difference in the effect sizes for DV between central OFC, medial OFC, and ventral striatum. Despite our efforts to decouple the GV, DV, and PE parameters, DV remained correlated to a certain degree with both GV ($r = 0.531$) and PE ($r = 0.573$). Thus, it is difficult to determine whether the effects of DV in medial OFC and ventral striatum are attributable to activity in those regions reflecting DV computations or shared variance between DV and GV or DV and PE.

We attempted to address the issue of correlation between value computations by examining how activity in each of these regions scaled with the separate and uncorrelated values of bid, price, and random gain that are used to compute GV, DV, and PE. Figure 6 shows that these three areas were selectively sensitive to particular aspects of the trial. Activity in the medial OFC reflected only the bid or goal value of the food item. Activity in central OFC was sensitive to both the bid and the price, which shows that this area is sensitive both to the value of the outcomes generated by actions and to the cost of taking the actions. Finally, activity in ventral striatum reflected the amount of the price and exogenous gain, which was to be expected because the total value

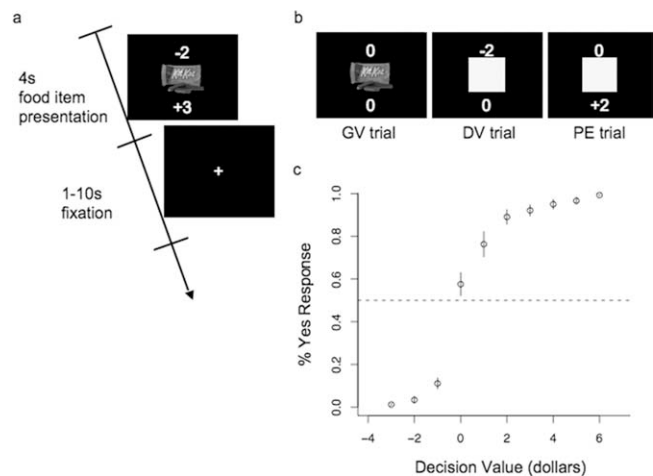


Figure 1. Task design. *a*, Subjects were presented with a high-resolution image of a food item or a yellow square with a price shown above and an action independent gain/loss shown below. Positive values indicated a gain and negative values indicated a cost or loss. Subjects indicated with a button press whether or not they would pay the price above the image for the right to eat the food item shown and the end of the experiment. The gain/loss below the image occurred regardless of the subject's choice. *b*, Trials including only food items, only price, and only gain were included to help separate the values for GV, DV, and PE. Trial order, food items, price, and gain were randomized within subjects. *c*, Percentage of "yes" responses as a function of decision value. The y -axis represents the percentage of trials on which subjects responded "yes" to purchase the food item shown. The x -axis represents the decision value of the trial. Decision value was calculated as the amount the subject bid for the food item minus the cost for that food item in the current trial. Error bars represent SEM.

of the trial, and thus prediction errors, are a linear function of them. (Recall that positive prices are compensations to the subject, whereas negative prices are subtractions from the subject's endowment.) In contrast, activity in the ventral striatum did not reflect the amount bid for the food item. Time courses for bid, price, and gain in medial OFC, central OFC, and ventral striatum are shown in supplemental Figure 1 (available at www.jneurosci.org as supplemental material).

These results indicate that goal value, decision value, and prediction error computations are supported by dissociable neural systems. This is illustrated in Figure 7, which shows the distinct areas where these computations are reflected.

Discussion

The design of the behavioral task is a critical component of the current study. In most experimental paradigms examining reward-based decision making, goal values, decision values, and prediction errors are highly correlated. This correlation makes it extremely difficult to isolate the neural basis of these computations and, thus, can lead to the misattribution of function. Although our task did not completely decouple the parameters for GV, DV, and PE, it did significantly reduce the correlations between values. This reduction was sufficient to show that activity in ventral striatum reflects prediction error and not goal value or decision value computations. The fact that some degree of correlation remains between the parameters for GV, DV, and PE does not detract from our findings because these correlations bias the data against our results.

Our data, together with the large number of studies that have reported prediction errors in the ventral striatum (McClure et al., 2003; O'Doherty et al., 2003; Abler et al., 2006; Li et al., 2006; Rodriguez et al., 2006; Tobler et al., 2006; Bray and O'Doherty, 2007; Murray et al., 2008; Rolls et al., 2008) and the studies of Plassman et al. (2007) and Rolls et al. (2008), which found goal

Table 1. Areas showing activity in the contrast of the parametric regressor for GV

Region	Side	BA	MNI coordinates			Peak Z
Inferior occipital gyrus	R	18	33	−81	−9	5.41
Middle frontal gyrus, orbital sulcus	R	11/47	30	27	−12	5.18
Middle frontal gyrus, orbital sulcus	L	11/47	−24	30	−18	4.8
Medial frontal gyrus, medial OFC	B	11/25	−1	27	−18	4.25
Anterior cingulate cortex	B	24/32	1	39	6	4.16
Inferior frontal gyrus	R	45	51	24	21	3.93
Inferior frontal gyrus	R	45	45	33	9	3.92
Medial orbital gyrus	R	28/47	21	9	−21	3.7
Parahippocampal gyrus	R	27	21	−30	−3	3.69
Medial frontal gyrus	R	10	6	51	0	3.61
Amygdala	L		−15	−9	−15	3.56

Height threshold, $T = 4.09$, false discovery rate corrected $p < 0.01$; extent threshold, $k = 5$ voxels. L, Left; R, right.

Table 2. Areas showing activity in the contrast of the parametric regressor for DV

Region	Side	BA	MNI coordinates			Peak Z
Middle frontal gyrus, posterior orbital gyrus	L	11	−27	36	−6	3.66
Middle frontal gyrus, posterior orbital gyrus	R	11	24	36	−3	3.16*

Height threshold, $T = 3.73$; $p < 0.05$, small volume false discovery rate (FDR) corrected, for anatomical mask of the ventral prefrontal cortex; extent threshold, $k = 10$ voxels. L, Left; R, right.

* $p < 0.06$, FDR corrected.

Table 3. Areas showing activity in the contrast of the parametric regressor for PE

Region	Side	BA	MNI coordinates			Peak Z
Ventral striatum, caudate/putamen	R		15	21	−6	3.32

Height threshold, $T = 3.73$; $p < 0.05$, small volume false discovery rate corrected, for 10 mm sphere centered on peak coordinates for PE in the ventral striatum in the study by Knutson et al. (2005); extent threshold, $k = 10$ voxels. L, Left; R, right.

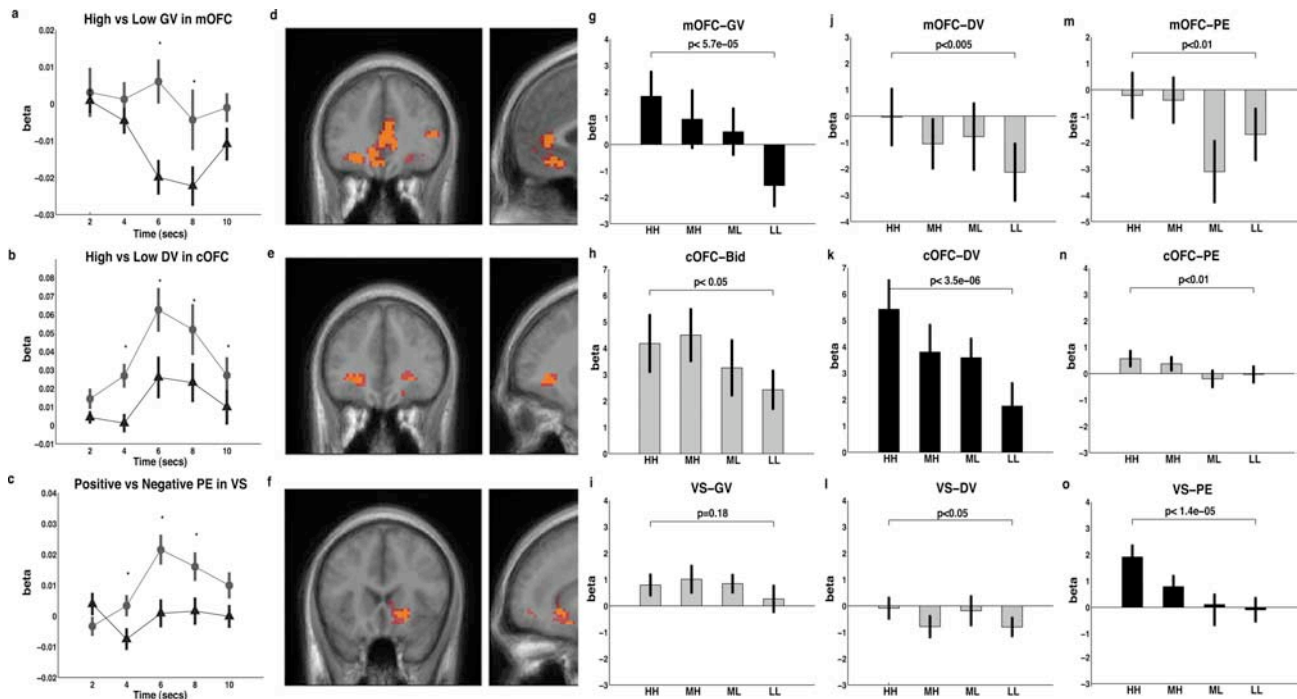


Figure 2. Regions associated with goal value, decision value, and prediction error. For time courses in **a–c**, high values are shown with circular markers connected by solid lines. Low values are marked by triangles connected by dotted lines. Markers represent the mean β value from a finite impulse response model at each time point. Error bars represent the SEM. Asterisks indicate points in time where there is a significant difference between high and low based on paired t tests at $p < 0.05$. Activation maps in **d, e, and f** are shown at a threshold of $p < 0.005$ uncorrected and an extent threshold of 10 voxels. Voxels in red are significant at $p < 0.005$ uncorrected, whereas voxels in yellow remain significant at $p < 0.0001$ uncorrected. **d**, Activity in the OFC and anterior cingulate cortex correlated with goal value. **e**, Activity in the central OFC correlated with decision value. **f**, Activity in the ventral striatum correlated with prediction error. **g–o**, Bar graphs represent the mean β value \pm SE at high (HH), medium high (MH), medium low (ML), and low (LL) levels of GV, DV, and PE. The p values displayed in the bar graphs were determined based on paired t tests between high and low trials. Bars in black show the value that best fits each region.

value signals in the OFC and medial prefrontal cortex but not in the striatum, suggest that activity in the ventral striatum reflects prediction error but not goal value computations. This is an im-

portant result for neuroeconomics and behavioral neuroscience because the interpretation of ventral striatal activity as reflecting goal values is fairly widespread (Aharon et al., 2001; Knutson et




Figure 3. Fit of each value computation to activity in medial OFC, central OFC, and ventral striatum. The y -axis shows the mean β for GV, DV, and PE. The error bars represent the 95% confidence interval of the mean. The different value computations (GV, DV, PE) are shown on the x -axis. Dark gray bars highlight value computations for which the fit to activity is greater than zero. The effects of GV, DV, and PE were compared in each region using Wilcoxon signed rank tests and those results are represented by the p values on each graph.

al., 2001, 2007; Hariri et al., 2006; Yacubian et al., 2006, 2007; Kable and Glimcher, 2007; Schaefer and Rotte, 2007). Here, we use two recent examples to argue that this misinterpretation is attributable to a limitation of previous task designs that confounded goal values with prediction PE signals. Knutson et al. (2007) showed that ventral striatal activity was correlated with subject's preference ratings for various consumer products. Based on this correlation, they suggested that ventral striatal activity coded for the goal value of the object. A problem with their interpretation is that in their design the prediction error at the time of product presentation (equal to the goal value of the current product, minus the predicted value derived from previously shown products) is very highly correlated with the goal value of the product. Therefore, another interpretation of their results, consistent with the data in this study, is that the ventral striatum is encoding prediction errors. Similarly, Kable and Glimcher (2007) reported that activity in medial prefrontal cortex, posterior cingulate and ventral striatum encode subjective goal value in a temporal discounting paradigm. As in the study by Knutson et al. (2007), in Kable and Glimcher's (2007) experiment, prediction errors and goal values are very highly correlated and, thus, it is difficult to distinguish between areas that encode each computation. Our data suggest that their interpretation misattributes the prediction error signals in the ventral striatum as goal value computations for the same reasons discussed above.

We emphasize that the current data by no means diminish the importance of the ventral striatum in reward learning and decision making, but rather serve to clarify its role in these processes. The computation of goal value, decision value, and prediction error are all necessary for economic decision making. Regions involved in these computations must work in concert to determine the optimal course of action. Knowing the neural basis of these computations provides a basis for determining the underlying dysfunction when specific deficits are seen in decision making. With regard to regional specificity, it is important to note




Figure 4. Fit for GV, DV, and PE to activity in the ventral striatum. The β values in this figure come from a secondary fMRI analysis model where PE was entered as the first parametric regressor for response giving it the maximum explanatory power. The y -axis shows the mean β for GV, DV, and PE. The error bars represent the 95% confidence interval of the mean. The different value computations (GV, DV, PE) are shown on the x -axis. Dark gray bars highlight value computations for which the fit to activity is greater than zero. The effects of GV, DV, and PE were compared using paired t tests and those results are represented by the p values on the graph.

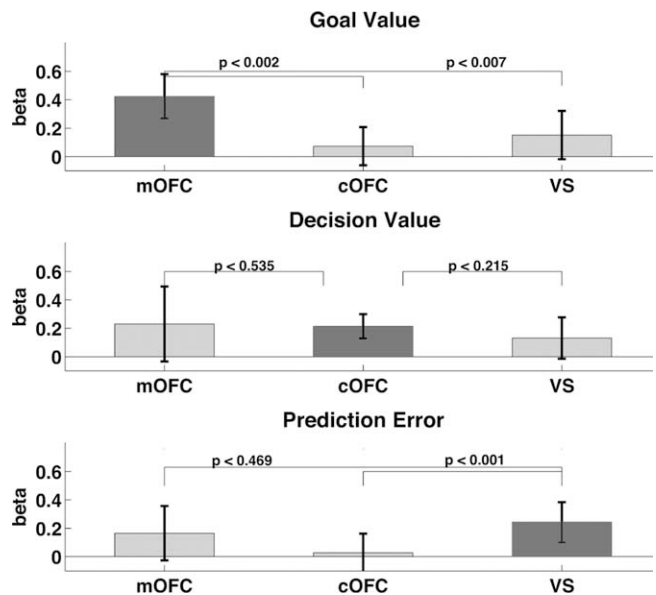


Figure 5. Comparison of the fits for GV, DV, and PE in mOFC, cOFC, and VS. The y -axis shows the mean β for value specified in the graph title. The error bars represent the 95% confidence interval of the mean. The regions of interest (mOFC, cOFC, and VS) are shown on the x -axis. Dark gray bars highlight regions where the fit to activity is greater than zero. The effects fits were compared between regions using Wilcoxon signed rank tests and those results are represented by the p values on each graph.

that we are using fMRI to measure changes in the BOLD signal, and this BOLD signal primarily reflects the input to and local processing within a brain region rather than the efferent signals from that region (Logothetis, 2002). For example, input from midbrain dopamine neurons that have been shown to express prediction error signals (Schultz et al., 1997; Hollerman and Schultz, 1998; Bayer and Glimcher, 2005; Bayer et al., 2007) is likely to contribute to the BOLD signal in the ventral striatum. Thus, although our data show dissociations in the neural basis of goal value, decision value, and prediction error computations, these computations are still supported by interconnected neural systems rather than isolated brain regions.

An important application of these results, and of related studies in neuroeconomics, is to neuropathologies involving decision-making deficits such as addiction. Work in animal models and human neuroimaging studies have shown differential involvement of prefrontal and ventral striatal regions as drug use moves from acute initial experimentation to addiction and relapse (Kalivas and Volkow, 2005). Knowing the precise value computations performed in each region could aid in the devel-

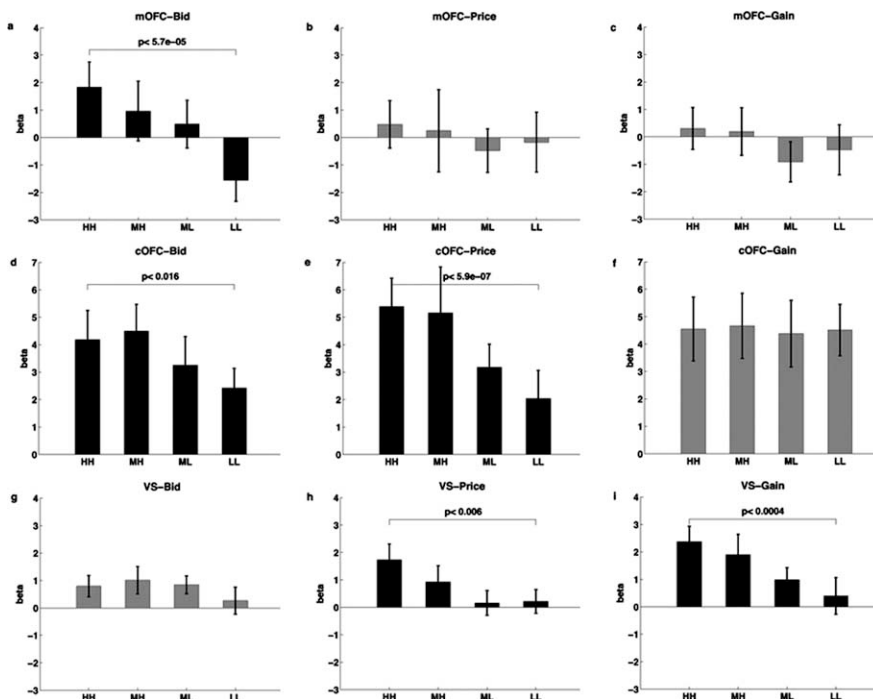


Figure 6. Selectivity of medial OFC, central OFC, and ventral striatum to bid, price, and random gain. Bar graphs represent the mean β value \pm SE on the y -axis and high (HH), medium high (MH), medium low (ML), and low (LL) levels of bid, price or random gain on the x -axis. Positive prices represented compensations to the subject whereas negative prices represented subtractions from the subject's endowment. **a–c**, Activity in medial OFC for the different levels of bid, price, and gain. **d–f**, Activity in central OFC for the different levels of bid, price, and gain. **g–i**, Activity in ventral striatum for the different levels of bid, price, and gain. The p values displayed in the bar graphs were determined based on paired t tests between high and low trials. Bars in black show values with a significant effect.




Figure 7. Combined activation maps for GV, DV, and PE. Activation maps are shown at an uncorrected threshold of $p < 0.001$. Activity correlated with GVs in the mOFC is shown in red, activity correlated with DVs in the cOFC is shown in yellow, and activity correlated with PEs in the ventral striatum is shown in green.

opment of behavioral and pharmacotherapeutic interventions targeted at specific stages of addiction. A better understanding of how the brain computes different value signals in the context of decision making might also help to determine the mechanisms underlying altered reward processing in neuropathologies like depression, and schizophrenia (Tremblay et al., 2005; Forbes et al., 2006; Juckel et al., 2006; Murray et al., 2008). These possibilities are clearly important avenues for future research.

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.
- Aharon I, Etcoff N, Ariely D, Chabris CF, O'Connor E, Breiter HC (2001) Beautiful faces have variable reward value: fMRI and behavioral evidence. *Neuron* 32:537–551.
- Bayer HM, Glimcher PW (2005) Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* 47:129–141.
- Bayer HM, Lau B, Glimcher PW (2007) Statistics of midbrain dopamine neuron spike trains in the awake primate. *J Neurophysiol* 98:1428–1439.
- Becker G, DeGroot M, Marshak J (1964) Measuring utility by a single response sequential method. *Behav Sci* 9:226–232.
- Bray S, O'Doherty J (2007) Neural coding of reward-prediction error signals during classical conditioning with attractive faces. *J Neurophysiol* 97:3036–3045.
- Deichmann R, Gottfried JA, Hutton C, Turner R (2003) Optimized EPI for fMRI studies of the orbitofrontal cortex. *NeuroImage* 19:430–441.
- Dom G, Sabbe B, Hulstijn W, van den Brink W (2005) Substance use disorders and the orbitofrontal cortex: systematic review of behavioural decision-making and neuroimaging studies. *Br J Psychiatry* 187:209–220.
- Duvernoy HM (1999) The human brain surface, blood supply, and three-dimensional sectional anatomy, Ed 2. New York: Springer.
- Forbes EE, Christopher May J, Siegle GJ, Ladouceur CD, Ryan ND, Carter CS, Birmaher B, Axelson DA, Dahl RE (2006) Reward-related decision-making in pediatric major depressive disorder: an fMRI study. *J Child Psychol Psychiatry* 47:1031–1040.
- Hariri AR, Brown SM, Williamson DE, Flory JD, de Wit H, Manuck SB (2006) Preference for immediate over delayed rewards is associated with magnitude of ventral striatal activity. *J Neurosci* 26:13213–13217.
- Hollerman JR, Schultz W (1998) Dopamine neurons report an error in the temporal prediction of reward during learning. *Nat Neurosci* 1:304–309.
- Juckel G, Schlagenhauf F, Koslowski M, Wustenberg T, Villringer A, Knutson B, Wrase J, Heinz A (2006) Dysfunction of ventral striatal reward prediction in schizophrenia. *NeuroImage* 29:409–416.
- Kable JW, Glimcher PW (2007) The neural correlates of subjective value during intertemporal choice. *Nat Neurosci* 10:1625–1633.
- Kalivas PW, Volkow ND (2005) The neural basis of addiction: a pathology of motivation and choice. *Am J Psychiatry* 162:1403–1413.
- Knutson B, Adams CM, Fong GW, Hommer D (2001) Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *J Neurosci* 21:RC159.
- Knutson B, Taylor J, Kaufman M, Peterson R, Glover G (2005) Distributed neural representation of expected value. *J Neurosci* 25:4806–4812.
- Knutson B, Rick S, Wimmer GE, Prelec D, Loewenstein G (2007) Neural predictors of purchases. *Neuron* 53:147–156.
- Li J, McClure SM, King-Casas B, Montague PR (2006) Policy adjustment in a dynamic economic game. *PLoS ONE* 1:e103.
- Logothetis NK (2002) The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. *Philos Trans R Soc Lond B Biol Sci* 357:1003–1037.
- McClure SM, Berns GS, Montague PR (2003) Temporal prediction errors in a passive learning task activate human striatum. *Neuron* 38:339–346.
- Murray GK, Corlett PR, Clark L, Pessiglione M, Blackwell AD, Honey G, Jones PB, Bullmore ET, Robbins TW, Fletcher PC (2008) Substantia nigra/ventral tegmental reward prediction error disruption in psychosis. *Mol Psychiatry* 13:267–276.
- 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 JP, Buchanan TW, Seymour B, Dolan RJ (2006) Predictive neural coding of reward preference involves dissociable responses in human ventral midbrain and ventral striatum. *Neuron* 49:157–166.
- Paulus MP (2007) Decision-making dysfunctions in psychiatry—altered

- homeostatic processing? *Science* 318:602–606.
- Plassmann H, O'Doherty J, Rangel A (2007) Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *J Neurosci* 27:9984–9988.
- Rodriguez PF, Aron AR, Poldrack RA (2006) Ventral-striatal/nucleus-accumbens sensitivity to prediction errors during classification learning. *Hum Brain Mapp* 27:306–313.
- Rolls ET, McCabe C, Redoute J (2008) Expected value, reward outcome, and temporal difference error representations in a probabilistic decision task. *Cereb Cortex* 18:652–663.
- Schaefer M, Rotte M (2007) Favorite brands as cultural objects modulate reward circuit. *NeuroReport* 18:141–145.
- Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. *Science* 275:1593–1599.
- Sutton RS, Barto AG (1998) Reinforcement learning. Cambridge, MA: MIT.
- Tobler PN, O'Doherty JP, Dolan RJ, Schultz W (2006) Human neural learning depends on reward prediction errors in the blocking paradigm. *J Neurophysiol* 95:301–310.
- Tremblay LK, Naranjo CA, Graham SJ, Herrmann N, Mayberg HS, Hevenor S, Busto UE (2005) Functional neuroanatomical substrates of altered reward processing in major depressive disorder revealed by a dopaminergic probe. *Arch Gen Psychiatry* 62:1228–1236.
- Yacubian J, Glascher J, Schroeder K, Sommer T, Braus DF, Buchel C (2006) Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. *J Neurosci* 26:9530–9537.
- Yacubian J, Sommer T, Schroeder K, Glascher J, Kalisch R, Leuenberger B, Braus DF, Buchel C (2007) Gene-gene interaction associated with neural reward sensitivity. *Proc Natl Acad Sci USA* 104:8125–8130.