

Differential Encoding of Losses and Gains in the Human Striatum

Ben Seymour,¹ Nathaniel Daw,² Peter Dayan,² Tania Singer,^{1,3} and Ray Dolan¹

¹Wellcome Trust Centre for Neuroimaging, UCL, London WC1N 3BG, United Kingdom, ²Gatsby Computational Neuroscience Unit, and ³Institute of Cognitive Neuroscience, London WC1N 3AR, United Kingdom

Studies on human monetary prediction and decision making emphasize the role of the striatum in encoding prediction errors for financial reward. However, less is known about how the brain encodes financial loss. Using Pavlovian conditioning of visual cues to outcomes that simultaneously incorporate the chance of financial reward and loss, we show that striatal activation reflects positively signed prediction errors for both. Furthermore, we show functional segregation within the striatum, with more anterior regions showing relative selectivity for rewards and more posterior regions for losses. These findings mirror the anteroposterior valence-specific gradient reported in rodents and endorse the role of the striatum in aversive motivational learning about financial losses, illustrating functional and anatomical consistencies with primary aversive outcomes such as pain.

Key words: basal ganglia; classical conditioning; reinforcement learning; neuroeconomics; opponent process; decision making

Introduction

A wealth of human and animal studies implicates ventral and dorsal regions of the striatum in aspects of the learned control of behavior in the face of rewards and punishments. In experiments involving primary rewards and punishments, the blood oxygenation level-dependent (BOLD) signal in the human striatum measured using functional magnetic resonance imaging (fMRI) covaries closely with key learning signals used by abstract learning models (O'Doherty et al., 2003; Haruno et al., 2004; Seymour et al., 2004; Tanaka et al., 2004, 2006; Yacubian et al., 2006). These algorithms originate in sound psychological learning accounts and are known to acquire normative predictions and affectively optimal behaviors (Sutton and Barto, 1981, 1990; Barto, 1995).

However, two related sets of findings, regarding the orientation of this signal and the relationship between rewards and punishments, remain difficult to accommodate fully under this interpretation. First, the BOLD signal seen in the striatum typically takes the form of a signed prediction error with baseline activity when outcomes match their predictions and above- and below-baseline excursions and when outcomes are more or less than expected, respectively. Of course, rewards and punishments have opposite valences, with a negative punishment (e.g., one expected but omitted) bearing a close computational and psychological

relationship with a “positive” reward. However, in experiments that involve cues that predict rewards exclusively (which can be presented or omitted), or primary punishments exclusively (which can also be presented or omitted), the BOLD signals are apparently oppositely oriented, with positive BOLD excursions accompanying both positive reward and positive punishment and below-baseline excursions accompanying both negative (or omitted) reward and punishment (Delgado et al., 2000; Knutson et al., 2000; Berridge et al., 2001; Breiter et al., 2001; Pagnoni et al., 2002; Elliott et al., 2003; O'Doherty et al., 2003; Jensen et al., 2003; Zink et al., 2003; Seymour et al., 2004, 2005; Tanaka et al., 2004; Nieuwenhuis et al., 2005; Yacubian et al., 2006).

Second, in the above experiments that involve financial costs (in contrast to those involving primary punishments such as physical pain), the striatal BOLD signal is typically observed to be oriented as in rewarding tasks, with monetary gains associated with positive BOLD activations, and losses with “sub-baseline” signals. Indeed, there are few reports of any brain areas showing a positive BOLD response to financial loss at all, and although this is not exclusively the case [for instance, in the amygdala (Yacubian et al., 2006) and insula cortex (Knutson et al., 2007)], it has been suggested that monetary losses and gains might be fully processed by a unitary (appetitive) system, centered on the striatum (Tom et al., 2007).

Potential explanations for these puzzles include the possibility that the striatal BOLD signal reflects the release of different neuromodulators (Daw et al., 2002; Doya, 2002) (one reporting prediction errors of each valence) or the possibility that neighboring regions of the striatum report on the different valences (Reynolds and Berridge, 2001, 2002). Indeed, there are sound psychological and neurophysiological reasons to think that separate, opponent systems are responsible for the two valences (Konorski, 1967; Dickinson and Dearing, 1979; Gray, 1991). But, on this interpretation, it remains unclear why different circum-

Received Jan. 29, 2007; revised Feb. 26, 2007; accepted March 18, 2007.

This work was supported by a grant from Wellcome Trust Programme to R.D. P.D. is supported by Gatsby Charitable Foundation. N.D. was supported by the Royal Society. T.S. was supported by the Medical Research Council. We thank Eric Featherstone for technical help. We declare no competing financial interests.

Correspondence should be addressed to Ben Seymour, Wellcome Trust Centre for Neuroimaging, UCL, 12 Queen Square, London WC1N 3BG, UK. E-mail: bseymour@fil.ion.ucl.ac.uk.

N. Daw's present address: Center for Neural Science and Department of Psychology, New York University, New York, NY 10003.

T. Singer's present address: Center for Social Neuroscience and Neuroeconomics, University of Zürich, Blümlisalpstrasse 10 8006 Zürich, Switzerland.

DOI:10.1523/JNEUROSCI.0400-07.2007

Copyright © 2007 Society for Neuroscience 0270-6474/07/274826-06\$15.00/0

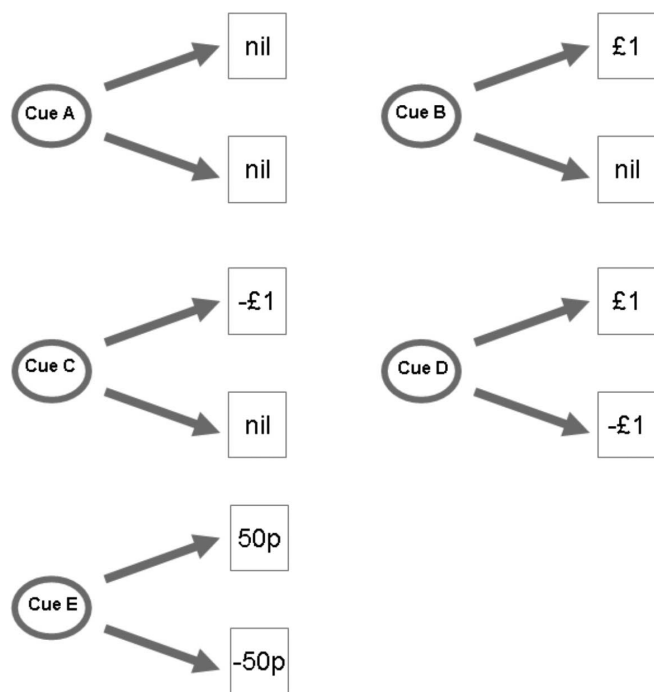


Figure 1. Experimental design. Visual cues were presented for 3.5 s and followed immediately with the outcome depicting the outcome amount, which was displayed for 1.5 s. For the analysis, events were marked at the time of the outcome, and linear contrasts were performed between the different outcome types. p, Pence.

stances implicate each signal (for instance, why pain is apparently reported by a punishment-oriented prediction error but monetary losses are not). We designed a Pavlovian conditioning experiment, involving mixed gain and loss outcomes, to address these underlying issues.

Materials and Methods

The key requirements for the task were to integrate monetary predictions about gains and losses and to avoid framing the problem entirely in terms of one valence. One strategy for mitigating the latter, at the potential expense of low experimental power and only subtle outcomes, is to make the task involve predictions alone, with no requirement for action, and thereby avoiding subjects having expectations that they will be able to win. Thus, we used fMRI to examine striatal representations of financial loss in tasks that involve mixed gains and losses, using a probabilistic first-order Pavlovian learning task with monetary outcomes. Importantly, the design included both mixed and nonmixed valence outcome probabilities, allowing us to look specifically at the influence on outcome representations (specifically, the prediction error) of the context provided by the nonexperienced outcome.

Subjects. Twenty four (11 female) subjects, age range 19–35, participated in the study. All were free of neurological or psychiatric disease and fully consented to participate. The study was approved by the Joint National Hospital for Neurology and Neurosurgery (University College London Hospitals NHS trust) and Institute of Neurology (University College London) Ethics Committee. Subjects were remunerated by amounts corresponding to their actual winnings during the task (mean zero) added to a fixed prestated amount for time and inconvenience (£20).

Stimuli and Task. We performed a probabilistic first-order Pavlovian delay conditioning task, with visual cues predictively paired with monetary outcomes, as demonstrated in Figure 1. Visual cues were presented on a computer monitor projected onto a screen, visible via an angled mirror on top of the fMRI headcoil. The visual stimuli were presented for 3.5 s, and on termination, were followed immediately by a 1.5 s duration image of their outcome, an empty circle (no outcome), a 50 pence coin, or a £1.00 coin, below which was written in bold letters the amount, and

Table 1. Cue-outcome contingencies

Cue	Outcome	Probability
Neutral	£0	1
Univalent reward	£0	0.5
	£1	0.5
Univalent loss	£0	0.5
	£1	0.5
Bivalent cue (£1)	£1	0.5
	£1	0.5
Bivalent cue (50 pence)	£0.50	0.5
	£0.50	0.5

whether they had won or lost (for example, “WIN £1.00”). The five cues predicted outcomes shown in Table 1.

The visual stimuli were abstract colored images, ~6 cm in diameter viewed on the projector screen from a distance of ~50 cm. They were fully balanced and randomized across subjects and matched for luminance. We presented 200 trials over two sessions, with each trial being presented with a jittered interval of 2–6 s.

Preference task. After the conditioning task, we assessed the acquisition of Pavlovian cue values using a preference task, involving forced choices between pairs of cues. Each cue was presented alongside (horizontally adjacent) each other cue, and subjects (still inside the fMRI scanner) made an arbitrary preference judgment between them, using a response keypad (no outcomes were delivered). Each possible combination was presented five times (making 50 trials), in random order, and with the position of each cue (on the left or right side of the screen) also randomized. The total number of preference choices for each cue was summed (in a similar manner to a league table) and nonparametric comparisons assessed statistically.

Pupillometry. Pupil diameter was measured online during fMRI scanning by an infrared eye tracker (Model 504; Applied Sciences Laboratories, Waltham MA) recording at 60 Hz. Pupil recordings were analyzed on an event-related trial basis and used to find evidence of basic conditioning among the reward, aversive, and neutral cues. We used the peak light reflex after presentation of the cue, which is a standard measure of autonomic arousal (Bitsios et al., 2004), and we performed analyses using a repeated-measures ANOVA and *post hoc t* tests. Technical problems led to the data not being collected for four of the 24 subjects.

fMRI. Subjects learned the task *de novo* in an fMRI scanner to allow us to record regionally specific neural responses. Functional brain images were acquired on a 1.5T Sonata Siemens AG (Erlangen, Germany) scanner. Subjects lay in the scanner with foam head-restraint pads to minimize any movement. Images were realigned with the first volume, normalized to a standard echo-planar imaging template, and smoothed using a 6 mm full-width at half-maximum Gaussian kernel. Realignment parameters were inspected visually to identify any potential subjects with excessive head movement, and none were found. Images were analyzed in an event-related manner using the general linear model, with the onsets of each outcome represented as a stick function to provide a stimulus function. Regressors of interest (10 in total) were then generated by convolving the stimulus function with a hemodynamic response function (HRF). Effects of no interest included the onsets of visual cues and realignment parameters from the image preprocessing to provide additional correction for residual subject motion.

From the outcomes, linear contrasts of the statistical parametric maps (SPMs) from the outcomes were taken to a group-level (random-effects) analysis by way of a one-sample *t* test. Montreal Neurological Institute coordinates and statistical *z*-scores are reported in figure legends.

Group level activations were localized according to the group averaged structural scan. Activations were checked on a subject-by-subject basis using individual normalized structural scans, acquired after the functional test-scanning phase, to ensure correct localization. We report activity in areas in which we had previous hypotheses, based on previous data, although without specification of laterality. These regions have established roles in both aversive and appetitive predictive learning, and included the putamen, caudate, nucleus accumbens, midbrain (substan-

tia nigra), amygdala, anterior insula cortex, and orbitofrontal cortex. We report activations at a threshold of $p < 0.001$, which survive false discovery rate (FDR) correction at $p < 0.05$ for multiple comparisons using an 8 mm sphere around coordinates based on previous studies. Note that in Figures 3 and 4, we use a threshold of $p < 0.005$ (with a five voxel extent threshold) for display purposes. No other activation was found outside our areas of interest that survived whole-brain correction for multiple comparisons using FDR correction at $p < 0.05$. Details and statistics of all significant activations appear in the figure legends of the appropriate contrasts.

We performed two central analyses. One involved trial-based contrasts for positive reward and positive loss prediction errors: (1) positive-reward prediction error, bivalent £1.00 win outcome minus univalent £1.00 win outcome; and (2) positive-loss prediction error, bivalent £1.00 loss minus univalent £1.00 loss outcome.

In the second analysis, we used a simple reinforcement learning model to generate a signal corresponding to the outcome prediction error, which, as in previous studies, was applied as a regressor to the imaging data (O'Doherty et al., 2003). Here, we used a temporal difference model with a learning rate of $\alpha = 0.3$. This was based primarily on our previous data from Pavlovian learning (Seymour et al., 2005), although it should be noted that converging evidence from a number of studies, examining both Pavlovian and instrumental learning, has supported a comparable learning rate. Furthermore, the results presented below are robust to changes in learning in realistic ranges (0.3–0.7) based on previous studies. In this model, the value v of a particular cue (referred to as a state s) is updated according to the following learning rule: $v(s) \leftarrow v(s) + \alpha\delta$, where δ is the prediction error. This is defined as follows: $\delta = r_t - v(s)_t$, where r is the return (i.e., the amount of money). We used a parametric design, in which the temporal difference prediction error modulated the stimulus functions on a stimulus-by-stimulus basis. The statistical basis of this approach has been described previously (Buchel et al., 1998; O'Doherty et al., 2003). Regressors corresponding to the outcome prediction errors were then generated by convolving the stimulus function with an HRF.

Finally, we considered two further trial-based contrasts. One contrast sought the representation of the negative prediction errors: (1) negative reward prediction error, univalent £1.00 win outcome minus bivalent £1.00 win outcome; and (2) negative loss prediction error, univalent £1.00 loss minus bivalent £1.00 loss outcome. These contrasts afforded no significant difference at our thresholds.

The second contrast considered residual activity in the striatum, when equal prediction errors are subtracted: (1) zero net prediction error, bivalent £0.50 win outcome minus univalent £1.00 win outcome; and (2) zero net prediction error, bivalent £0.50 loss outcome minus univalent £1.00 loss outcome. As expected from standard models, none of these contrasts yielded a significant difference.

To address the possibility that cue-related responses might confound identification of prediction error-related responses, we repeated all analyses (both trial based and model based), with the inclusion of a single cue-related regressor. Inspection of the regressor covariance matrix relating to parameter estimability after convolution of the design matrix with the HRF suggested that the models were not over specified. Indeed, for the model-based analysis, there was no correlation between the cue regressor and the prediction error. In keeping with this, the SPMs for both trial-based and model-based analyses showed minimal changes in results. Second, we repeated the trial-based analysis with full specification of the identity of the cue (i.e., with five separate cue regressors). As above, this did not alter the results to any substantial degree. Third, we orthogonalized the outcome regressors with respect to the cue regressors,

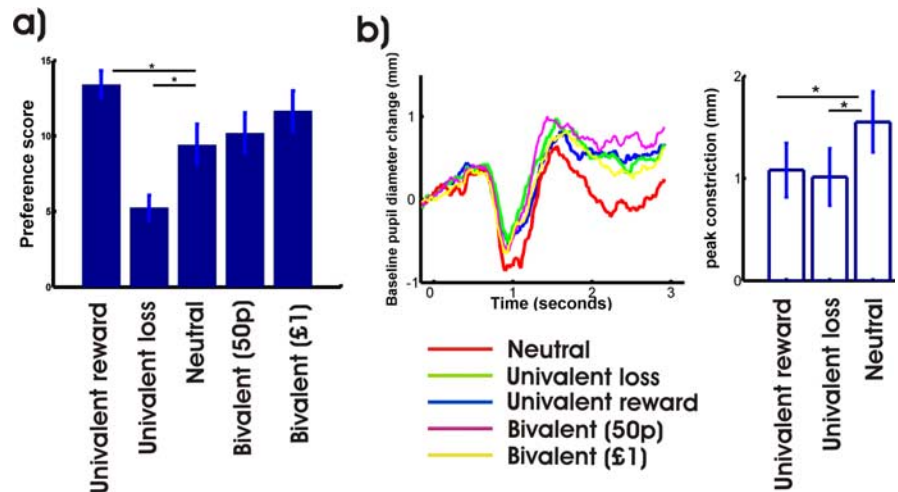


Figure 2. Behavioral results. *a*, Preference scores. One-way repeated-measures ANOVA; $F_{(4,92)} = 5.572$; $p = 0.0005$; *post hoc* two-tailed *t* test yielded significant differences between univalent reward and neutral and univalent loss and neutral ($p < 0.05$). *b*, Mean pupillometry, average across all trials across learning, in a trial-specific manner. We looked for a basic effect of conditioning between the rewarding, aversive, and neutral cue, which is a standard measure of conditioning. Repeated-measures ANOVA revealed a significant effect of trial type. $F_{(2,19)} = 3.342$; $p < 0.05$. The *post hoc t* tests showed a significant effect (increased amplitude of light reflex) for both rewarding and aversive cues when compared with the neutral cue ($p < 0.05$). Error bars indicate SEM. *p*, Pence.

and again, the results changed only minimally (in either direction). No significant correlations were found with the cue-related regressors.

Results

Behavioral results

The postconditioning preference task demonstrated significant preference for the cue associated with univalent reward cue over the neutral cue, in turn preferred to the univalent loss cue. Preference scores for the bivalent cues were slightly above those of the neutral cue, for which the expected value is equivalent (Fig. 2, see figure legend for statistics).

Pupil diameter, which is an autonomic measure of arousal, also provided evidence of basic conditioning of the rewarding and aversive cues, compared with the neutral cue (Fig. 2, see figure legend for statistics).

fMRI results

The experimental design allowed comparison of neural responses to winning money in two conditions: one in which the alternative was winning nothing, and one in which the alternative was losing. Similarly, it allows comparison of neural activity corresponding to losing money when the alternative was nothing or winning. Thus, the key BOLD contrasts were between the univalent and bivalent outcomes for both gain and loss outcomes, because these reveal appetitive and aversive prediction errors, respectively, specifically relating to the outcomes associated with mixed-valence predictions.

In the appetitive case (bivalent cue followed by a £1 reward, univalent reward cue followed by £1 reward), this corresponds to a positive relative reward prediction error of 50 pence and was associated with activation in the ventral striatum (Fig. 3*a*). In the aversive case (bivalent cue followed by a £1 loss, univalent loss cue followed by £1 loss), this corresponds to a positive aversive prediction error of -50 pence and was also associated with activation in the ventral striatum (Fig. 3*b*). The peak of the aversive prediction error was slightly posterior to the appetitive prediction error, as shown in the sagittal section displayed in Figure 3*c*.

However, the magnitude of these peaks was such that this




Figure 3. fMRI simple bivalent, univalent contrasts. **a**, Aversive prediction error, right ventral striatum [Montreal Neurological Institute (MNI) coordinates (x, y, z) : $-16, 0, -10$; $z = 3.74$; 46 voxels at $p < 0.005$]. This contrast also revealed a peak in the right anterior insula [data not shown; MNI coordinates (x, y, z) : $30, 18, -12$; $z = 3.60$]. Yellow corresponds to $p < 0.005$; magenta corresponds to $p < 0.001$. **b**, Reward prediction error, right ventral striatum [MNI coordinates (x, y, z) : $-16, 6, -6$; $z = 3.38$; 28 voxels at $p < 0.005$]. Yellow corresponds to $p < 0.005$; magenta corresponds to $p < 0.001$. **c**, Sagittal view showing the two peaks, reward (green) and aversive (red).




Figure 4. fMRI temporal difference (TD) model. **a**, Left, Aversive TD error, right mid striatum [Montreal Neurological Institute (MNI) coordinates (x, y, z) : $-20, -4, 6$; $z = 3.89$; $p < 0.005$; 21 voxels]. Yellow corresponds to $p < 0.005$; magenta corresponds to $p < 0.001$. Right, The image is also shown in sagittal section (in red). **b**, Left, Appetitive TD error, right ventral striatum (nucleus accumbens) [MNI coordinates (x, y, z) : $10, 6, -1$; $z = 3.13$; shown at $p < 0.005$; 15 voxels] and left ventral striatum (nucleus accumbens) [MNI coordinates (x, y, z) : $-12, 6, -18$; $z = 3.62$; 14 voxels]. Yellow corresponds to $p < 0.005$; magenta corresponds to $p < 0.001$. Right, The image is also shown in sagittal section (in green).

analysis could not reliably differentiate the location of appetitive and aversive prediction errors, with the activity in each peak being only insignificantly greater than activity associated with the contrast that defined the other peak. Furthermore, the trial-based contrasts (three and four) testing for negative prediction errors of either valence showed no significant effects. This could reflect an asymmetry reported at the spiking level for dopaminergic neurons (Fiorillo et al., 2003; Bayer and Glimcher, 2005; Niv et al., 2005; Morris et al., 2006), where positive errors are coded more strongly than negative ones. It may additionally be because of the relatively crude trial-based measures.

Therefore, we considered a more sensitive analysis based on a temporal difference learning model. This model is known to offer a good account of the neurophysiological responses of dopami-

nergic cells associated with Pavlovian learning about rewards in monkeys (Montague et al., 1996; Schultz et al., 1997) and has been successfully used in human fMRI to probe prediction-error components of the BOLD signal from the striatal targets of these cells (O'Doherty et al., 2003; Haruno et al., 2004; Seymour et al., 2004; Tanaka et al., 2004, 2006). We applied the model as in previous studies and used the prediction error occurring at the time of the outcomes generated by this model as a parametric regressor in the fMRI data analysis. This model incorporates both positive and negative prediction errors and thus identifies valence-specific responses. Aversive prediction errors

should be negatively correlated with this signal; appetitive prediction errors should be positively correlated with it. Therefore, unlike the trial-based contrasts, this analysis should identify areas that are specific to either valence.

In other words, this analysis identifies subject-specific, trial-specific activity that correlates with the prediction errors fitted by the temporal difference learning model. This analysis was applied solely to the bivalent cues (because it is during these trials that we expected to find opponent prediction error representations).

Activity associated with an aversive temporal difference outcome prediction error was observed posteriorly in the mid putamen (Fig. 4a). Activity associated with an appetitive temporal difference outcome prediction error was observed in more anterior ventral striatum, in close proximity to the nucleus accumbens (Fig. 4b). These activations are presented in sagittal sections (Fig. 4c, green and red, respectively), to permit comparison with the simple prediction error contrasts shown in Figure 3c.

Given a recent report of identification of an aversive prediction error in the amygdala (Yacubian et al., 2006), we looked at a reduced threshold (uncorrected $p < 0.01$) specifically in that region. However, no correlated activity was identified.

Discussion

Our results suggest a partial resolution to the puzzles outlined in the introduction. The data suggest that aversive and appetitive prediction errors may be represented in a similar manner, albeit somewhat spatially resolvable along an axis of the striatum. The appetitive prediction error appears to direct the BOLD signal in more anterior and more ventral regions than the aversive prediction error. Furthermore, it appears that the prevalence of each sort of coding may depend on the affective context.

Although one should be cautious regarding the topographic spatial resolution of fMRI, the anterior–posterior gradient resembles that seen in stimulation studies of the ventral striatum in rats, in which microinjecting a GABA agonist or a glutamate antagonist into more anterior regions produces appetitive responses (feeding) and into more posterior regions produces aversive responses (paw treading, burying) (Reynolds and Berridge, 2001, 2002, 2003). These studies are characteristic of a growing body of evidence pointing to a role of the ventral striatum in aversive motivation and with distinct neuronal responses associated with appetitive and aversive events (Ikemoto and Panksepp, 1999; Horvitz, 2000; Schoenbaum and Setlow, 2003; Setlow et al., 2003; Jensen et al., 2003; Seymour et al., 2004, 2005; Roitman et al., 2005; Wilson and Bowman, 2005).

Aversive learning is well recognized to involve the amygdala.

Interestingly, a recent gambling study involving mixed gains and losses of money, at differing amounts and probabilities, identified loss prediction errors in the amygdala, but only gain-related prediction error in the striatum (Yacubian et al., 2006). Although it is difficult to place too much emphasis on the respective negative findings for this and our study, it is noteworthy that these two areas are richly interconnected, both directly and indirectly (Russchen et al., 1985).

The anatomical separation within the striatum could well be accompanied by a separation in terms of the relevant neuro-modulators (Daw et al., 2002; Doya, 2002). A substantial body of data points to the role of dopamine in striatal reward-related activity (Everitt et al., 1999; Montague et al., 2004; Wise, 2004), specifically relating to the representation of prediction errors that guide learning in Pavlovian and instrumental learning tasks (Montague et al., 1996; Schultz et al., 1997). Furthermore, dopamine has been observed to modulate striatal reward prediction errors in human monetary gambling tasks selectively (Pessiglione et al., 2006). If dopamine is involved in the appetitive prediction error observed here, this raises the question as to the nature of the aversive prediction error signal, given previous observations and current controversies concerning dopaminergic involvement in aversive behaviors (Ikemoto and Panksepp, 1999; Horvitz, 2000; Ungless, 2004). One possibility is that serotonin released from the dorsal raphe nucleus plays this role (Daw et al., 2002). Consistent with this hypothesis, there is evidence of a serotonin-dopamine gradient along a caudal-rostral axis in the striatum (Heidbreder et al., 1999; Brown and Molliver, 2000). However, because our study was not pharmacological, we cannot rule out the possibility that instead of there being a separate, nondopaminergic opponent, dopamine provides a valence-independent signal that interacts with valence-specific activity intrinsically coded in the striatum (Seymour et al., 2005).

From the perspective of studies into financial decision making and prediction, it is noteworthy that we see striatal BOLD signals above the baseline associated with prediction errors for financial losses, whereas most previous imaging studies involving positive and negative financial returns show only decreases below baseline for unexpected losses. This result is important because it makes the findings for financial losses consonant with those for primary aversive outcomes such as pain. It also reinforces caution in the interpretation of striatal activity in human decision-making tasks, which as noted in the past (Poldrack, 2006), are sometimes prone to the reverse inference that striatal activity implies the operation of reward mechanisms.

One possible reason for the difference between our results and previous results is that in experimental monetary decision-making tasks, subjects make choices under the reasonable expectation (perhaps based on implicit knowledge of the mores of ethical committees) of a net financial gain. This establishes an appetitive context or frame within which all outcomes are judged. In contrast, most decisions in day-to-day life involve risks that span positive and negative outcomes; we hoped that mixed-outcome prediction, with no opportunity for choice, would avoid such a frame. Empirical work in finance and economics has suggested that such mixed-outcome decisions fit rather awkwardly within the descriptive framework usually applied to decisions that involve pure gains and losses. Constructs such as Prospect theory suggest a strong dependence of decision making on valence context (positive or negative) in which options are judged (Levy and Levy, 2002). The absence of a positive orientation for loss prediction errors in previous studies may thus have arisen from such positive frames. Our results hint that more naturalistic

human studies that involve genuine risk of financial loss may be critical to gain further insights into the role of the striatum and other structures in the judgment and integration of gains and losses.

References

- Barto AG (1995) Adaptive critic and the basal ganglia. In: Models of information processing in the basal ganglia (Houk JC, Davis JL, Beiser DG, eds), pp 215–232. Cambridge: MIT.
- Bayer HM, Glimcher PW (2005) Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* 47:129–141.
- Becerra L, Breiter HC, Wise R, Gonzalez RG, Borsook D (2001) Reward circuitry activation by noxious thermal stimuli. *Neuron* 32:927–946.
- Bitsios P, Szabadi E, Bradshaw CM (2004) The fear-inhibited light reflex: importance of the anticipation of an aversive event. *Int J Psychophysiol* 52:87–95.
- Breiter HC, Aharon I, Kahneman D, Dale A, Shizgal P (2001) Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron* 30:619–639.
- Brown P, Molliver ME (2000) Dual serotonin (5-HT) projections to the nucleus accumbens core and shell: relation of the 5-HT transporter to amphetamine-induced neurotoxicity. *J Neurosci* 20:1952–1963.
- Buchel C, Holmes AP, Rees G, Friston KJ (1998) Characterizing stimulus-response functions using nonlinear regressors in parametric fMRI experiments. *NeuroImage* 8:140–148.
- Daw ND, Kakade S, Dayan P (2002) Opponent interactions between serotonin and dopamine. *Neural Netw* 15:603–616.
- Delgado MR, Nystrom LE, Fissell C, Noll DC, Fiez JA (2000) Tracking the hemodynamic responses to reward and punishment in the striatum. *J Neurophysiol* 84:3072–3077.
- Dickinson A, Dearing MF (1979) Appetitive-aversive interactions and inhibitory processes. In: Mechanisms of learning and motivation (Dickinson A, Boakes RA, eds), pp 203–231. Hillsdale, NJ: Erlbaum.
- Doya K (2002) Metalearning and neuromodulation. *Neural Netw* 15:495–506.
- Elliott R, Newman JL, Longe OA, Deakin JF (2003) Differential response patterns in the striatum and orbitofrontal cortex to financial reward in humans: a parametric functional magnetic resonance imaging study. *J Neurosci* 23:303–307.
- Everitt BJ, Parkinson JA, Olmstead MC, Arroyo M, Robbins TW (1999) Associative processes in addiction and reward. The role of amygdala-ventral striatal subsystems. *Ann N Y Acad Sci* 877:412–438.
- Fiorillo CD, Tobler PN, Schultz W (2003) Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299:1898–1902.
- Gray JA (1991) The psychology of fear and stress, Ed 2. Cambridge, UK: Cambridge UP.
- Haruno M, Kuroda T, Doya K, Toyama K, Kimura M, Samejima K, Imamizu H, Kawato M (2004) A neural correlate of reward-based behavioral learning in caudate nucleus: a functional magnetic resonance imaging study of a stochastic decision task. *J Neurosci* 24:1660–1665.
- Heidbreder CA, Hedou G, Feldon J (1999) Behavioral neurochemistry reveals a new functional dichotomy in the shell subregion of the nucleus accumbens. *Prog Neuropsychopharmacol Biol Psychiatry* 23:99–132.
- Horvitz JC (2000) Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neuroscience* 96:651–656.
- Ikemoto S, Panksepp J (1999) The role of nucleus accumbens dopamine in motivated behavior: a unifying interpretation with special reference to reward-seeking. *Brain Res Brain Res Rev* 31:6–41.
- Jensen J, McIntosh AR, Crawley AP, Mikulis DJ, Remington G, Kapur S (2003) Direct activation of the ventral striatum in anticipation of aversive stimuli. *Neuron* 40:1251–1257.
- Knutson B, Westdorp A, Kaiser E, Hommer D (2000) fMRI visualization of brain activity during a monetary incentive delay task. *NeuroImage* 12:20–27.
- Knutson B, Rick S, Wimmer GE, Prelec D, Loewenstein G (2007) Neural predictors of purchases. *Neuron* 53:147–156.
- Konorski J. (1967) Integrative activity of the brain: an interdisciplinary approach. Chicago: University of Chicago.
- Levy M, Levy H (2002) Prospect theory: much ado about nothing? *Manage Sci* 48:1334–1349.
- Montague PR, Dayan P, Sejnowski TJ (1996) A framework for mesence-

- phalic dopamine systems based on predictive Hebbian learning. *J Neurosci* 16:1936–1947.
- Montague PR, Hyman SE, Cohen JD (2004) Computational roles for dopamine in behavioural control. *Nature* 431:760–767.
- Morris G, Nevet A, Arkadir D, Vaadia E, Bergman H (2006) Midbrain dopamine neurons encode decisions for future action. *Nat Neurosci* 9:1057–1063.
- Nieuwenhuis S, Heslenfeld DJ, von Geusau NJ, Mars RB, Holroyd CB, Yeung N (2005) Activity in human reward-sensitive brain areas is strongly context dependent. *NeuroImage* 25:1302–1309.
- Niv Y, Duff MO, Dayan P (2005) Dopamine, uncertainty and TD learning. *Behav Brain Funct* 1:6.
- 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.
- Pagnoni G, Zink CF, Montague PR, Berns GS (2002) Activity in human ventral striatum locked to errors of reward prediction. *Nat Neurosci* 5:97–98.
- 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.
- Poldrack RA (2006) Can cognitive processes be inferred from neuroimaging data? *Trends Cogn Sci* 10:59–63.
- Reynolds SM, Berridge KC (2001) Fear and feeding in the nucleus accumbens shell: rostrocaudal segregation of GABA-elicited defensive behavior versus eating behavior. *J Neurosci* 21:3261–3270.
- Reynolds SM, Berridge KC (2002) Positive and negative motivation in nucleus accumbens shell: bivalent rostrocaudal gradients for GABA-elicited eating, taste “liking”/“disliking” reactions, place preference/avoidance, and fear. *J Neurosci* 22:7308–7320.
- Reynolds SM, Berridge KC (2003) Glutamate motivational ensembles in nucleus accumbens: rostrocaudal shell gradients of fear and feeding. *Eur J Neurosci* 17:2187–2200.
- Roitman MF, Wheeler RA, Carelli RM (2005) Nucleus accumbens neurons are innately tuned for rewarding and aversive taste stimuli, encode their predictors, and are linked to motor output. *Neuron* 45:587–597.
- Russchen FT, Bakst I, Amaral DG, Price JL (1985) The amygdalostriatal projections in the monkey. An anterograde tracing study. *Brain Res* 329:241–257.
- Schoenbaum G, Setlow B (2003) Lesions of nucleus accumbens disrupt learning about aversive outcomes. *J Neurosci* 23:9833–9841.
- Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. *Science* 275:1593–1599.
- Setlow B, Schoenbaum G, Gallagher M (2003) Neural encoding in ventral striatum during olfactory discrimination learning. *Neuron* 38:625–636.
- 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.
- Seymour B, O'Doherty JP, Koltzenburg M, Wiech K, Frackowiak R, Friston K, Dolan R (2005) Opponent appetitive-aversive neural processes underlie predictive learning of pain relief. *Nat Neurosci* 8:1234–1240.
- Sutton RS, Barto AG (1981) Toward a modern theory of adaptive networks: expectation and prediction. *Psychol Rev* 88:135–170.
- 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.
- Tanaka SC, Doya K, Okada G, Ueda K, Okamoto Y, Yamawaki S (2004) Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops. *Nat Neurosci* 7:887–893.
- Tanaka SC, Samejima K, Okada G, Ueda K, Okamoto Y, Yamawaki S, Doya K (2006) Brain mechanism of reward prediction under predictable and unpredictable environmental dynamics. *Neural Netw* 19:1233–1241.
- Tom SM, Fox CR, Trepel C, Poldrack RA (2007) The neural basis of loss aversion in decision-making under risk. *Science* 315:515–518.
- Ungless MA (2004) Dopamine: the salient issue. *Trends Neurosci* 27:702–706.
- Wilson DI, Bowman EM (2005) Rat nucleus accumbens neurons predominantly respond to the outcome-related properties of conditioned stimuli rather than their behavioral-switching properties. *J Neurophysiol* 94:49–61.
- Wise RA (2004) Dopamine, learning and motivation. *Nat Rev Neurosci* 5:483–494.
- 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.
- Zink CF, Pagnoni G, Martin ME, Dhamala M, Berns GS (2003) Human striatal response to salient nonrewarding stimuli. *J Neurosci* 23:8092–8097. Vol. 27, No. 18