

Temporal Prediction Errors in a Passive Learning Task Activate Human Striatum

Samuel M. McClure,¹ Gregory S. Berns,²
and P. Read Montague^{1,*}

¹Human Neuroimaging Lab
Center for Theoretical Neuroscience
Division of Neuroscience
Baylor College of Medicine
Houston, Texas 77030

²Department of Psychiatry and Behavioral Sciences
Emory University School of Medicine
Atlanta, Georgia 30322

Summary

Functional MRI experiments in human subjects strongly suggest that the striatum participates in processing information about the predictability of rewarding stimuli. However, stimuli can be unpredictable in character (what stimulus arrives next), unpredictable in time (when the stimulus arrives), and unpredictable in amount (how much arrives). These variables have not been dissociated in previous imaging work in humans, thus conflating possible interpretations of the kinds of expectation errors driving the measured brain responses. Using a passive conditioning task and fMRI in human subjects, we show that positive and negative prediction errors in reward delivery time correlate with BOLD changes in human striatum, with the strongest activation lateralized to the left putamen. For the negative prediction error, the brain response was elicited by expectations only and not by stimuli presented directly; that is, we measured the brain response to nothing delivered (juice expected but not delivered) contrasted with nothing delivered (nothing expected).

Introduction

As a mobile organism navigates through the world, sensory data flow into a vast array of parallel systems. Changes in the predictability of these data streams act as important markers, flagging epochs during which attention should be redirected or learning should occur. Indeed, formal learning theory has elevated such observations into a set of prescriptions that describe how animals learn when expectations about the world are violated (Dickinson, 1981; Mackintosh, 1983). These findings in behavioral learning are paralleled by work showing how specific neural systems respond to similar violations of expectations about reward delivery. For example, midbrain dopamine neurons, whose activity is implicated in reward processing, give transient responses to deviations in expectations about rewarding stimuli (Schultz, 1998). This fact can be captured in the hypothesis that dopamine neurons encode a *prediction error* in the time and amount of reward delivery (Montague et al., 1996; Schultz et al., 1997). Although this function of midbrain dopaminergic systems is not suffi-

cient to capture all responses in these neurons (Berridge and Robinson, 1998), it does provide a quantitative basis for the design and interpretation of fMRI experiments in human subjects where reward expectations can be parametrically manipulated.

It is known that changes in the predictability of sequential gustatory stimuli cause increased activation in traditional brain reward structures. In Berns et al. (2001), squirts of juice and water were delivered to human subjects in two separate sequences while continuous blood oxygenation level-dependent (BOLD) measurements were acquired. In one sequence, juice and water were delivered in a fashion where both the time and type of the next stimulus was completely predictable. That is, juice and water were alternated, and the time of arrival of the next stimulus was fixed (Figure 1A). In a separate unpredictable sequence, the order of juice and water squirts and the time between squirts were randomized; however, the average time between stimuli remained fixed (Figure 1A). The differential brain response to the predictable and unpredictable sequences revealed significant activity differences in the ventral striatum and ventromedial frontal cortex (Berns et al., 2001).

This experiment leaves open two important issues. First, for the sequential gustatory stimuli, the juice and water played two roles: (1) stimuli that predict reward and (2) stimuli that are rewards. This issue is easily rectified by using a neutral stimulus (like a light) to act as the predictor of reward (juice or water). Second, this experiment changed two sources of predictability concurrently: (1) the time (when) and (2) type (what) of the next stimulus. This issue is rectified by separating the temporal prediction errors from the stimulus prediction errors (errors in expectations of what comes next). Fortunately, standard instrumental and Pavlovian conditioning paradigms are sufficient to address these two problems.

Pagnoni et al. (2002) carried out an instrumental task, and a region of interest (ROI) analysis revealed that the temporal prediction error in reward delivery was locked to increased activation in the nucleus accumbens. The literature on instrumental and Pavlovian conditioning is vast; however, there is compelling evidence that distinct processes mediate these different forms of learning (Berridge, 2000; Dickinson and Balleine, 2002). This situation suggested that different neural structures would be involved in instrumental and Pavlovian assays of temporal prediction errors in reward delivery times. Accordingly, we performed a separate experiment looking for brain structures with activity changes correlated with temporal prediction errors when no actions are required.

In this paper, we employed a simple classical conditioning paradigm in human subjects in which a light predicted the time of reward delivery. During training (*normal events*; Figure 1B), the light consistently preceded delivery of a juice squirt by 6 s. The amount of juice remained constant (0.8 ml). After 50 such pairings, six *catch events* were randomly inserted in the pairing sequence. In these catch events, the delivery of juice is delayed 4 s beyond the time expected from training

*Correspondence: read@bcm.tmc.edu

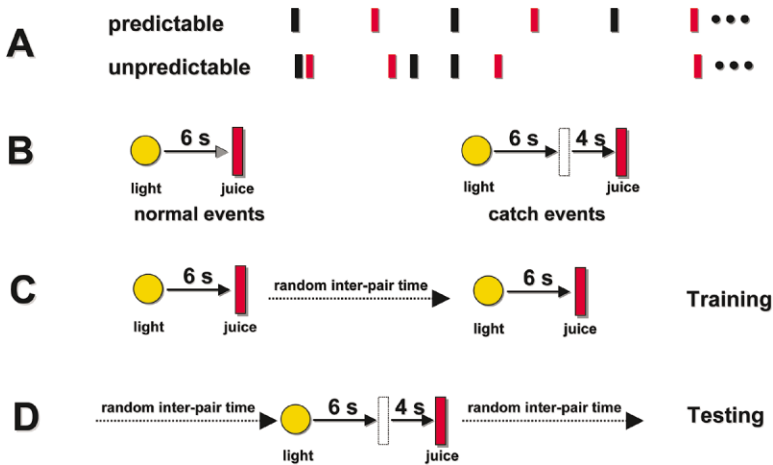


Figure 1. Experimental Design

(A) In a previous experiment, juice and water were delivered to subjects in two separate (predictable and unpredictable) sequences. Stimuli delivered during the unpredictable sequence were associated with greater changes in brain activity in the ventral striatum compared with stimuli delivered during the predictable sequence.

(B) However, stimuli can be unpredictable in character (*what* stimulus arrives next), unpredictable in time (*when* the stimulus arrives), and unpredictable in amount (*how much* arrives). We sought to separate the effects of temporal prediction errors only. Subjects were trained to expect juice at a fixed time following a flash of light (normal events) and then changes in brain response were probed when juice was delivered at an unexpected time (catch events).

(C) Normal events consisted of brief (1 s) flashes of a yellow light centered in their visual field and orally delivered fruit juice (in 0.8 ml boluses). The time between individual events was randomly selected from between 4 and 14 s.

(D) After 49 consecutive light-juice pairings, several catch events were randomly inserted among normal events. For the catch events, the time of juice delivery was extended to 10 s beyond the preceding flash of light. A separate group of subjects was given a control experiment, where the yellow cue light predicted another light (red). All other aspects of the experiment were held the same.

(Figure 1D). From single-unit recordings in monkeys, catch events are expected to induce two prediction errors: less reward than predicted at the trained time (negative prediction error) and more reward than predicted at the unexpected delayed time (positive prediction error). To control for effects related solely to the timing of events, we repeated the experiment in a separate group of subjects using a neutral predicted stimulus. In these experiments, the light cue (yellow in color) predicted a different light (red).

Results

In the experiment, the first two scanning runs consisted entirely of normal events in which juice delivery consistently followed the light cue at a 6 s delay. During the third scanning run, several catch events were inserted during which juice delivery was delayed to 10 s following the light cue. Catch events engendered two prediction errors: (1) a positive prediction error for juice delivered at the untrained time relative to the trained time and (2) a negative prediction error for the absence of expected juice delivery at 6 s during catch events versus periods of no juice delivery during normal events. We determined which brain regions showed a correlation with these two prediction errors in turn. Brain regions indicated as significant were composed of five or more contiguous voxels each significant to $p < 0.001$ (see Table 1).

For each of the contrasts we discuss, no regions were significant in the control experiment (cue predicts light) at $p < 0.001$. One possible reason for this may have been that subjects were not as attentive during the control experiment as during the main experiment. However, analysis of BOLD responses to light stimuli in both the main and control experiment showed that primary visual cortex showed a significant random effect across all light presentations ($p < 0.001$). Analysis of variance of the mean percent change in BOLD signal across all visual event types (cue in main experiment, cue and

predicted light in control experiment) showed no significant differences ($F_{3,44} = 0.55$, $p = 0.65$).

Head motion during the experiment was unavoidable; however, motion artifacts did not contribute to any of the results to be discussed. Event-related head movements were calculated for all of our main experimental effects and were not found to be significantly different for any effect. Juice delivery during normal and catch events, for example, were 0.071 ± 0.030 mm and 0.068 ± 0.033 mm, respectively. For events with no juice delivery, there was likewise no significant difference between catch and normal events. Following the absence of juice delivery during catch events, there was an average deviation of 0.068 ± 0.035 mm. For the control period in normal events, head motion was 0.137 ± 0.089 mm.

Positive Prediction Error: Unexpected versus Expected Juice Delivery

For juice delivery during catch and normal events, the sensory input was equivalent. The only difference was the relative predictability of the juice delivery. Unpredicted delivery (catch event) was associated with greater changes in BOLD response in the left putamen ($p < 0.001$) than predicted juice delivery (normal event) (the right putamen significant at $p < 0.005$). There were no other areas of significant activation for this contrast, and no regions showed significantly greater response to predicted delivery compared to unpredicted delivery.

Negative Prediction Error: Absence of Expected Juice Delivery versus Periods of No Juice Delivery during Normal Events

Hemodynamic response functions were fit to the data beginning at 6 s following the light cue during catch events and at 10 s following the cue during normal events. In both of these cases, there were no stimuli delivered to the subject. The only difference between the two conditions was in the subjects' expectations. That is, the absence of juice delivery during the catch

Table 1. Summary of Brain Regions Displaying Differential Activation

Brain Region	MNI Coordinates (x, y, z)	Peak t Value	Cluster Size
Juice delivered (unexpected) > juice delivered (expected) L putamen	-18, 4, 8	4.68	11
Juice delivered (expected) > juice delivered (unexpected) None			
Juice not delivered (expected) > juice not delivered (unexpected) L putamen	-18, 1, 8	4.18	5
Juice not delivered (unexpected) > juice not delivered (expected) None			

All regions are composed of voxels significant with $p < 0.001$.

events was less juice than expected: a negative prediction error. No juice could have been expected at 10 s following the light during normal events, so there should be no prediction error associated with these events. We found significantly less activity following the absence of juice delivery during catch events than during normal events restricted to the left putamen (Figure 3, $p < 0.001$). As was the case for the positive prediction error (Figure 2), the right putamen was significant at a lesser

threshold ($p < 0.005$). No brain regions demonstrated significantly greater changes in brain response during catch events minus normal events.

These findings are further exemplified by a region of interest analysis performed over those voxels in the left putamen found to be significantly affected by the negative prediction error event. As can be seen in Figure 4, catch events evinced a decreased BOLD signal following the time of expected juice delivery and show an

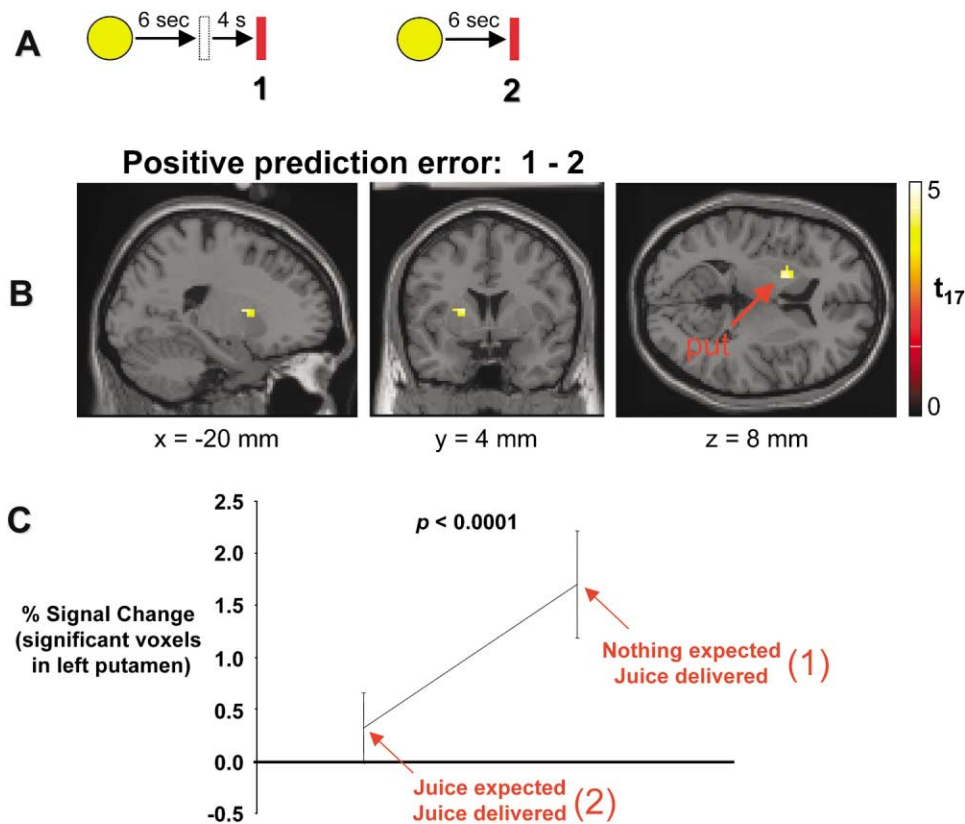


Figure 2. Positive Prediction Error Causes Increased Response in Left Putamen

(A) Comparing the brain response to juice delivered during catch and trained events reveals the effect of predictability on the induced brain response.

(B) Unpredictable juice delivery is associated with significantly greater activity in the left putamen (put) and parts of the left globus pallidus. Regions shown are thresholded at $p < 0.001$ with an extent threshold of 5 voxels. The opposite contrast (predicted juice delivery – unpredicted juice delivery) revealed no significant brain regions.

(C) Average hemodynamic response amplitudes were calculated for each subject for voxels found significant in (A). Combining the average response amplitude across subjects shows that *predicted* juice delivery induces essentially no change in fMRI signal, whereas *unpredicted* juice delivery induces a significantly positive change.

increased BOLD signal following the delivery of juice and the delayed (unexpected) time.

Region of Interest Analysis: Nucleus Accumbens

We found no significant differential activity in the nucleus accumbens for any contrast using regression to a generic hemodynamic response (and a threshold of $p < 0.001$). Based on our previous work (Berns et al., 2001; Pagnoni et al., 2002), we had a strong a priori hypothesis about this region. We therefore undertook a separate region of interest analysis focusing specifically on the nucleus accumbens.

The average impulse response function was derived separately for catch events and normal events. No assumptions were made about the temporal characteristics of evoked hemodynamic responses. As shown in Figure 5, there was a much greater change in BOLD signal following unexpected juice delivery than following expected juice delivery ($p < 0.005$, paired t test over signal at 4, 6, and 8 s following time of unexpected juice delivery as indicated by red points). Fitting each individual's impulse response function with a generic hemodynamic response (HRF) and performing a t test over HRF amplitudes indicated a trend but was not significant ($p = 0.086$). This was true in spite of the fact that the HRF provided an excellent fit to the average impulse response function across subjects ($r^2 = 0.915$). The average impulse response function showed no sign of a negative prediction error signal for time points following the unexpected absence of juice during catch events. The cue was associated with a weakly positive response at 4 s, but this was not significant ($p = 0.13$, paired t test versus BOLD amplitude at 0 s combined over normal and catch events). There was no difference in the average impulse response function between catch and normal events in the control experiment.

Discussion

The use of a classical conditioning paradigm provides a framework for isolating the neural effect of temporal prediction errors in reward delivery. After training subjects to expect juice squirts at a fixed time (6 s) following a light cue, two prediction errors were induced through the presentation of catch events in which juice was delivered at a delayed time relative to the predictive light cue (10 s delay): (1) a *negative prediction error* was caused by the absence of juice at the expected time, and (2) a *positive prediction error* resulted from the unexpected delivery of juice at the delayed time. For the negative prediction error, no stimuli were delivered to the subjects in order to produce the effect. This may prove to be an important experimental manipulation for generating responses related to potentially complex stimuli, since there is no required control stimulus.

The use of a classical conditioning paradigm provided a framework for isolating the neural effect of temporal prediction errors in reward delivery. After training subjects to expect juice squirts at a fixed time (6 s) following a light cue, two prediction errors were induced through the presentation of catch events in which juice was delivered at a delayed time relative to the predictive light cue (10 s delay): (1) a negative prediction error

caused by the absence of juice at the expected time and (2) a positive prediction error resulting from the unexpected delivery of juice at the delayed time. For the negative prediction error, the change in BOLD signal was due to subjects' expectations, since no stimuli were delivered to the subjects in order to create the effect. We found that both of these prediction errors correlated with differences in BOLD signal exclusively in the left putamen (with lesser changes in the right putamen). The positive prediction error was associated with increased brain activity in this structure (Figure 2), while the negative prediction error correlated with decreased activity (Figure 3). Neither of these effects remained when a separate light flash was substituted for juice in the experiment (Figure 5B). This indicated that the changes in brain activity were not due to the temporal arrangement of the stimuli alone, but some property of the juice was required for the effect. One possibility is that the juice events were behaviorally salient, i.e., the subjects knew they would have to swallow, whereas the control experiment required no response, and striatal responses occur only to prediction errors about salient events (Horvitz, 2000).

The fact that only one brain region showed significantly different responses that correlated with prediction errors is likely due to two main causes. First, the experiment was very simple. There were not many different ways for subjects to perform during the task since nothing was required except to watch for a yellow dot and to swallow juice. Second, the limited number of significant brain regions may be due to deviations from the assumptions inherent in our analysis. In particular, the whole-brain analysis we performed assumed that events induced a change in BOLD signal of a specific form and that these responses summed linearly. This concern is particularly relevant in this experiment due to the fact that the majority of events exist as part of a compound stimulus (light followed at a consistent time by reward). Methods exist for estimating BOLD signal changes for individual events within a compound (Ollinger et al., 2001a, 2001b). However, these methods require a significant number of events in which events are delivered in isolation (light and subsequent withholding of juice on 25%–40% of trials, Ollinger et al., 2001b). We could not employ these methods in this experiment due to concerns over extinction of the learned light-reward association. The raw data support our current interpretation that we are observing strong brain responses only to temporal prediction errors. The raw hemodynamic responses during the catch and control trials (Figure 4), which are not affected by assumptions about the hemodynamic response, showed significant deviation from baseline only at the time of the negative and positive prediction errors (labeled on Figure 4), and the raw time course showed no response to the initial cue.

Changes in temporal predictability have previously been observed to affect activity in the nucleus accumbens during an operant conditioning task, where the time of juice delivery was changed relative to an action performed by the subject (Pagnoni et al., 2002). This experiment examined the passive case, where no action was required aside from swallowing the juice reward. While in the active case the subjects must learn a *response-stimulus* association, in this experiment a *stimu-*

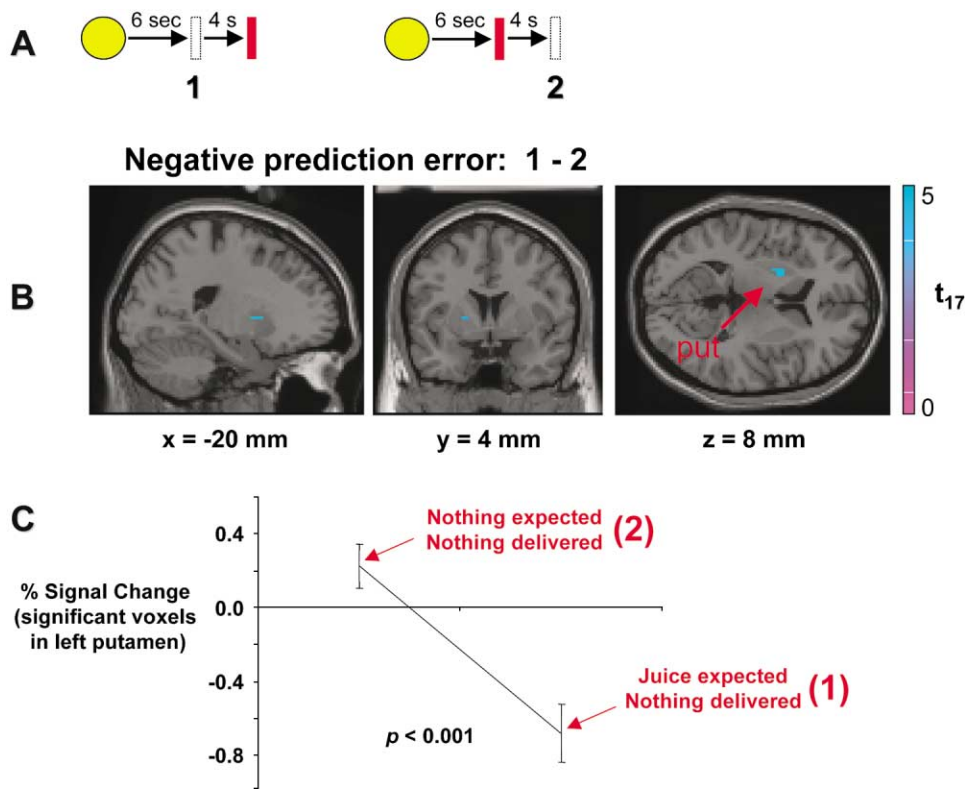


Figure 3. Negative Prediction Error Correlated with Decreased Activity in Left Putamen

After training, juice is expected at 6 s following light flashes.

(A) Comparing the brain response at 6 s following the light in catch events versus a period of no juice delivery in normal events (10 s following light) reveals brain regions that correlate with the negative prediction error.

(B) The failure of juice delivery during catch events correlated with decreased activity selectively in the left putamen. No brain regions showed a greater response to failed expectation in catch events versus the control time during normal events.

(C) Best-fit hemodynamic response amplitudes for the significant voxels demonstrate that there is a significant decrement in BOLD signal following the absence of juice delivery at expected times.

lus-stimulus association must be learned. Interestingly, this experimental difference was enough to cause a different activation pattern. Unpredictability in our study correlated with greatest activity changes in the dorsal striatum, particularly in the left putamen. In Pagnoni et al., the focus was on the ventral striatum, including the nucleus accumbens. We observed significantly elevated

activity changes in catch versus normal events in the nucleus accumbens as well (Figure 5), but the differences were not significant in a regression analysis using a standard hemodynamic response function (Friston et al., 1994).

Until the past few years, activation of reward pathways had been observed using fMRI only during intravenous

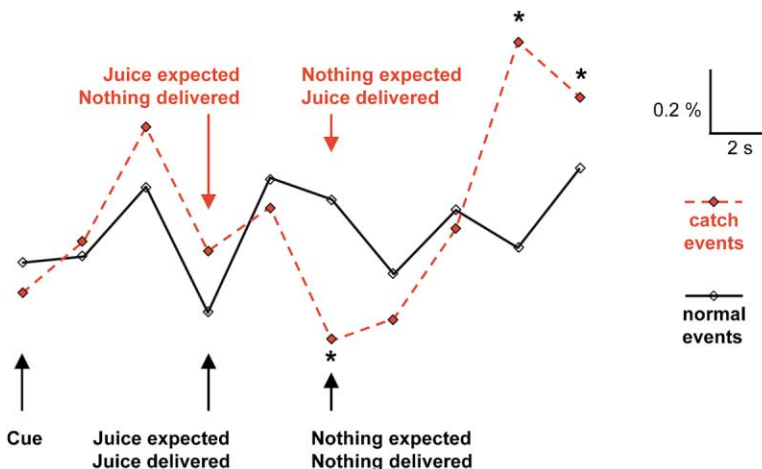


Figure 4. Impulse Response Function in the Left Putamen

The average change in BOLD signal was calculated for normal (solid lines) and catch (dashed lines) events for those voxels found to have a significant brain response to the negative prediction error signal. Paired t tests were performed for all time points. Asterisks (*) indicate those time points which are significantly different ($p < 0.05$) between catch and normal events. The BOLD signal is significantly depressed following the absence of expected juice delivery and significantly elevated following the delivery of juice and the unexpected time during catch events.

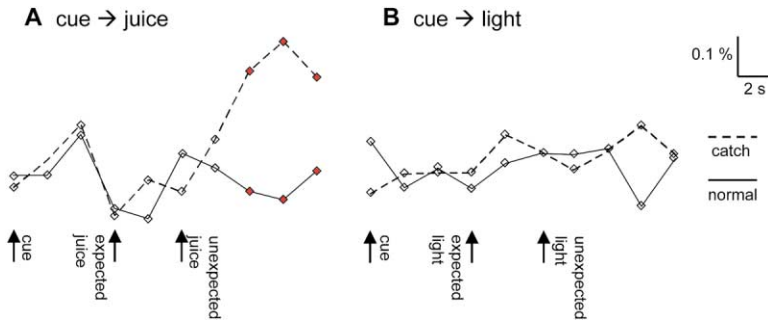


Figure 5. Response in the Nucleus Accumbens—Bilateral

Average impulse response functions were calculated for normal (solid lines) and catch (dashed lined) events.

(A) Unpredictable juice delivery (catch events) was associated with a significantly greater ($p < 0.005$) BOLD signal than juice delivered during normal events (paired t test over points indicated in red).

(B) This effect was not evident in the control experiments where juice was replaced with flashes of differently colored light.

cocaine infusion (Breiter et al., 1997). Since then, a variety of physiological stimuli have been discovered that induce activation in these same neural structures. Our group has found activation following oral fruit juice delivery by looking for correlations with the changes in predictability of the time and character of its delivery (Berns et al., 2001) as well as following instrumental conditioning (Pagnoni et al., 2002). Activation of the caudate, putamen, ventral striatum, orbitofrontal, and ventromedial frontal cortex has also been observed using monetary rewards (Delgado et al., 2000; Elliott et al., 2000; Breiter et al., 2001; Critchley et al., 2001; Knutson et al., 2001; O'Doherty et al., 2001) as well as to more abstract stimuli such as inviting human faces (Kampe et al., 2001; Aharon et al., 2001), emotionally positive words (Hamann and Mao, 2002), and social rewards (Rilling et al., 2002). This suggests that human reward pathways are involved in the processing of rewards independent of modality and thereby may provide a means for comparing the value of disparate stimuli (Montague and Berns, 2002).

As more evidence accumulates on human reward processing using fMRI, we can test the hypothesis that reward pathway activation parallels the prediction error signal seen in animal experiments (Schultz et al., 1997). To date, the evidence is strongly in favor of this hypothesis. We have aimed at testing it directly (Berns et al., 2001; Pagnoni et al., 2002) and found results consistent with nearly equivalent experiments performed in monkeys. O'Doherty and colleagues (2002) used fruit juice in a similarly styled experiment in which a neutral stimulus predicted juice delivery at a variable time delay. They found increased activation in the ventral striatum and midbrain to the predictive stimulus, again consistent with the prediction error hypothesis. A similar response to the conditioned stimulus was observed both in Pagnoni et al. (2002) and in this experiment but did not reach significance in either experiment. Further support for the prediction error hypothesis comes from other groups, including Knutson and colleagues (2000, 2001). Knutson showed that presentation of abstract visual stimuli can be trained to induce scaled activity changes in the nucleus accumbens based on the amount of monetary reward they predict (Knutson et al., 2001). The receipt of the monetary reward itself induced no significant activity changes. This is consistent with the finding that the prediction error signal shifts to the time of the conditioned stimulus with learning.

It is important to note that not all evidence directly supports the idea that activation of classic human reward processing structures is dependent on reward pre-

dictability. Elliott and colleagues found that in a simple decision-making task, activity in the ventral striatum was predicted by the value of current monetary earnings (Elliott et al., 2000). This finding is reminiscent of the idea that dopaminergic projections encode the absolute reward value of the environment (Wise and Rompre, 1989). It is not known in this experiment what subjects' reward predictions were during the experiment. However, given the predominance of differential striatal activity to reward prediction errors and not the reward itself, it is likely that these brain regions are more closely linked to the prediction error and that it is temporally specific.

Experimental Procedures

FMRI Experiment

Twenty-eight normal subjects were studied, including 21 females and 7 males. Subjects varied in age from 22 to 38, averaging 28 years old. Eighteen subjects participated in the main (juice) experiment, and ten were used for the control experiment. All subjects were right-handed and were physically and psychologically normal, as determined through interviews conducted prior to the experiment. Written consent was obtained from all subjects, and the study was approved by the Institutional Review Board at Emory University.

Stimulus Paradigm

Subjects were instructed that they were involved in a task designed to study reward processing. They were told that they would not be required to do anything in the scanner except to watch the visual display and swallow juice as it was delivered.

The task consisted of three scanning sessions of 6 min duration each. During the first two sessions, all light/juice pairings were presented with 6 s light-to-juice time (normal events). The time between individual pairings was randomly selected from between 4 and 14 s (at 2 s increments). This gave 23 pairings during the first scanning runs and 22 pairings during the second scanning run. For the third run, six pairings were randomly selected to be catch events. For these (catch event) pairings, the light-to-juice time was increased to 10 s. In total, there were 49 training pairings, spanning all three runs, given before the first catch event. During the final scanning run, the number of trained (6 s) events outnumbered the number of catch events 2 to 1.

Control Experiment

For the control experiment, the timing of all events was the same as in the main experiment. The single change was that the (yellow) light cue was followed by a (red) light instead of a squirt of juice.

FMRI Acquisition

Scanning was performed on a 1.5 Tesla Philips Intera scanner located at Emory University. Prior to functional scanning, high-resolution T1-weighted structural images ($0.9375 \times 0.9375 \times 5$ mm resolution, no slice gap) were taken of the subjects' brains. Each functional session consisted of 150 whole-brain gradient-echo echo-planar T2*-weighted functional scans acquired using blood oxygenation level-dependent (BOLD) contrast. Volumes were acquired once every 2 s, with each volume consisting of 24 horizontal sections having an in-plane resolution 3.75×3.75 mm (no slice gap, echo time of

40 ms, flip angle of 90°). Horizontal sections were acquired parallel to the anterior commissure-posterior commissure (AC-PC) axis.

Juice Delivery

Juice delivery was accomplished using a computer-controlled syringe pump (Harvard Apparatus, Holliston, MA). Plastic tubes ~10 meters in length were run from the syringe pump in the scanner control room to the subjects' mouths. They were held in the subjects' mouths using plastic mouthpieces. Individual juice squirts were of 1 s duration, delivering a total of 0.8 ml juice. This was just enough to taste and not too much to cause swallowing difficulties while lying prone. Padding and restraints were used to minimize head movements that necessarily occur while swallowing. Image realignment performed after the experiment indicated that head movements remained within an acceptable range (less than 1.5 mm). Following the experiment, all subjects reported enjoying the taste of the juice.

Data Analysis

All data were analyzed using Statistical Parametric Mapping (SPM99; Wellcome Department of Cognitive Neurology, London, UK) (Friston et al., 1995). Prior to analysis, functional images were aligned to correct for head movements using a six-parameter rigid body transformation. These parameters were used to determine whether head movements may account for any of results, by determining the mean head movement associated with each event type. Event-related head movements were calculated as the mean translation during the 2 s following each event relative to the head position during the previous scan. The images were then coregistered with the subjects' anatomical image using a 12 parameter affine transformation and corrected for slice-timing artifacts resulting from nonsimultaneous slice acquisition. Both functional and anatomical images were then normalized to a Montreal Neurological Institute (MNI) template image using a 12 parameter affine transformation followed by nonlinear warping using basis functions (Ashburner and Friston, 1999). To improve the signal-to-noise ratio, an 8 mm (FWHM) isotropic Gaussian kernel was applied to the functional images, which were then band-pass filtered in the temporal domain.

Analysis was performed using an event-related, random effects statistical model. Individual design matrices were generated for each subject, in which each separate experimental effect was modeled as a point stimulus inducing a generic hemodynamic response of unknown amplitude. First-order Taylor series coefficients were included in the model to account for slight discrepancies in juice delivery time and noninstantaneous duration (since juice squirts were of 1 s duration and were delayed approximately half a second from the delivery command due to compliance in the plastic delivery tubes). The model was then regressed to the data giving best-fitting (L2) amplitude estimates for each effect. The first two scanning runs were modeled to consist of two separate effects: light flash and juice delivery. For the third scanning run, the independent effects included the light cue, juice delivery during normal events, juice delivery during catch events, the absence of juice delivery at 6 s during catch events, and the absence of juice delivery during normal events (10 s after light). Individual contrast images were created for each contrast of interest (for example, juice delivery during catch events minus juice delivery during normal events) and each subject. Separate one-sample t tests were then performed across equivalent contrast images and results displayed for voxels showing significance at $p < 0.001$ (uncorrected for multiple comparisons).

A separate ROI analysis was performed using the software package AFNI (Cox, 1996). The onset times of the light cue were entered into a vector, and best fitting impulse response functions (IRF) were extracted for periods encompassing entire pairings (using AFNI plug-in 3dDeconvolve). ROI masks were created bilaterally for the nucleus accumbens as described previously (Pagnoni et al., 2002), and IRFs averaged over the entire volume. Paired t tests were then performed across subjects for time points following event onsets in the IRFs.

Acknowledgments

We would like to thank Megan Martin, Giuseppe Pagnoni, and Mike Wiest for their assistance with this experiment. Additionally, we benefited greatly from discussions with Ron Fisher and Peter Dayan.

This work was supported by grants from the Kane Family Foundation (P.R.M.) and the National Institutes for Health RO1 MH52797 (P.R.M.), RO1 DA11723 (P.R.M.), RO1 MH61010 (G.S.B.), and K08 DA00367 (G.S.B.).

Received: August 22, 2002

Revised: February 28, 2003

Accepted: March 7, 2003

Published: April 23, 2003

References

- Aharon, I., Etcoff, N., Ariely, D., Chabris, C.F., O'Connor, E., and Breiter, H.C. (2001). Beautiful faces have variable reward value: fMRI and behavioral evidence. *Neuron* 32, 537–551.
- Ashburner, J., and Friston, K.J. (1999). Nonlinear spatial normalization using basis functions. *Hum. Brain Mapp.* 7, 254–266.
- Berns, G.S., McClure, S.M., Pagnoni, G., and Montague, P.R. (2001). Predictability modulates human brain response to reward. *J. Neurosci.* 21, 2793–2798.
- Berridge, K.C. (2000). Reward learning: Reinforcement, incentives, and expectations. In *The Psychology of Learning and Motivation*, D.L. Medin, ed. (New York: Academic Press), pp. 223–278.
- Berridge, K.C., and Robinson, T.E. (1998). What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? *Brain Res. Brain Res. Rev.* 28, 309–369.
- Breiter, H.C., Gollub, R.L., Weisskoff, R.M., Kennedy, D.N., Makris, N., Berke, R.T., Rosen, B.R., and Hyman, S.E. (1997). Acute effects of cocaine on human brain activity and emotion. *Neuron* 19, 591–611.
- Breiter, H.C., Aharon, I., Kahneman, D., Dale, A., and Shizgal, P. (2001). Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron* 30, 619–639.
- Cox, R.W. (1996). AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Comput. Biomed. Res.* 29, 162–173.
- Critchley, H.D., Mathias, C.J., and Dolan, R.J. (2001). Neural activity in the human brain relating to uncertainty and arousal during anticipation. *Neuron* 29, 537–545.
- Delgado, M.R., Nystrom, L.E., Fissell, C., Noll, D.C., and Fiez, J.A. (2000). Tracking the hemodynamic responses to reward and punishment in the striatum. *J. Neurophysiol.* 84, 3072–3077.
- Dickinson, A. (1981). *Contemporary animal learning theory* (Cambridge, UK: Cambridge University Press).
- Dickinson, A., and Balleine, B.W. (2002). The role of learning in motivation. In *Learning, Motivation and Emotion*, Volume 3 of Steven's Handbook of Experimental Psychology, Third Edition, C.R. Gallistel, ed. (New York: John Wiley & Sons), pp. 26–72.
- Elliott, R., Friston, K.J., and Dolan, R.J. (2000). Dissociable neural responses in human reward systems. *J. Neurosci.* 20, 6159–6165.
- Friston, K.J., Jezzard, P., and Turner, R. (1994). Analysis of fMRI time-series revisited. *Hum. Brain Mapp.* 2, 45–53.
- Friston, K.J., Holmes, A.P., Worsley, K., Poline, J.B., Frith, C.D., and Frackowiak, R.S.J. (1995). Statistical parametric maps in functional brain imaging: a general linear approach. *Hum. Brain Mapp.* 2, 189–210.
- Hamann, S., and Mao, H. (2002). Positive and negative emotional verbal stimuli elicit activity in the left amygdala. *Neuroreport* 13, 15–19.
- Horvitz, J.C. (2000). Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neuroscience* 96, 651–656.
- Kampe, K.K.W., Frith, C.D., Dolan, R.J., and Frith, U. (2001). Reward value of attractiveness and gaze. *Nature* 413, 589.
- Knutson, B., Westdorp, A., Kaiser, E., and Hommer, D. (2000). FMRI visualization of brain activity during a monetary incentive delay task. *Neuroimage* 12, 20–27.
- Knutson, B., Adams, C.M., Fong, G.W., and Hommer, D. (2001). Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *J. Neurosci.* 21, RC159.

- Mackintosh, N.J. (1983). *Conditioning and associative learning* (Oxford, UK: Oxford University Press).
- Montague, P.R., and Berns, G.S. (2002). Neural economics and the biological substrates of valuation. *Neuron* 36, 265–284.
- Montague, P.R., Dayan, P., and Sejnowski, T.J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J. Neurosci.* 16, 1936–1947.
- O’Doherty, J., Kringelbach, M.L., Rolls, E.T., Hornak, J., and Andrews, C. (2001). Abstract reward and punishment representations in the human orbitofrontal cortex. *Nat. Neurosci.* 4, 95–102.
- O’Doherty, J.P., Deichmann, R., Critchley, H.D., and Dolan, R.J. (2002). Neural responses during anticipation of a primary taste reward. *Neuron* 33, 815–826.
- Ollinger, J.M., Shulman, G.L., and Corbetta, M. (2001a). Separating processes within a trial in event-related functional MRI: I. The method. *Neuroimage* 13, 210–217.
- Ollinger, J.M., Corbetta, M., and Shulman, G.L. (2001b). Separating processes within a trial in event-related functional MRI: II. Analysis. *Neuroimage* 13, 218–229.
- Pagnoni, G., Zink, C.F., Montague, P.R., and Berns, G.S. (2002). Activity in human ventral striatum locked to errors of reward prediction. *Nat. Neurosci.* 5, 97–98.
- Rilling, J., Gutman, D., Zeh, T., Pagnoni, G., Berns, G., and Kilts, C. (2002). A neural basis for social cooperation. *Neuron* 35, 395–405.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *J. Neurophysiol.* 80, 1–27.
- Schultz, W., Dayan, P., and Montague, P.R. (1997). A neural substrate of prediction and reward. *Science* 275, 1593–1599.
- Wise, R.A., and Rompre, P.P. (1989). Brain dopamine and reward. *Annu. Rev. Psychol.* 40, 191–225.