Activity in human ventral striatum locked to errors of reward prediction

Giuseppe Pagnoni[1], Caroline F. Zink[1], P. Read Montague[2] and Gregory S. Berns[1]

[1] *Department of Psychiatry and Behavioral Sciences, Emory University School of Medicine, Atlanta, Georgia 30322, USA*

[2] *Division of Neuroscience, Baylor College of Medicine, Houston, Texas 77030, USA*

Correspondence should be addressed to G.P. (gpagnon@emory.edu)

The mesolimbic dopaminergic system has long been known to be involved in the processing of rewarding stimuli[1], although recent evidence from animal research has suggested a more specific role of signaling errors in the prediction of rewards[2,3]. We tested this hypothesis in humans, using functional magnetic resonance imaging (fMRI) and an operant conditioning paradigm for the discrete delivery of small quantities of fruit juice, along with a control experiment in which juice was substituted with a neutral visual stimulus. A local estimation of the activity in the ventral striatum showed a significant differentiation when the juice was withheld at the expected time of delivery; this finding was not replicated in the case of visual stimulation, providing evidence for time-locked processing of reward prediction errors in human ventral striatum.

In a previous fMRI study we found an enhancement of striatal activity when a primary rewarding stimulus (fruit juice or water) was delivered in an unpredictable fashion compared to a predictable one[4]. The goal of the present study was to determine whether this enhancement was due to errors of reward prediction in accordance with theoretical models of mesolimbic activity[3].

Two versions of an operant conditioning task were used in 32 normal subjects who were undergoing event-related fMRI. All the subjects gave informed consent for a protocol approved by the Emory University Institutional Review Board. In experiment 1 ($n = 17$, age = 27.5 ± 6.7, mean ± s.d.), a green disc appeared on a screen for 1 second, indicating that fruit juice was available; subjects were instructed to press a key on a button box when they wanted to receive the juice. After the subject's response, a small squirt of juice (0.6 ml) was delivered to the subject's mouth by a plastic tube connected to a computer-controlled syringe pump. The possibility that any observed response might be related to the predictability of a salient stimulus[5], not necessarily a rewarding one, motivated us to carry out experiment 2 with a second

group of subjects ($n = 15$, age = 27.7 ± 6.1), in which the juice was substituted with a simple visual stimulus (a blue disc). Three functional imaging runs sensitive to blood oxygenation level–dependent (BOLD) contrast were collected for each subject with a 1.5-Tesla Philips (Best, the Netherlands) Intera scanner (T_2*-weighted echo-planar imaging, $TR = 2000$ ms, TE = 40 ms, 64×64 matrix, 24 5-mm axial slices, 150 scans per run), following a T_1-weighted acquisition protocol for anatomical reference (256×256 matrix). To establish the conditioning between the button press and the reward, the delay between the subject's response and the juice delivery was fixed at 4 seconds during the first two functional runs. In the third run, two types of trials were randomly selected for presentation: "Regular" trials (with probability = 2/3), in which the delay between button press and juice delivery remained fixed at 4 seconds as in the preceding runs, and "Delayed" trials (with probability = 1/3), in which this delay was increased to 8 seconds. This created an event in which the prediction of the juice delivery at 4 seconds would prove wrong—that is, a 'prediction error' (**Fig. 1**). To minimize the induction of secondary conditioning patterns, the interval between the juice delivery and the next visual cue was randomly varied from 4 to 7 seconds. The average number of Regular and Delayed trials per subject in the third run was *a posteriori* found to be 15.5 ± 3.1 and 7.0 ± 2.2, respectively (mean ± s.d.). Because we were interested in the change in the ventral striatum response when a prediction error was introduced, only the fMRI data from the third run were used in the analysis.

The data were preprocessed and analyzed with the software AFNI[6]. The images of the third run were realigned to the first acquired scan and time-corrected for slice acquisition order. A region of interest (ROI) mask was drawn on each subject's nucleus accumbens (NAc) using the registered anatomical and functional images and a reference brain atlas[7]. We standardized the ROI as a square of 3×3 voxels in each hemisphere, in a single axial slice, thus conservatively allocating a volume of $3.75 \times 3.75 \times 5$ mm (70 mm[3]) for the investigated region on each side of the brain (**Fig. 2**). In some subjects, the ROI partially overlapped with the susceptibility artifact in the subcallosal region; to exclude the voxels showing an artifact-related signal drop, a combined threshold/cluster-growing algorithm (AFNI, plug-in Threshold) was applied to the mean of the functional images to compute an effective whole-brain mask. This screened out non-brain voxels and voxels falling within the artifact region. A logical 'and' calculation was then computed between the ROI and this mask, to yield the corrected ROI mask to be used in the analysis. The final ROI size
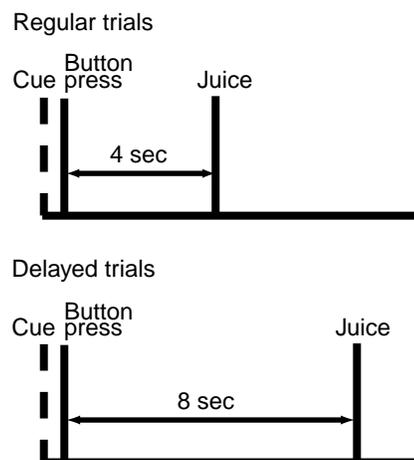

Fig. 1. Experimental design. Two trial types were used (Regular and Delayed). In experiment 1, subjects pressed a button after a visual cue signaled the availability of juice, with an average time distance from the visual cue presentation of 1.5 ± 1.0 s (mean ± s.d.). In the Regular trials, 0.6 ml of fruit juice was delivered orally 4 s after the button press. After behavioral conditioning (two functional runs of Regular-only trials), one-third of the trials were substituted with a Delayed trial in which the juice was delivered 8 s after the button press. This resulted in the occurrence of a prediction error at 4 s, the expected time of juice delivery. In experiment 2 the delivery of juice was substituted with the presentation of a neutral visual stimulus (a blue disc).
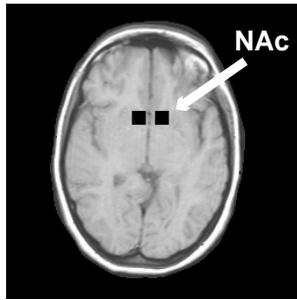
Fig. 2. The location of the region of interest for the nucleus accumbens (NAc) for one subject, as outlined on the subject's T_1-weighted image.

thus varied across subjects, with an average value of 16.9 ± 2.1 voxels (mean ± s.d.) and a range of 10 to 18 voxels.

For each voxel in the ROI, a general linear model (GLM) estimation was carried out on a subject-wise basis. The conditioned stimulus (button press) onsets were taken as the reference starting points, separately for the Regular and the Delayed trials. The time window for the estimation was 18 seconds, allowing 10 seconds after the last stimulus (the delayed juice, occurring 8 seconds after the button press) for the signal to return to the baseline. Because the spacing between the delivery of juice and the subsequent cue ranged from 4 to 7 seconds, this resulted in overlapping BOLD responses from consecutive trials; therefore, a simple selective averaging to yield a peri-stimulus trial-specific time course would not have been appropriate. Instead, each trial type was modeled as a sum of shifted and scaled copies of the delta function representing the reference starting points of the trials (the button presses), with time lags extending from 0 to 18 seconds. These covariates, along with a baseline model consisting of a constant and a linear trend, were entered in a GLM that was estimated in a voxel-wise fashion for each subject. This procedure permitted a deconvolution of overlapping BOLD responses without requiring a predetermined shape for the hemodynamic response[8]. The estimated BOLD responses for each voxel in the ROI were normalized to represent the percentage signal change with respect to the estimated baseline, and averaged across the ROI for each subject. These subject-level estimates for the Delayed and the Regular responses were entered into a two-way repeated-measures ANOVA with trial type (Regular and Delayed) and time as within-subject factors, followed by *post-hoc* comparison of the two trial type effects at each time point.

In experiment 1, there was a significant effect of trial type ($F_{1,16} = 5.532$, $p < 0.0318$). The interaction between trial type and time was not statistically significant; however, this was not surprising, because the experimental design used a differential event only at the expected time of juice delivery. This made a putative interaction effect fairly small when assessed over the entire trial length. However, *post-hoc* examination of the simple main effects at each time point did reveal significant differences between the Regular and the Delayed trials at 10 and 12 seconds after the button press ($p = 0.0036$ and $p = 0.0489$, respectively) (**Fig. 3**); there was also a tendency towards significance for the 8-second time point ($p = 0.0602$). To control for the potentially confounding effect of a differential residual movement of the subject's head in the two types of trials, we repeated the GLM analysis with the six motion parameters estimated in the realignment of the functional images included. The results were similar to those obtained earlier: a significant effect of trial type ($F_{1,16} = 4.512$, $p = 0.0496$) and no point-wise difference between the Delayed and the Regular response surviving the $p < 0.05$ threshold except for the point at 10 seconds after the button press ($p = 0.0116$), with a tendency towards significance for the 8-second time point ($p = 0.0888$). The same analysis performed on the data from experiment 2 did not yield any significant effect.
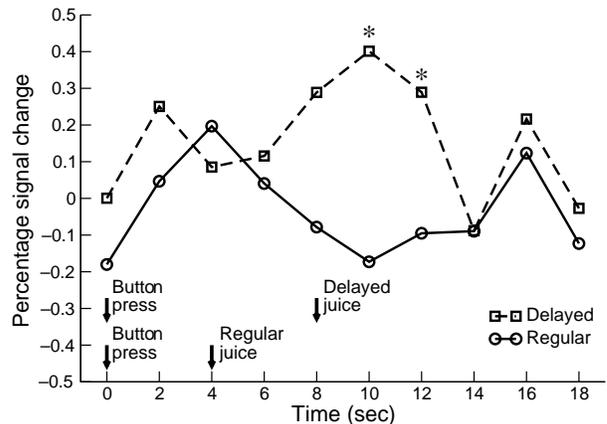

Fig. 3. Average BOLD responses across subjects for the Regular and Delayed trials. A statistically significant difference between the two curves (marked with an asterisk in the figure) was found only at 10 and 12 s after the button press ($p = 0.0036$ and $p = 0.0489$, respectively). The two curves diverge significantly only at the time corresponding to the prediction error (assuming an hemodynamic delay of 6–8 s).

Considering a typical hemodynamic delay of 6–8 seconds[9], the position of the latter points corresponded to the missing delivery of the expected reward 4 seconds after the button press. That is, the NAc responses to the Regular and Delayed trials diverged precisely at the point in time when the conditioned reward should have been received. This effect did not extend to the case of a neutral visual stimulation.

The finding of a prediction-error computation performed in the NAc in humans is consistent with a large body of primate physiological data and has considerable consequences for understanding the dynamics and plasticity of motivated behavior[10]. Although the involvement of ventral striatum in reward mechanisms has recently been investigated in humans using functional brain-mapping techniques[11–13], to our knowledge this is the first reported evidence of time-locked prediction-error coding in human ventral striatum during the processing of rewarding stimuli.

1. Spanagel, R. & Weiss, F. *Trends Neurosci.* **22**, 521–527 (1999).
2. Ljungberg, T., Apicella, P. & Schultz, W. *J. Neurophysiol.* **67**, 145–163 (1992).
3. Schultz, W., Dayan, P. & Montague, P. R. *Science* **275**, 1593–1599 (1997).
4. Berns, G. S., McClure, S. M., Pagnoni, G. & Montague, P. R. *J. Neurosci.* **21**, 2793–2798 (2001).
5. Horvitz, J. C. *Neuroscience* **96**, 651–656 (2000).
6. Cox, R. W. *Comput. Biomed. Res.* **29**, 162–173 (1996).
7. Duvernoy, H. M. *The Human Brain* (Springer-Verlag, Vienna, 1999).
8. Ward, B. D. in *AFNI 3dDeconvolve Documentation* (Medical College of Wisconsin, 2001).
9. Cohen, M. S. *Neuroimage* **6**, 93–103 (1997).
10. Ikemoto, S. & Panksepp, J. *Brain Res. Brain. Res. Rev.* **31**, 6–41 (1999).
11. Elliott, R., Friston, K. J. & Dolan, R. J. *J. Neurosci.* **20**, 6159–6165 (2000).
12. Delgado, M. R., Nystrom, L. E., Fissell, C., Noll, D. C. & Fiez, J. A. *J. Neurophysiol.* **84**, 3072–3077 (2000).
13. Breiter, H. C. & Rosen, B. R. *Ann. NY Acad. Sci.* **877**, 523–547 (1999).