European Journal of Neuroscience, Vol. 18, pp. 2069–2081, 2003

Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task

Jonathan D. Wallis and Earl K. Miller
The Picower Center for Learning and Memory, RIKEN-MIT Neuroscience Research Center and Department of Brain and Cognitive Sciences, E25-236, 45 Carleton Street, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

Keywords: dorsolateral, neurophysiology, orbitofrontal, prefrontal, reward, rhesus monkey

Abstract

An important function of the prefrontal cortex (PFC) is the control of goal-directed behaviour. This requires information as to whether actions were successful in obtaining desired outcomes such as rewards. While lesion studies implicate a particular PFC region, the orbitofrontal cortex (OFC), in reward processing, neurons encoding reward have been reported in both the OFC and the dorsolateral prefrontal cortex (DLPFC). To compare and contrast their roles, we recorded simultaneously from both areas while two rhesus monkeys (*Macaca mulatta*) performed a reward preference task. The monkeys had to choose between pictures associated with different amounts of a juice reward. Neuronal activity in both areas reflected the reward amount. However, neurons in the DLPFC encoded both the reward amount and the monkeys' forthcoming response, while neurons in the OFC more often encoded the reward amount alone. Further, reward selectivity arose more rapidly in the OFC than the DLPFC. These results are consistent with reward information entering the PFC via the OFC, where it is passed to the DLPFC and used to control behaviour.

Introduction

The actions of patients with damage to the prefrontal cortex (PFC) so often seem at odds with their goals that the prefrontal syndrome has been characterized as 'goal neglect' (Duncan *et al*., 1996). As an overarching goal for all organisms is to obtain rewards and avoid punishments, it makes sense that neurons encoding reinforcers are found throughout the PFC (Niki & Watanabe, 1979; Rosenkilde *et al*., 1981; Thorpe *et al*., 1983; Ono *et al*., 1984; Watanabe, 1996; Leon & Shadlen, 1999; Tremblay & Schultz, 1999).

One of the more enduring ideas regarding the functional organization of the PFC is that reward and affect information is more heavily processed in the orbitofrontal cortex (OFC) than other PFC regions (Goldman-Rakic, 1987; Gaffan & Murray, 1990; Baylis & Gaffan, 1991; Gaffan *et al*., 1993; Passingham, 1993; Damasio, 1994; Fuster, 1997; Baxter *et al*., 2000; Roberts & Wallis, 2000; Schultz, 2000; Wallis *et al*., 2001). The OFC is more directly interconnected with limbic structures than the lateral PFC, particularly the dorsolateral prefrontal cortex (DLPFC). The DLPFC is instead more directly interconnected with sensory and motor system structures and is thought to be more involved in cognitive functions such as attention, working memory and response selection than in dealing with the external world (Goldman-Rakic, 1987; Passingham, 1993; Fuster, 1997). For example, humans with DLPFC damage were impaired on a spatial working memory task, but showed normal performance on a 'gambling task' which required assessing different reward schedules; the opposite pattern of deficits was seen in patients with OFC damage (Bechara *et al*., 1998). A double dissociation has also been observed in monkeys: DLPFC lesions impair the ability to shift between rules, but not alternation of choices between rewarded vs. nonrewarded stimuli, while OFC-lesioned monkeys show the opposite pattern of deficits (Dias *et al*., 1996; Wallis *et al*., 2001).

Thus far neurophysiological studies have not found much in the way of such functional dissociations between these prefrontal regions; neurons encoding rewards are readily apparent in both the OFC (Thorpe *et al*., 1983; Tremblay & Schultz, 1999; Kawasaki *et al*., 2001) and the DLPFC (Watanabe, 1996; Leon & Shadlen, 1999; Kobayashi *et al*., 2002). One possible explanation for this discrepancy is that their neuronal properties have not been directly compared within the same study. Thus, their respective properties have to be inferred across different animals with different training histories performing different tasks. This can obscure differences and confound any comparison between them.

We therefore recorded neuronal activity from multiple electrodes simultaneously implanted in the DLPFC and OFC while monkeys performed a task in which they learned to associate novel pictures with different reward values. Two pictures were successively presented on alternate sides of a central fixation point, after which monkeys made a saccade to the remembered location of one or the other (Fig. 1A). Each picture was associated with different amounts of juice reward and the association was learned by trial and error. Once the monkey was consistently choosing the location of the picture associated with the larger reward, the contingencies were reversed and relearned (Fig. 1B). This allowed us to dissociate neuronal activity related to the picture from that related to the reward. We also ensured that the saccade direction (to the left or right) was balanced across conditions, allowing us to dissociate reward activity from that related to the monkey's action.

Correspondence: Dr Jonathan Wallis, as above.
E-mail: wallis@ai.mit.edu

Received 20 May 2003, revised 2 July 2003, accepted 25 July 2003

doi:10.1046/j.1460-9568.2003.02922.x

Materials and methods

Subjects and physiological procedures

The subjects were two adult female rhesus monkeys (*Macaca mulatta*) weighing 5.0–6.0 kg. All procedures were in accord with National Institutes of Health guidelines and the recommendations of the MIT Animal Care and Use Committee. The monkeys were anaesthetised [ketamine induction, 10 mg/kg i.m. with anaesthetic maintenance Isofluorane (1–3%) in balance oxygen] and fitted with a head restraint to ensure that their head could be immobilized during recording. A recording chamber was attached to the monkey's skull and positioned at an angle such that electrodes could be positioned in both the DLPFC and OFC (Fig. 2A). The anterior–posterior location of the centre of the recording chamber was +27 mm relative to the interaural line in monkey A, and +29 mm in monkey B. The lateral–medial position was centred on the principal sulcus. The recording chamber was horizontal in the anterior–posterior axis, and in the lateral–medial axis the chamber was positioned at an angle of 30° from the vertical in monkey A and 25° in monkey B. This arrangement allowed us to simultaneously record from both brain areas by lowering the electrodes to different depths. Recording locations were determined using a 1.5-T magnetic resonance imaging (MRI) scanner. Recordings were made using arrays of eight tungsten dura-puncturing microelectrodes (FHC Instruments, Bowdoin, ME, USA) using a grid (Crist Instruments, Damascus, MD, USA) with 1-mm spacing.

The approximate distance to lower the electrodes was determined from the MRI images, and the electrodes were advanced using custom-built manual microdrives until they were located just above the appropriate cell layer. The electrodes were then slowly lowered into the cell layer until a neuronal waveform was obtained. We then waited 3 h for the brain to settle and thus ensure stability during the recording session, which typically lasted 2 h. Neurons were randomly sampled; no attempt was made to select neurons on the basis of responsiveness. Waveforms were digitized and analysed offline (Plexon Instruments, Dallas, TX, USA). Typically one or two neurons could be discriminated on each electrode; a similar number of neurons per electrode were obtained in the OFC (mean 1.5) and the DLPFC (mean 1.6).

Recording locations were reconstructed by measuring the position of the recording chambers using stereotaxic methods. These were then plotted onto the MRI sections using commercial graphics software (CorelDraw, Ottawa, Canada). The distance of each recording location along the cortical surface from the principal sulcus was then traced and measured. The positions of the other sulci relative to the principal sulcus were also measured in this way, allowing the construction of the unfolded cortical maps shown in Fig. 2B. We confirmed the position of the electrodes in OFC by neurophysiologically mapping the depth of the white matter between the DLPFC and the OFC (a lack of large spikes indicated that the electrode was positioned in the white matter).

Behavioural task

Trials began when the monkey fixated a central fixation point (Fig. 1A). Eye position was monitored throughout the session using an infrared monitoring system (ISCAN, Burlington, MA, USA). Monkeys were required to keep their gaze within 1.5° of a central fixation point. If the monkey's gaze deviated outside this window the trial was immediately terminated; breaks of fixation were not included in the overall error rates. If the monkey maintained fixation for 800 ms

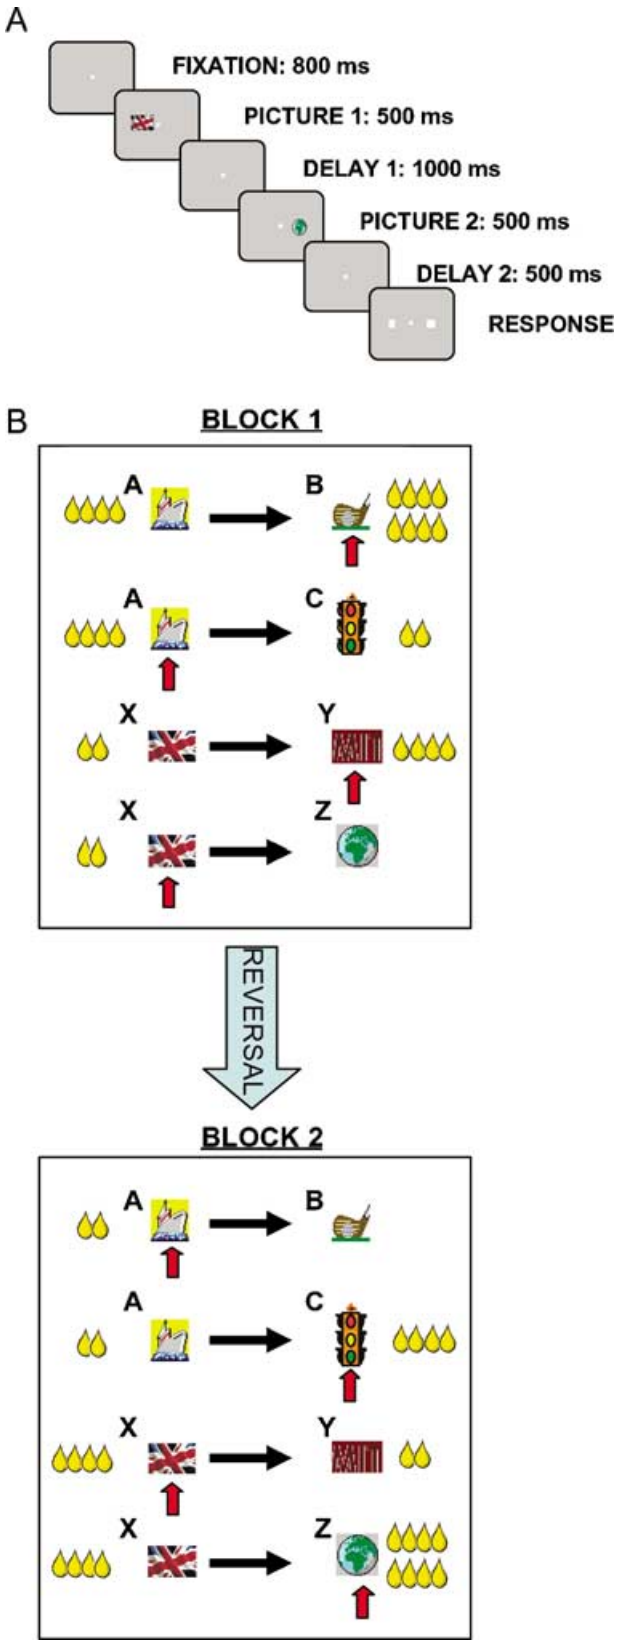

FIG. 1. The behavioural task. (A) Temporal sequence of events in the trial. (B) Stimulus–reward contingencies used in each block of trials. The red arrow indicates the correct response for the monkey, i.e. choosing the picture associated with the larger reward. The current block changed whenever the monkey had chosen the larger reward on 27 of the previous 30 trials. Picture 1 was stimulus A on half the trials and stimulus X on the other half. Pictures A and X appeared equally often on the left and right of the screen. Stimulus A was followed by either stimulus B or C as picture 2, while stimulus X was followed by either stimulus Y or Z. Picture 2 always appeared on the opposite side to picture 1.

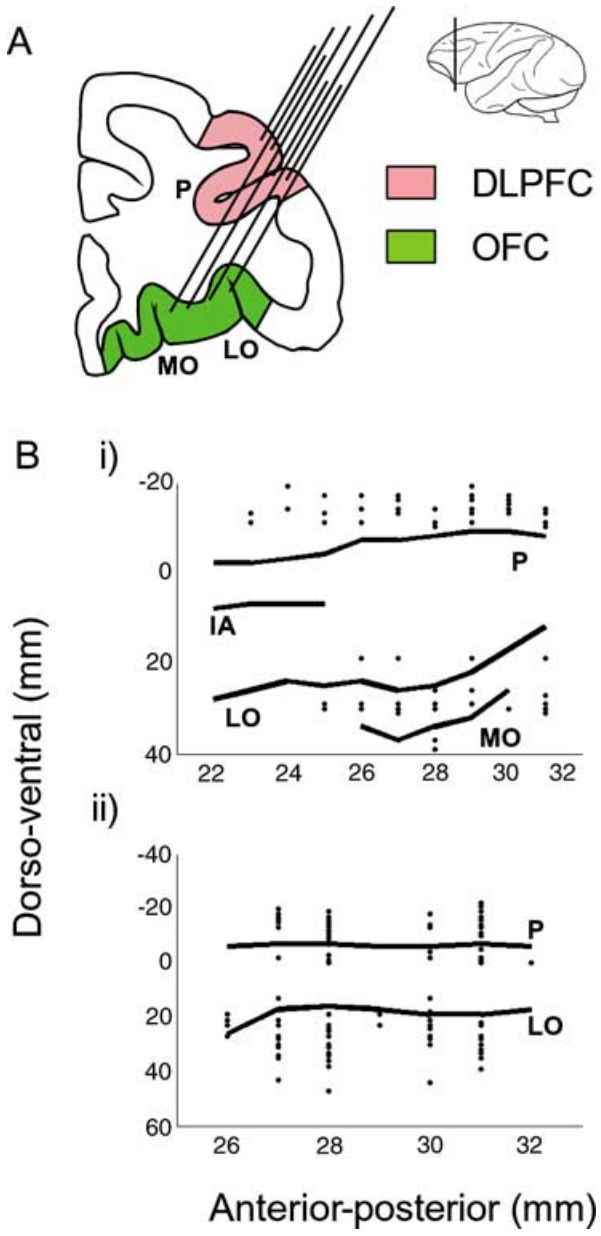

FIG. 2. (A) By positioning the recording chamber at the correct angle, neurons from both the DLPFC and OFC could be recorded simultaneously simply by lowering the electrodes the appropriate distance. The shading indicates the approximate extent from which neurons were sampled from both areas, through an illustrative coronal section of the frontal lobe. Neurons from the DLPFC were recorded from areas 9, 9/46 and 46, while neurons from the OFC were recorded from areas 11, 13 and 14 (Petrides & Pandya, 1994). (B) Flattened representation of the prefrontal cortex of both monkeys illustrating the position of the recording locations. P, principal sulcus; IA, inferior arcuate sulcus; LO, lateral orbital sulcus; MO, medial orbital sulcus.

a picture was presented to one side (at 5° eccentricity) of the fixation point for 500 ms. There was then a delay of 1000 ms before a second picture was presented for 500 ms on the opposite side to the first picture. After that, there was a second delay of 500 ms followed by two white squares which were simultaneously presented at the locations where the pictures had appeared. The monkey had to choose one of the pictures by making a direct saccade to its remembered location. The monkeys then received a fixed amount of juice depending on which picture they had chosen (from zero to eight drops, one drop of juice measuring ≈ 0.08 mL).

Stimuli and reward contingencies

We used two sets of three pictures (A, B and C; and X, Y and Z). For each set, one stimulus (A or X) always appeared as the first picture, and was presented to either the left or right of the fixation point following a random schedule. The second stimulus (B or C following A, Y or Z following X), was associated with either a smaller or larger amount of juice than the first picture, and appeared on the opposite side to the first (Fig. 1B). This design ensured that the monkey could not determine the saccade direction which would yield the greater reward based on the first picture alone. On half of the trials stimulus A (associated with four drops of juice) was presented first, followed by either stimulus B (eight drops) or stimulus C (two drops). On the other half of the trials stimulus X (associated with two drops of juice) was presented first, followed by either stimulus Y (four drops) or stimulus Z (zero drops). All trials were randomly interleaved.

New stimuli were used each day and the monkeys had to learn the associated reward values by trial and error. The pictures were complex and multicoloured. They were chosen from the internet, reduced in size to 1.8° and randomly assigned initial juice values. Once the monkey was consistently choosing the stimulus associated with the larger reward (i.e. on >27 out of the preceding 30 trials), the stimulus–reward contingencies were changed such that the 'correct' pictures, that is, the ones which yielded the larger reward, were now 'incorrect', yielding the smaller reward (Fig. 1B). This enabled us to dissociate neuronal activity related to the stimulus identity from activity related to the reward value. For example, on block 1, choosing B instead of A and choosing A instead of C resulted in a larger reward (Fig. 1B, top). Then, after the reversal, a larger reward was obtained when A was chosen instead of B and when C was chosen instead of A (Fig. 1B, bottom).

Data analysis

To determine the selectivity of a neuron, trials where the monkey was first learning the stimulus–reward contingencies were excluded by restricting the analysis to the 27 out of 30 criterion trials, as well as excluding those trials where the monkey chose the 'incorrect' picture (the one associated with the smaller amount of reward). There were frequently too few error trials to permit a meaningful analysis of the neuronal properties during the learning of the stimulus–reward contingencies (i.e. those trials in each block prior to the criterion trials). The trial was divided into four epochs: the presentation of the first picture (500 ms), the first delay (1000 ms), the second picture presentation (500 ms) and the second delay (500 ms).

Selectivity during the presentation of the first picture and the first delay epoch was analysed using a three-way ANOVA on the mean firing rate of the neuron with the factors being the amount of juice with which the picture was associated (two vs. four drops; see Fig. 1B), the location of the picture (left/right), and the identity of the picture (stimulus A or X). The analysis was performed by collapsing neuronal activity across blocks, thereby permitting the separation of the reward amount and the identity of the picture. A reward-selective neuron was defined as a neuron that had a main effect of reward (at $P < 0.05$) and no other significant main effects or interactions with either of the other two factors (at $P > 0.05$). Location- and stimulus-selective neurons were similarly defined. Neurons showing linear effects of two factors were defined as those with significant main effects of both factors but no significant interaction. Neurons showing nonlinear effects were defined as those showing a significant interaction between the two factors. A three-way ANOVA was also used to analyse neuronal activity during the presentation of the second picture and the second delay epoch. The factors were the location of the second picture (left/right),

the identity of the picture (stimulus B, C, Y or Z), and whether the second picture predicted more or less juice than the first picture. Because we restricted the analysis to 'correct' trials (defined as those trials where the monkey chose the larger reward) this last factor also corresponds to whether the monkey chose the first or second picture. Again, the analysis was collapsed across blocks. The selectivity of the neurons was defined as for the previous ANOVA. Differences in the prevalence of neurons between the two areas were assessed using a χ^2 test. For all statistical tests, significance was assessed at $P < 0.05$.

Once the second picture was presented, the monkey was able to predict how much juice it would receive at the end of the trial as well as plan the leftward or rightward saccade. We examined the influence of these factors on neuronal activity during the second picture and delay epochs using a two-way ANOVA. The factors were the direction of the forthcoming saccade (left vs. right) and the amount of juice the monkey expected to receive (two, four or eight drops).

The strength of the neuronal selectivity in each epoch was examined by performing a receiver operating characteristic (ROC) analysis (Green & Swets, 1966; Newsome *et al*., 1989; Rainer & Miller, 2000), which measures the degree of overlap between two response distributions. For each neuron, the preferred (more activity) and nonpreferred (less activity) conditions were compared, giving two distributions, P and N, respectively, of neuronal activity. For example, for a location-selective neuron these distributions might be the neuron's firing rate when a stimulus is presented on the left, compared to its firing rate when a stimulus is presented on the right. An ROC curve was then generated by taking each observed firing rate of the neuron and plotting the proportion of P which exceeded the value of that observation against the proportion of N which exceeded the value of that observation. The area under this curve was then calculated. A value of 0.5 would indicate that the two distributions completely overlap (because for each value of the neuron's firing rate the proportion of P and N exceeding that value is equal) and as such the neuron would not be selective. A value of 1.0, on the other hand, would indicate that the two distributions are completely separate (i.e. every value drawn from N is exceeded by the entire distribution of P, while none of the values of P are exceeded by any of the values in N) and so the neuron would be very selective. This method of analysis has the advantage that it is independent of the neuron's firing rate, and so can be used to compare neurons with different baseline firing rates and dynamic ranges. Furthermore, the ROC value can be thought of as the probability that an independent observer could identify, on the basis of the neuron's firing rate, which condition had been presented.

We analysed differences between the mean ROC values of different neuronal populations using *t*-tests assessed at $P < 0.05$. Nonparametric tests (Wilcoxon's rank-sum tests) yielded the same pattern of significant differences. Because we did not preselect the neurons from which we recorded, the mean ROC value of the population as a whole was often quite low, because many of the neurons were not necessarily selective for the factor under consideration. To ensure that the ROC values were above those expected by chance we performed a bootstrap analysis. For each neuron, the trials were randomly assigned to the different behavioural conditions, and the ROC value calculated. This process was repeated 1000 times for each neuron, and the mean ROC value was determined. A *t*-test then compared these mean ROC values derived from the shuffled data to the actual ROC values. Furthermore, because the ROC analysis only compares two distributions, in those situations where there were more than two conditions we chose to compare the conditions yielding the most and least neuronal activity. However, we also computed the ROC values for each pair of conditions and calculated the mean of these ROC values. While this method, by definition, yielded lower overall ROC values, the pattern of significant differences between the DLPFC and the OFC did not change.

To examine the time-course of the neurons' selectivity, we performed a sliding ROC analysis. Starting from the baseline period of the trial (500 ms of fixation prior to presentation of the first picture), an ROC value was calculated for a 200-ms epoch. This window was then stepped forward in 10-ms increments until the entire trial had been analysed. We used this analysis to compare the latency which selectivity appeared. A large window was necessary to ensure sufficient spikes to permit an ROC analysis. While this necessarily reduced the temporal resolution of our analysis (requiring our absolute latency values to be interpreted cautiously) it did not affect our conclusions regarding the relative latency of selectivity between the two areas. The latency for selectivity to appear was defined as the point at which the ROC curve exceeded 0.6. This criterion was chosen as one which yielded a close approximation to the time at which we judged selectivity to appear from the spike density histograms. For those neurons which exceeded this criterion we also looked at when the peak ROC value was reached. The differences that we observed between the areas were not dependent on the precise criterion that we used. However, if the criterion was set too low (< 0.58) the analysis became too noisy (due to spurious crossings of the threshold) while if the criterion was set too high (≥ 0.7) too few neurons reached the threshold to permit a meaningful statistical analysis. Other criteria (e.g. an ROC value in excess of 3 SD of the ROC values during the fixation period) yielded the same pattern of results.

Results

Behaviour

Monkey A performed an average of 11 reversals per day (range 7–15) over 16 recording sessions, while monkey B completed an average of 12 reversals per day (range 8–17) over 23 recording sessions. Prior to reaching criterion performance, Monkey B made a few more errors on each reversal than monkey A (11 vs. 7; $t_{446} = 4.1$, $P < 0.00005$). Analysing all the trials (including those prior to reaching criterion) revealed that the monkeys found the discrimination between eight and four drops of juice significantly easier than either the discrimination between four and two drops of juice or between two and zero drops of juice (monkey A, 93, 79 and 77%, respectively, one-way ANOVA and *post hoc* analysis using a Student–Newman–Keuls test, $F_{2,45} = 94$, $P < 5 \times 10^{-16}$; monkey B, 85, 77 and 79%, $F_{2,66} = 8$, $P < 0.001$).

After it reached criterion, monkey A's reaction times were not significantly different regardless of whether it would receive two drops of juice (mean 213 ms), four drops (mean 208 ms) or eight drops of juice (mean 207 ms; one-way ANOVA, $F_{2,45} = 2.7$, $P > 0.05$). By contrast, monkey B's reaction times were significantly slower when it would receive two drops (mean 259 ms) and four drops (mean 257 ms) than when it would receive eight drops (mean 241 ms, one-way ANOVA and *post hoc* analysis using a Student–Newman–Keuls, $F_{2,66} = 8.6$, $P < 0.00005$). The accuracy of the saccades (defined by the relative eye position immediately after the saccade relative to the target position) did not significantly differ for either monkey across the three different juice amounts (one-way ANOVA, monkey A, $F_{2,45} < 1$, $P > 0.1$; monkey B, $F_{2,66} < 1$, $P > 0.1$).

In principle, it was possible for the monkeys to perform the task above chance level while attending only to the second picture; they could learn to saccade either towards or away from the location of the second picture, depending on whether it was associated with a large or small amount of juice, respectively. To examine whether this was the case, we ran separate behavioural tests. Monkeys performed 10 sessions in which we alternated between sessions of the original task

and probe sessions in which the first picture was replaced by a blue square (and thus conveyed no information about which choice was correct). A comparison of the performance across these sessions revealed that their performance was much poorer on the probe sessions. Over the five probe sessions, monkey A solved 15 discriminations and made an average of 37 errors per discrimination while over the five sessions of the normal task the monkey solved 58 discriminations with an average of 8 errors per discrimination ($t_{71} = 5.49$, $P < 1 \times 10^{-6}$). On the probe sessions, Monkey B solved 17 discriminations while making an average of 37 errors per discrimination compared to solving 51 discriminations and making an average of 15 errors per discrimination on the standard task sessions ($t_{66} = 3.06$, $P < 0.005$). This indicates that that the monkeys did indeed attend to both pictures and based their choices on their relative reward value rather than solving the task using just the second picture.

Neuronal properties

A total of 301 neurons were recorded (Fig. 2): 167 from the DLPFC (72 from monkey A and 95 from monkey B) and 134 from the OFC (43 from monkey A and 91 from monkey B). In order to obtain sufficient statistical power, the neurons from the two monkeys were pooled. For all significant results, there were no qualitative differences between the two monkeys (i.e. the effects were in the same direction), unless otherwise noted.

The task involved the monkeys identifying the pictures, assessing their reward value and choosing the location of the picture which had the higher reward value. First, we examine neuronal activity reflecting attributes of the two pictures: their identity, location, and reward amount associated with each picture. Then, we report on neuronal activity which reflected the outcome of the picture comparison: the behavioural response and the expected reward value.

Encoding stimulus attributes: the location, identity and reward value of the pictures

Location selectivity

Some neurons had activity which varied with picture location. Figure 3A shows the activity of such a neuron across the entire trial, sorted according to the location and identity of the first picture, as well as the amount of juice which is associated with this picture. During the presentation of the first picture, the neuron had a slightly higher firing rate when the first picture was presented on the left, but not when it was presented on the right. During the subsequent delay and the presentation of the second picture, the neuron had a higher firing rate when the first picture had appeared on the right (and therefore the second picture appeared on the left).

We performed a three-way ANOVA on the neurons' firing rates during the picture and delay epochs using the factors of the picture's location, identity and the amount of juice with which it was associated (see Materials and methods). Neurons selective for just the location (left vs. right) of the pictures (main effect of the location factor assessed at a significance level of $P < 0.05$, with no other significant main effects or interactions, assessed at $P > 0.05$) were more numerous in the DLPFC than the OFC during the first picture, first delay and second picture epochs (Table 1).

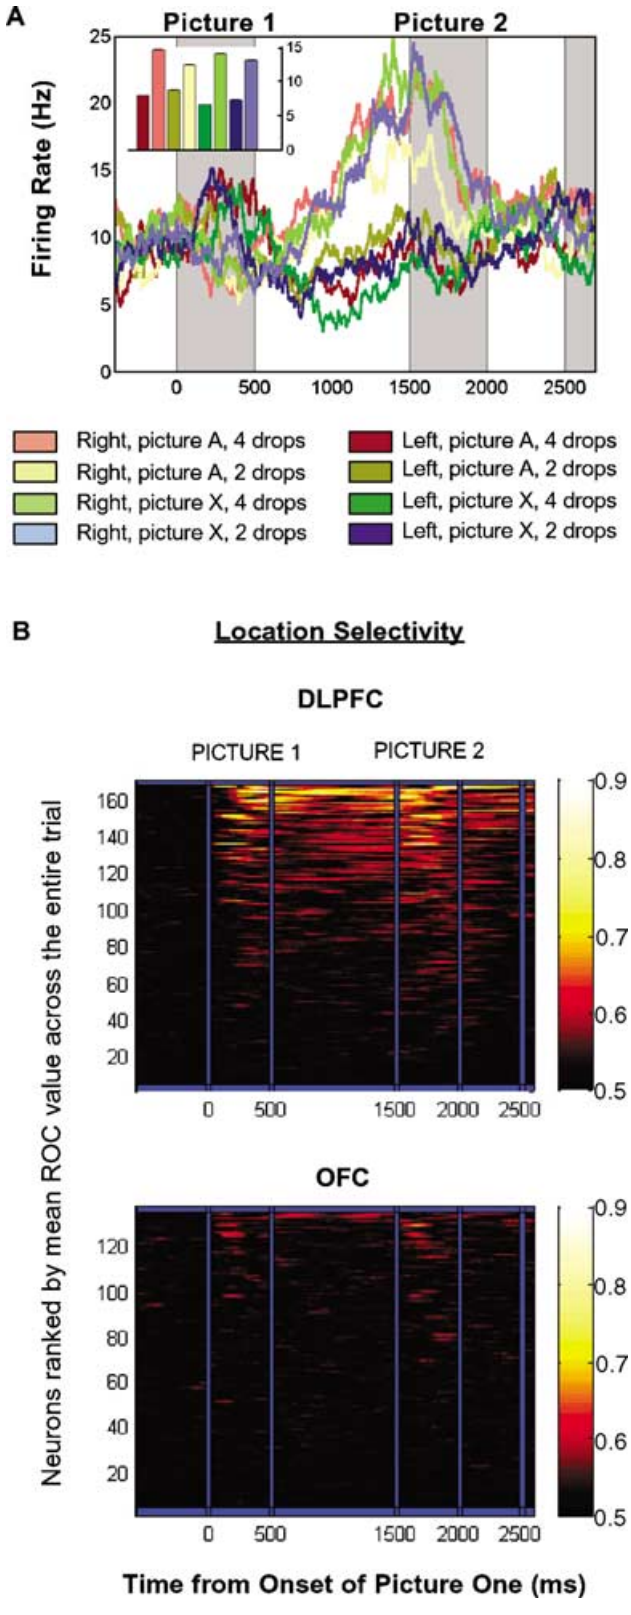

FIG. 3. (A) Example of a location-selective neuron recorded from the DLPFC. The neuron's activity is sorted according to the identity, location and reward assignment of the first picture. During the presentation of the first picture, the neuron was weakly selective for pictures appearing on the left but then, during the first delay and second picture epoch, the neuron showed a higher firing rate when the first picture had appeared on the right, and the second picture would appear on the left irrespective of the identity of the pictures or the number of drops of juice which they predicted. Inset bar graph illustrates the mean and SEM of the neuron's activity during the first delay epoch across the eight different conditions. (B) Time-course of location selectivity, using a sliding ROC analysis across all the neurons. The figures were constructed by sorting the neurons by their mean ROC value across both picture and both delay epochs. Location selectivity was stronger in the DLPFC than the OFC.

TABLE 1. Percentage of neurons encoding different stimulus attributes during the four epochs

	First picture		First delay		Second picture		Second delay	
	DLPFC	OFC	DLPFC	OFC	DLPFC	OFC	DLPFC	OFC
Main effects								
Reward	7	5	4	9	2	5	7	6
Location	14*	4*	18*	8*	16*	6*	5	3
Picture	5	11	4	7	1***	19***	4	7
Interactions								
Reward × Location								
Linear	4	4	7	3	4	1	5	2
Nonlinear	2	1	9	4	8	4	11	9
Total	7	5	16*	7*	11	4	17	11
Location × Picture								
Linear	2	2	3	3	5	4	2	3
Nonlinear	19	13	15	13	20	17	15	19
Total	21	15	18	16	25	21	17	22
Reward × Picture								
Linear	3	4	3	2	2	4	2	7
Nonlinear	4	6	5	3	7	4	0	4
Total	7	10	8	5	9	8	2**	12**

Neurons were defined, using a three-way ANOVA, according to whether they showed a main effect of a single factor with no significant nonlinear interactions, a main effect of two factors with no significant nonlinear interactions (that is, a linear interaction), or a significant nonlinear interaction; $n = 167$ (DLPFC), 134 (OFC). $^{*}P < 0.05$, $^{**}P < 0.005$, $^{***}P < 5 \times 10^{-8}$ (χ^2 test) for differences in the proportion of neurons between the two areas.

To examine the time-course of this selectivity, we performed a sliding ROC analysis using a 200-ms time window which was incremented across the trial in 10-ms steps (see Materials and methods). This was conducted for each and every recorded neuron and the results are shown in Fig. 3B. Plotted are each neuron's ROC values for every time step. ROC values are derived from the absolute difference in firing rate between left vs. right picture presentations, so values range from 0.5 (no location information) to nearly 1.0 (perfect discrimination between the locations). We sorted the values according to each neuron's mean ROC value across both picture and delay epochs. These plots suggest that the effect of location was stronger in the DLPFC. To quantify this, we computed an ROC analysis for each neuron using their average activity in the picture and delay epochs (see Materials and methods). Table 2 shows that the mean ROC values for location selectivity were significantly higher in the DLPFC than the OFC in all epochs.

TABLE 2. Mean strength of location and picture selectivity, as determined by the ROC analysis

	DLPFC	OFC	*P*-value
Location			
Picture one	0.549	0.526	$<5 \times 10^{-6}$
Delay one	0.566	0.532	$<5 \times 10^{-7}$
Picture two	0.574	0.531	$<1 \times 10^{-10}$
Delay two	0.552	0.526	$<5 \times 10^{-7}$
Picture			
Picture one	0.523	0.531	$<0.05^{B}$
Delay one	0.527	0.528	>0.1
Picture two	0.563	0.585	<0.00005
Delay two	0.567	0.580	$<0.05^{B}$

Significant differences between the two areas were assessed using a *t*-test; $n = 167$ (DLPFC), 134 (OFC). The superscript B indicates that the difference was only apparent in monkey B. In all of the cases, the mean ROC value was significantly different from that expected by chance (determined by bootstrap analysis, and assessed using a *t*-test with the criterion $P < 0.05$; see Materials and methods for details).

Picture selectivity

Some neurons exhibited selectivity for the identity of the picture. An example is shown in Fig. 4. This neuron was strongly selective for the presentation of picture C during the second picture epoch, as well as weakly selective for the presentation of picture A during the first picture epoch.

The results of the three-way ANOVA (see above) showed a bias towards the encoding of the picture in the OFC. Picture-selective neurons were defined as those showing a significant main effect of the picture identity factor, with no other significant main effects or interactions. They were significantly more abundant in the OFC than the DLPFC during the second picture epoch (Table 1).

To examine the time-course of picture selectivity, we performed a sliding ROC analysis on the neuron's preferred and nonpreferred pictures during the first picture epoch (Fig. 5A) and second picture epoch (Fig. 5B). In both cases picture selectivity appeared to be marginally stronger in the OFC than the DLPFC. To quantify this, we performed an ROC analysis using the neurons' average activity to the first picture, during the first picture and delay epochs, and the neurons' average activity to the second picture during the second picture and delay epochs. The mean picture-identity ROC values were significantly higher in OFC than DLPFC during both picture epochs but not the delay epochs (Table 2).

Associated reward value

Relatively few neurons encoded just the reward value (number of drops of juice) currently associated with each picture (i.e. showed a significant main effect of the reward factor with no other significant main effects or interactions according to the three-way ANOVA described above; see Table 1). However, there were significantly more neurons in the DLPFC than the OFC whose activity simultaneously reflected the reward value and location of the pictures during the first delay epoch (see Table 1; a significant main effect of both reward and location, or a significant interaction between the two factors according to the three-way ANOVA). Further, there were more neurons in the OFC whose activity simultaneously reflected the reward value and identity of the pictures during the second delay epoch (see Table 1; a significant main effect of both reward and picture identity, or a significant interaction between the two factors according to the three-way ANOVA). The

relatively few reward value-selective neurons precluded a detailed statistical comparison between the areas. However, we will see below that, in contrast to the relatively sparse encoding of the exact reward value associated with each picture, many neurons encoded the higher of the two reward values, which was the value the monkey could expect to receive upon successful completion of the trial.

In summary, there was a difference between the DLPFC and the OFC in selectivity for the identity and for the location of the picture. In terms of both the prevalence of selective neurons and the strength of selectivity, the DLPFC was biased towards processing location while the OFC was biased towards the processing of the picture's identity. However, it should also be noted that there were many neurons in both areas which encoded both types of information. Indeed, the most frequently observed interaction was between the Picture and Location factors (see Table 1).

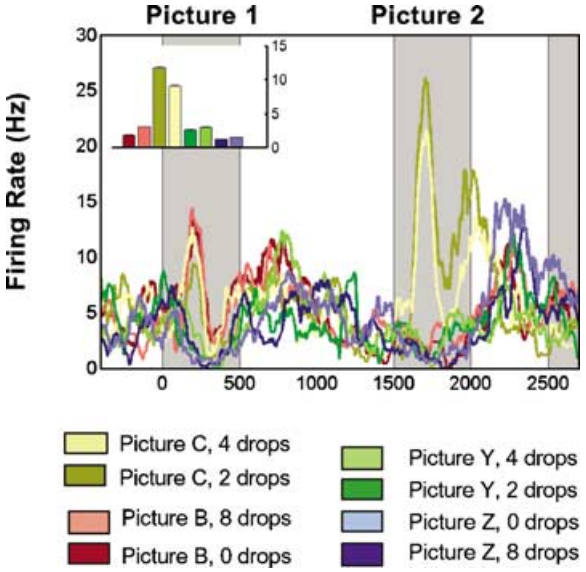

FIG. 4. (A) Example of a picture-selective neuron recorded from the OFC. The neuron's activity was sorted according to the identity and reward assignment of the second picture (collapsed across location for clarity). This neuron exhibited a higher firing rate when picture C was the second picture, rather than the other three pictures, regardless of whether picture C predicted four drops of juice or, after the stimulus–reward contingencies had been reversed, two drops of juice. The neuron was also weakly selective to picture A during the presentation of the first picture. Inset bar graph illustrates the mean and SEM of the neuron's activity during the presentation of the second picture across the different pictures and reward values.

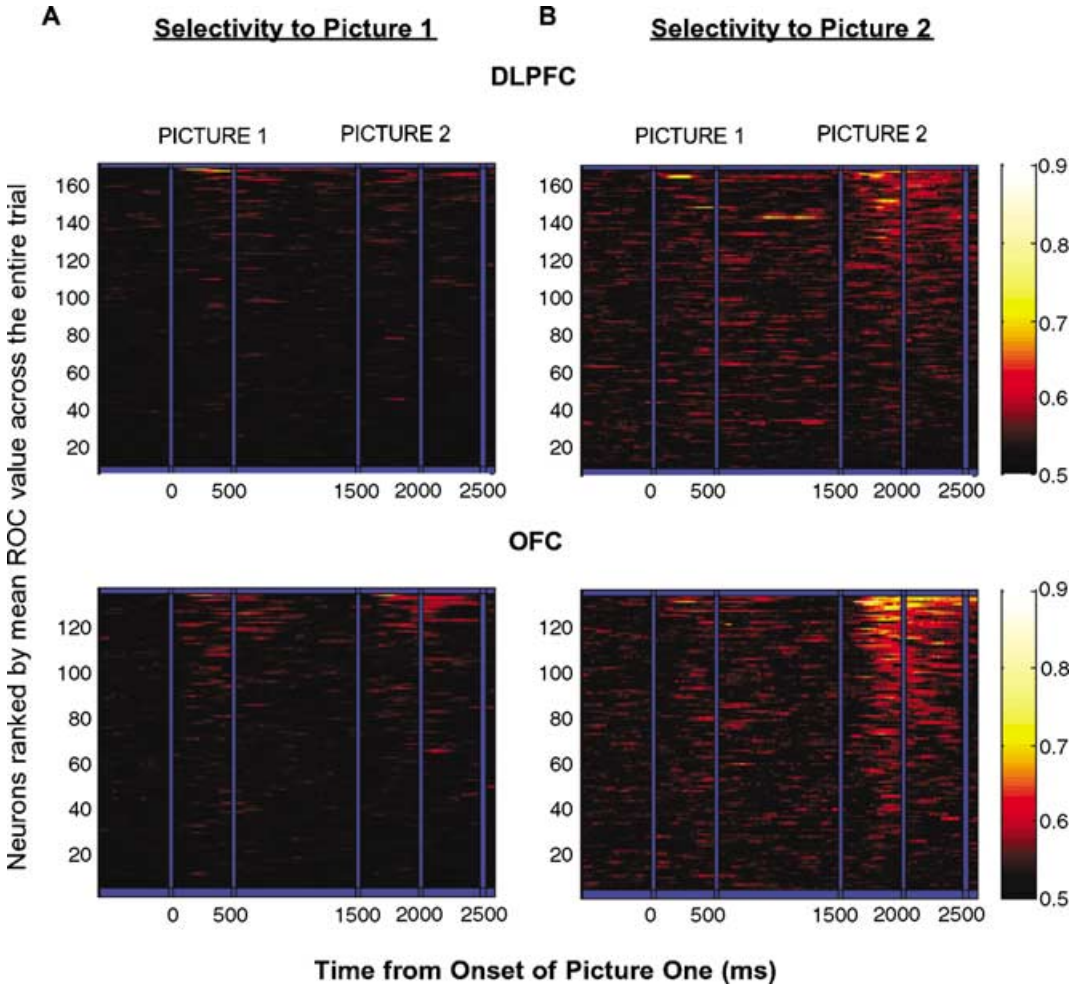

FIG. 5. Time-course of selectivity to (A) the first picture and (B) the second picture as determined by the ROC analysis. The figures were constructed by sorting the neurons by their mean ROC value across both picture and delay epochs. Picture selectivity was stronger for the second picture, and was marginally stronger in the OFC compared to the DLPFC.

Activity encoding the forthcoming behavioural response and the expected reward

Once the second picture was presented, the monkey could plan its behavioural response (a saccade to the left or right) as well as predict how many drops of juice it would receive for doing so. Correspondingly, information about both the forthcoming saccade and the expected reward began to appear in neuronal activity after onset of the second picture. Note that encoding of the expected reward value is different from encoding the reward value associated with individual pictures (discussed above). Here, we focus on the amount of reward value which the animal would receive at the end of the trial which, on correct trials, would be the larger of the reward values associated with the two pictures. In contrast to the few PFC neurons which reflected the reward value of the individual pictures (see above), there was a robust representation of the expected reward value in PFC neuronal activity.

Expected reward value

An example of a neuron whose activity reflected the value of the to-be-delivered reward is shown in Fig. 6A. Beginning shortly after onset of the second picture, it had a higher firing rate if four drops of juice would be delivered (for a correct response) on that trial and a lower rate if two or eight drops would be delivered. Another example is shown in Fig. 6B. This neuron showed a depression in its firing rate to the second picture which was most pronounced when eight drops of juice would be delivered and showed a progressively less pronounced depression for four and two drops of juice, respectively. An example of a neuron which reflected both the to-be-delivered reward amount and the forthcoming saccade is shown in Fig. 6C. It showed a higher firing rate when two drops of juice would be delivered and a lower firing rate to eight drops, but fired most when a leftward saccade would result in two drops of juice.

Neurons selective for the forthcoming behavioural response (saccade right or saccade left) and/or the to-be-delivered reward amount (two, four or eight drops) were identified using a two-way ANOVA on the average activity in the second picture and delay epochs. The results are listed in Table 3. More OFC than DLPFC neurons encoded the to-be-delivered reward only (defined as those showing a significant main effect of Reward, assessed at $P < 0.05$, with no significant main effect of Saccade and no significant Reward × Saccade interaction; e.g.

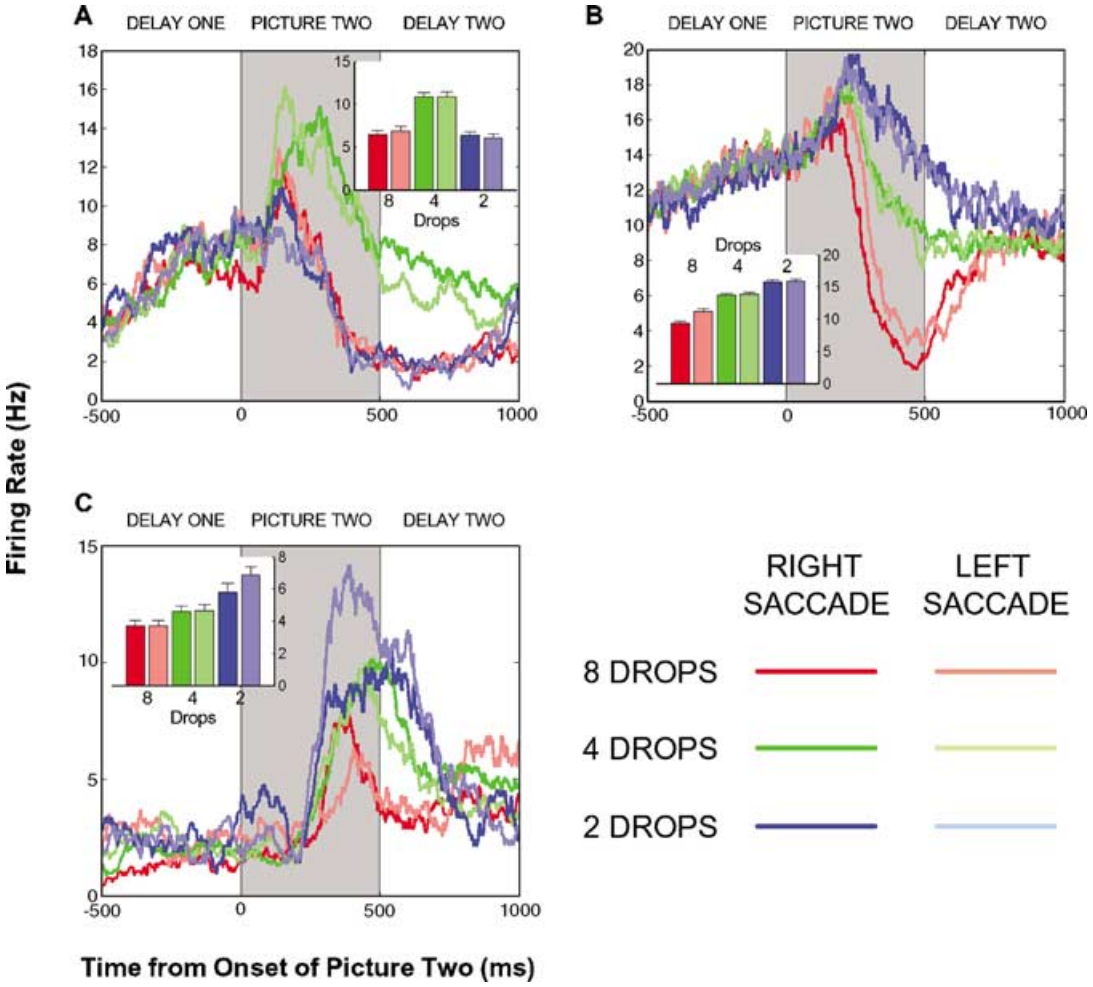

FIG. 6. Examples of neurons predicting the amount of juice the monkey would receive, and/or the monkey's response. (A) A reward-predicting OFC neuron which showed a higher firing rate when four drops of juice were predicted, as opposed to eight drops or two drops. (B) A reward-predicting OFC neuron which encoded the reward parametrically, showing a depression in its firing rate which was strongest for eight drops of juice and progressively weaker for smaller amounts of reward. (C) A DLPFC neuron which encoded the reward parametrically by showing its highest firing rate when two drops of juice were predicted, a lower firing rate to four drops and its lowest firing rate to eight drops, but also responded more strongly when two drops were predicted on the left as opposed to on the right. In all three figures the inset bar graph illustrates the mean and SEM of the neuron's activity during the presentation of the second picture for the two different responses and three different reward values.

TABLE 3. Percentage of neurons predicting the monkey's saccade and/or the expected reward during the presentation of the second picture and the following delay

	Picture		Delay	
	DLPFC	OFC	DLPFC	OFC
Main Effects				
Saccade	1	3	5	4
Reward	13^{**}	28^{**}	26^{*}	40^{*}
Interaction				
Saccade × Reward				
Linear	1	1	7	1
Nonlinear	42^{***}	18^{***}	32^{***}	11^{***}
Total	43^{***}	19^{***}	39^{****}	13^{****}

Examples are shown in Fig. 6; $n = 167$ (DLPFC), 134 (OFC). $^{*}P < 0.05$, $^{**}P < 0.005$, $^{***}P < 0.000\,05$, $^{****}P < 0.000\,001$ for differences in the proportion of neurons between the two areas (χ^2 test).

Figure. 6A and B). They were significantly more abundant in the OFC during both the presentation of the second picture (OFC 38/134 or 28%, DLPFC 22/167 or 13%; $\chi^2 = 9.8$, $P < 0.005$) and the second delay epoch (OFC 53/134 or 40%, DLPFC 44/167 or 26%; $\chi^2 = 5.3$, $P < 0.05$).

By contrast, the opposite pattern was seen for neurons whose activity simultaneously reflected the forthcoming saccade and the expected reward value; they were more abundant in the DLPFC than the OFC. These neurons were defined as showing a significant Reward × Saccade interaction, or a significant main effect of both Reward and Saccade, according to the two-way ANOVA (see Table 3 and Fig. 6C). They were more numerous in the DLPFC than the OFC during both the second picture epoch (DLPFC 72/167 or 43%, OFC 25/134 or 19%; $\chi^2 = 19$, $P < 0.00005$) and the second delay epoch (DLPFC 65/167 or 39%, OFC 17/134 or 13%; $\chi^2 = 25$, $P < 0.000001$).

We compared the time-course of selectivity for the to-be-delivered reward using the sliding ROC analysis (see above and Materials and methods). Both the OFC and the DLPFC contained neurons which predicted the reward (Fig. 7A) although some of the DLPFC neurons were also modulated by the saccade (Fig. 7B). To determine the time at which selectivity for the to-be-delivered reward first appeared we calculated the point at which the sliding ROC value first exceeded a criterion of 0.6 (see Materials and methods). This measure did not differ between the two areas; 36% (60/167) of the DLPFC neurons reached criterion in a mean time of 467 ms while 39% (52/134) of the OFC neurons reached criterion in a mean time of 426 ms ($t_{110} = 1.0$, $P > 0.1$).

However, while selectivity for the reward tended to appear at about the same time in both areas, it then rose more rapidly and peaked earlier in the OFC than in the DLPFC (Fig. 8A). So, for those neurons

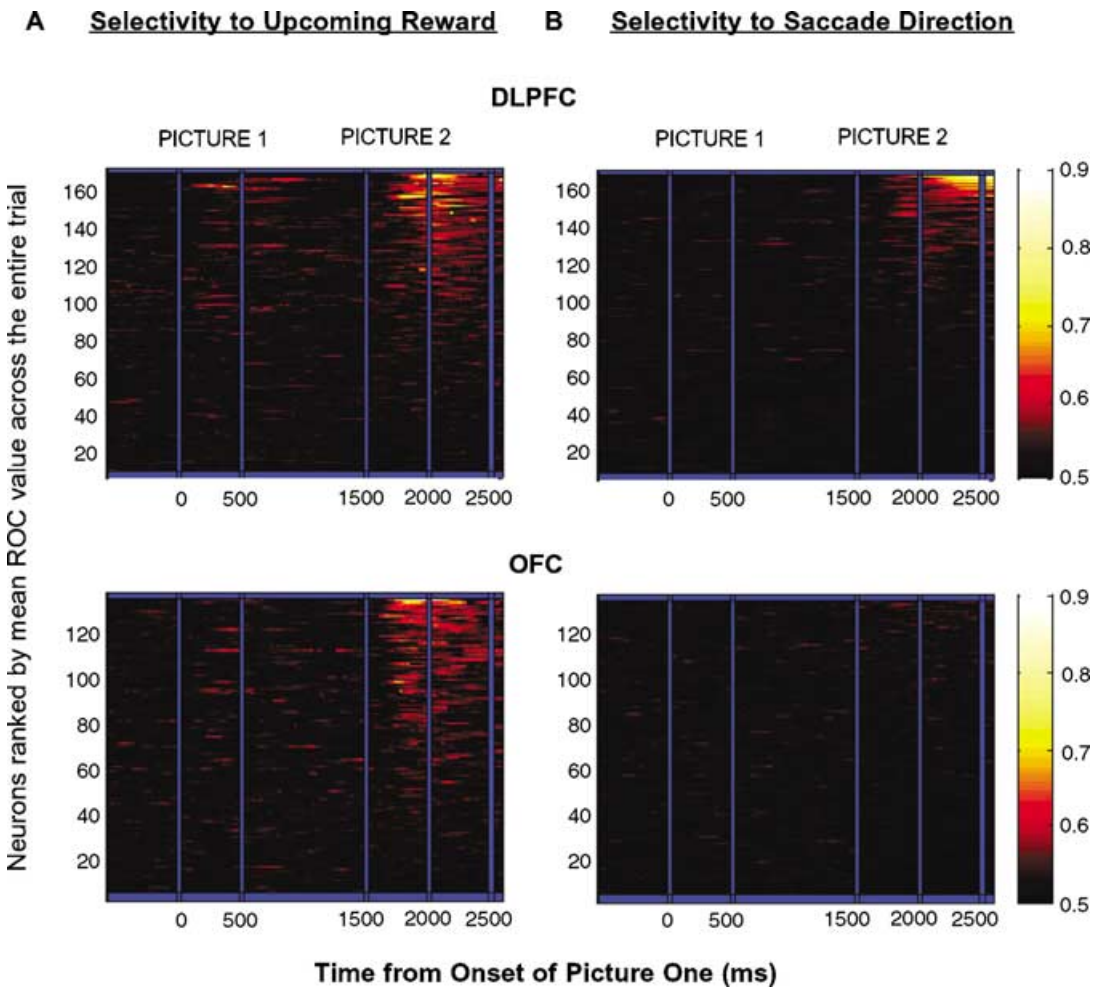

FIG. 7. Time-course of selectivity for predicting (A) the eventual reward amount and (B) the upcoming saccade direction as determined by the ROC analysis. The figures were constructed by sorting the neurons by their mean ROC value across both picture and both delay epochs.

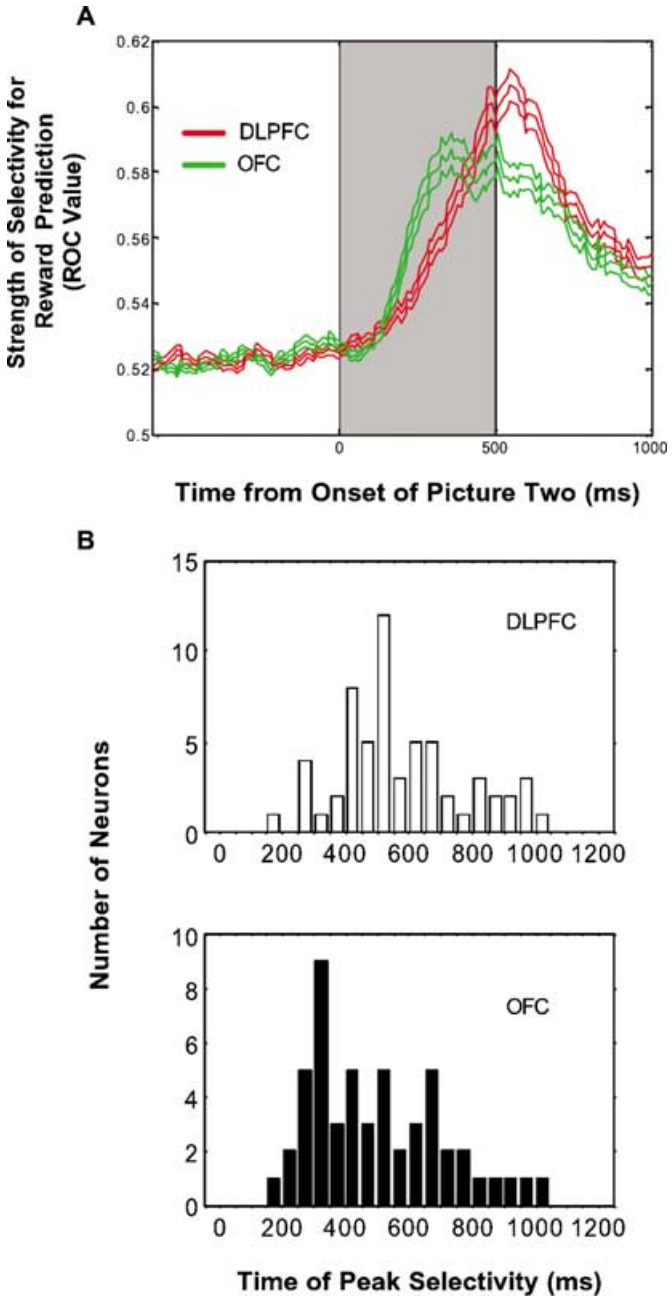

FIG. 8. (A) Time-course of reward selectivity as determined by the sliding ROC analysis across the DLPFC and OFC populations of reward-selective neurons. Reward selectivity increased in both areas, ≈300 ms after the presentation of the second picture. Although it increased to about the same level in both areas, the peak selectivity was reached earlier in the OFC than the DLPFC. The thick line indicates the mean selectivity of the neurons, while the thinner lines indicate the SEM. (B) Distribution of time of peak reward selectivity across the DLPFC and OFC populations of reward-selective neurons. The mean time of peak reward selectivity was 510 ms after the onset of picture two in the OFC, and 592 in the DLPFC.

which reached criterion we calculated the value and time of the peak ROC value between the onset of the second picture and the start of the behavioural response. There was no difference between the two areas in the mean peak ROC value (DLPFC = 0.654, OFC = 0.646, $t_{110} = 0.89$, $P > 0.1$), but the peak was reached significantly earlier in the OFC than the DLPFC. Figure 8B shows a distribution of the times at which each neuron reached its peak ROC value for the to-be-delivered reward. On average, this occurred ≈80 ms earlier in the OFC (510 ms after the onset of picture two) than the DLPFC (592 ms after the onset of picture two; $t_{110} = 2.1$, $P < 0.05$).

TABLE 4. Mean strength of selectivity for predicting the upcoming reward amount and saccade direction, as determined by the ROC analysis across the entire population of neurons from which we recorded

	DLPFC Mean ROC value	OFC Mean ROC value	*P*-value
Reward			
Picture two	0.565	0.565	>0.1
Delay two	0.592	0.573	<0.01
Saccade			
Picture two	0.527	0.517	<0.00005
Delay two	0.545	0.521	<0.000001

Significant differences between the two areas were assessed using a *t*-test; $n = 167$ (DLPFC), 134 (OFC). In all of the cases the mean ROC values were significantly higher than expected by chance, except for the Saccade ROC values in the OFC which were not significantly greater than chance (determined by bootstrap analysis, and assessed using a *t*-test at $P < 0.05$; see Materials and methods for details).

We also compared the strength of selectivity for the to-be-delivered reward by computing ROC values on the average activity across the second picture epoch and second delay epoch (Table 4). Above, we reported that more OFC neurons encode the reward alone while more DLPFC encoded the combination of the reward and saccade. Consequently, the overall strength of selectivity for the expected reward was similar in the OFC and DLPFC during the presentation of the second picture. However, the DLPFC values were significantly higher than the OFC for the second delay epoch.

Finally, we examined whether neurons encoded the to-be-delivered reward parametrically (i.e. had the most activity when eight drops would be delivered and least to two drops, or vice versa, such as the neuron in Fig. 6B) or whether they encoded reward nonparametrically (i.e. did not rank the forthcoming reward amounts in order, such as the neuron in Fig. 6A). Because there were six different combinations in which the neuron could rank the three reward values, two of which were parametric, a third of the neurons would be expected to encode the reward parametrically by chance. During the second picture epoch, 8/39 (21%) of OFC neurons and 10/23 (43%) of DLPFC neurons encoded the reward parametrically. During the second delay, 13/42 (31%) of OFC neurons and 23/35 (66%) of DLPFC encoded the reward parametrically, this last proportion being the only one which significantly differed from chance level (binomial test, $P < 0.00005$).

Behavioural response

Selectivity for the forthcoming saccade direction also differed between the DLPFC and OFC. The two-way (Saccade × Reward) ANOVA indicated that neurons selective for saccade direction were largely located in the DLPFC (Table 3; a significant main effect of Saccade or a significant Saccade × Reward interaction, assessed at $P < 0.05$). This was mirrored in the ROC analysis of the strength of saccade direction selectivity (Fig. 7B), whereas saccade selectivity was virtually absent in the OFC. ROC values were calculated for every recorded neuron using their average activity in the second picture epoch and the second delay epoch. Table 4 shows that the mean DLPFC ROC values were higher than the OFC values in both the second picture and delay epochs.

Discussion

Our results indicate that neuronal signals related to critical task variables: reward value, the location of the pictures, identity of the

pictures and saccadic eye movements were evident in both the DLPFC and the OFC. This first and foremost suggests a good deal of overlap in their properties and that both PFC regions are capable of representing a diverse range of information. However, our results also indicate differences between the areas which suggest some relative functional specializations. Information about the location of the pictures was more evident (i.e. more neurons were significantly selective and population selectivity was stronger) in the DLPFC than the OFC, while information about the identity of the pictures was more evident in the OFC than DLPFC. In encoding forthcoming, end-of-trial events, DLPFC neurons were more likely to encode a combination of the to-be-delivered reward value and the saccade direction while OFC neurons were more likely to encode the reward value alone. While information about the to-be-delivered reward was of similar strength in both areas, it peaked earlier in the OFC than DLPFC.

Functional implications

The finding that OFC activity has a greater tendency to encode reward value independently of the behavioural response is consistent with its anatomy. The OFC is heavily and reciprocally connected with gustatory and olfactory cortices (Morecraft *et al*., 1992; Carmichael & Price, 1995b), as well as the basolateral amygdala which might provide the OFC with information as to the value of the reward (Baxter & Murray, 2002; Cardinal *et al*., 2002). Thus, the OFC is conceivably the first prefrontal region which would receive information about the value of the forthcoming juice reward. Our observation that reward value information peaks sooner in the OFC than the DLPFC is consistent with that notion.

Our finding that the DLPFC has a greater tendency to encode the to-be-delivered reward along with the behavioural response (e.g. a given neuron might show the highest activity when, say, a saccade to the left would result in, say, a reward of four drops) also corresponds with known PFC anatomy. The DLPFC does not receive direct inputs from gustatory and olfactory cortex (Mesulam & Mufson, 1982; Barbas, 1993), but it is heavily interconnected with areas of the frontal lobe responsible for oculomotor control (Petrides & Pandya, 1999). Thus, the DLPFC may be an area where information about reward value converges with information about the animal's actions, allowing the monkey to make a comparison between the two different reward amounts. Similarly, the DLPFC has also been implicated in the comparison of different somatosensory stimuli (Romo *et al*., 1999; Romo & Salinas, 2003). This is not to say that reward information and motor actions are not integrated in other cortical areas, such as the anterior cingulate cortex (Paus, 2001) and the parietal cortex (Platt & Glimcher, 1999).

This does not mean that the DLPFC is the only route by which reward information from the OFC can influence behaviour; the OFC is also connected to the striatum, which in turn is interconnected with motor system structures (Yeterian & Pandya, 1991), as well as the hypothalamus and amygdala (Carmichael & Price, 1995a) where it can exert influence over the autonomic nervous system. It is possible that we may have seen more integration of reward and motor information in the OFC if we had required the monkey to make a different motor response, such as an arm movement, rather than an eye movement (Passingham, 1993). But it may be that reward value information enters the PFC through the OFC and is then relayed to the DLPFC. If this hypothesis is true, inactivation of the DLPFC should not affect reward information in the OFC whereas inactivation of the OFC should attenuate reward information in the DLPFC.

Differences between the DLPFC and OFC were not restricted to the way in which they processed rewards. These areas also differed in the degree to which they encoded the physical properties of the stimuli. Neurons in the DLPFC had a greater tendency to encode the location of the stimuli, whereas the OFC had a greater tendency to encode their identity (although many neurons in both areas were capable of encoding both types of information). This too is consistent with the anatomical connections of the two areas; the DLPFC is more heavily interconnected to areas processing spatial information, such as the parietal cortex (Pandya & Barnes, 1987), while the OFC is more heavily interconnected to areas processing visual form information, such as the inferior temporal cortex (Carmichael & Price, 1995b).

Of course, we should keep in mind that, even with these differences, these two prefrontal regions are also characterized by a large degree of overlap in their neuronal properties. This suggests a great deal of exchange and communication between them, which is consistent with the extensive interconnections between different prefrontal areas (Barbas & Pandya, 1989). Models of the role of the PFC in executive brain functions suggest that such overlap of disparate information is essential for learning and representing the myriad, arbitrary relationships which describe our knowledge of how the world works (Norman & Shallice, 1986; Dehaene & Changeux, 1995; Miller, 2000; Shimamura, 2000; Miller & Cohen, 2001). Such knowledge is central to the role of executive functions in planning and executing goal-directed behaviour (Miller & Cohen, 2001).

Furthermore, although we have focused on the general properties of the neurons, such as whether they are influenced by the location or identity of the pictures, the actual pattern of selectivity observed in individual neurons is often quite complex. For example, the neuron in Fig. 3A shows a higher firing rate during the first picture epoch when the stimulus is presented on the left, but a higher firing rate during the delay when the first picture is presented on the right (and thus the second picture will appear on the left). However, such complexity in the neuron's individual responses does not alter the general bias towards processing different aspects of the task in the two different PFC areas.

Relation to prior studies

Our findings are in line with both prior neurophysiological and neuropsychological investigations of the PFC. Previous studies have reported neurons encoding reward value and type (e.g. raisin vs. cabbage) in the DLPFC (Watanabe, 1996; Leon & Shadlen, 1999; Kobayashi *et al*., 2002) and the OFC (Thorpe *et al*., 1983; Tremblay & Schultz, 1999). Furthermore, our results suggesting that the OFC may be a source of reward signals to the DLPFC may explain why neuropsychological studies have pointed to the relative importance of the OFC for tasks that require assessment of reward value (Gaffan & Murray, 1990; Baylis & Gaffan, 1991; Gaffan *et al*., 1993; Dias *et al*., 1996; Bechara *et al*., 1998; Baxter *et al*., 2000). Lesions of the DLPFC have less effect on such tasks (Dias *et al*., 1996; Bechara *et al*., 1998), perhaps because the reward information from the OFC is able to control the monkey's behaviour via other outputs, such as those to the striatum and the autonomic nervous system.

Our results are similar to studies by Schoenbaum and colleagues which have investigated the properties of OFC neurons in the rat. They showed, using an odour discrimination task, that OFC neurons predict the outcome of the trial, specifically whether the rat will receive a sweet (positive) taste or a bitter (negative) taste (Schoenbaum *et al*., 1998). Furthermore, the neurons showed a rapid reversal in their selectivity which paralleled the reversal in the odour–outcome associations, similar to the OFC neurons recorded in the present study. It has been argued that the OFC is critically important for the flexible encoding of stimulus–reward associations (Rolls, 1996) based on the finding that OFC neurons (Thorpe *et al*., 1983), but not amygdala neurons (Sanghera *et al*., 1979), could rapidly reverse these associations. However, Scheonbaum showed that amygdala neurons could,

under the right circumstances, reverse even more rapidly than OFC neurons (Schoenbaum *et al*., 1999). Our results suggest that the DLPFC is a third area which is capable of such reversals.

In sum, the number of areas of the primate brain which have neurons responsive to reward is growing rapidly (Schultz, 2000). This includes the caudate nucleus (Hikosaka *et al*., 1989; Apicella *et al*., 1991; Apicella *et al*., 1997; Hassani *et al*., 2001), the nucleus accumbens (Apicella *et al*., 1991; Schultz *et al*., 1992), the subthalamic nucleus (Matsumura *et al*., 1992), the substantia nigra (Schultz, 1986; Schultz & Dickinson, 2000; Waelti *et al*., 2001), the thalamus (Komura *et al*., 2001), the anterior cingulate cortex (Niki & Watanabe, 1976), the amygdala (Nishijo *et al*., 1988), the hypothalamus (Burton *et al*., 1976), the perirhinal cortex (Liu & Richmond, 2000), the parietal cortex (Platt & Glimcher, 1999), the hippocampus (Watanabe & Niki, 1985) and the supplementary eye fields (Amador *et al*., 2000) as well as the PFC. The ultimate purpose of the brain is to help us obtain those things which enhance our survival, and avoid those things which threaten us, and so it should not be surprising that so many regions are influenced by rewards. Many of these regions are thought to have specific roles in reward processing (e.g. Schultz, 2000), but it is unclear to what extent these roles overlap; this will require recording from multiple areas under the same experimental conditions (Schoenbaum *et al*., 1998, 1999). To that end, we have shown that two prefrontal areas encode reward value, but also found differences in the nature and timing of this encoding which suggests the possibility that the information is serving different functions. It may be that OFC primarily encodes the reward per se, while the DLPFC uses this information to control behaviour.

Acknowledgements

This work was supported by NINDS grant NS35145 and the RIKEN-MIT Neuroscience Research Center. The Wellcome Trust and the McDonnell-Pew Foundation supported J.D.W. The authors thank Marlene Wicherski for valuable comments on this manuscript.

Abbreviations

DLPFC, dorsolateral prefrontal cortex; MRI, magnetic resonance imaging; OFC, orbitofrontal cortex; PFC, prefrontal cortex; ROC, receiver operating characteristic.

References

Amador, N., Schlag-Rey, M. & Schlag, J. (2000) Reward-predicting and reward-detecting neuronal activity in the primate supplementary eye field. *J. Neurophysiol.*, **84**, 2166–2170.

Apicella, P., Legallet, E. & Trouche, E. (1997) Responses of tonically discharging neurons in the monkey striatum to primary rewards delivered during different behavioral states. *Exp. Brain Res.*, **116**, 456–466.

Apicella, P., Ljungberg, T., Scarnati, E. & Schultz, W. (1991) Responses to reward in monkey dorsal and ventral striatum. *Exp. Brain Res.*, **85**, 491–500.

Barbas, H. (1993) Organization of cortical afferent input to orbitofrontal areas in the rhesus monkey. *Neuroscience*, **56**, 841–864.

Barbas, H. & Pandya, D.N. (1989) Architecture and intrinsic connections of the prefrontal cortex in the rhesus monkey. *J. Comp. Neurol.*, **286**, 353–375.

Baxter, M.G. & Murray, E.A. (2002) The amygdala and reward. *Nature Rev. Neurosci.*, **3**, 563–573.

Baxter, M.G., Parker, A., Lindner, C.C., Izquierdo, A.D. & Murray, E.A. (2000) Control of response selection by reinforcer value requires interaction of amygdala and orbital prefrontal cortex. *J. Neurosci.*, **20**, 4311–4319.

Baylis, L.L. & Gaffan, D. (1991) Amygdalectomy and ventromedial prefrontal ablation produce similar deficits in food choice and in simple object discrimination learning for an unseen reward. *Exp. Brain Res.*, **86**, 617–622.

Bechara, A., Damasio, H., Tranel, D. & Anderson, S.W. (1998) Dissociation of working memory from decision making within the human prefrontal cortex. *J. Neurosci.*, **18**, 428–437.

Burton, M.J., Rolls, E.T. & Mora, F. (1976) Effects of hunger on the responses of neurons in the lateral hypothalamus to the sight and taste of food. *Exp. Neurol.*, **51**, 668–677.

Cardinal, R.N., Parkinson, J.A., Hall, J. & Everitt, B.J. (2002) Emotion and motivation: the role of the amygdala, ventral striatum, and prefrontal cortex. *Neurosci. Biobehav. Rev.*, **26**, 321–352.

Carmichael, S.T. & Price, J.L. (1995a) Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *J. Comp. Neurol.*, **363**, 615–641.

Carmichael, S.T. & Price, J.L. (1995b) Sensory and premotor connections of the orbital and medial prefrontal cortex of macaque monkeys. *J. Comp. Neurol.*, **363**, 642–664.

Damasio, A.R. (1994) *Descartes' Error: Emotion, Reason, and the Human Brain*. Putman, New York.

Dehaene, S. & Changeux, J.P. (1995) Neuronal models of prefrontal cortical functions. *Ann. NY Acad. Sci.*, **769**, 305–319.

Dias, R., Robbins, T.W. & Roberts, A.C. (1996) Dissociation in prefrontal cortex of affective and attentional shifts. *Nature*, **380**, 69–72.

Duncan, J., Emslie, H., Williams, P., Johnson, R. & Freer, C. (1996) Intelligence and the frontal lobe: the organization of goal-directed behavior. *Cognit. Psychol*, **30**, 257–303.

Fuster, J.M. (1997) *The Prefrontal Cortex: Anatomy, Physiology and Neuropsychology of the Frontal Lobe*. Lippincott, Williams & Wilkins PA, USA.

Gaffan, D. & Murray, E.A. (1990) Amygdalar interaction with the mediodorsal nucleus of the thalamus and the ventromedial prefrontal cortex in stimulus-reward associative learning in the monkey. *J. Neurosci.*, **10**, 3479–3493.

Gaffan, D., Murray, E.A. & Fabre-Thorpe, M. (1993) Interaction of the amygdala with the frontal lobe in reward memory. *Eur. J. Neurosci.*, **5**, 968–975.

Goldman-Rakic, P.S. (1987) Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In Plum, F. (Ed.), *Handbook of Physiology, the Nervous System, Higher Functions of the Brain*. American Physiological Society, Bethesda, MD, pp. 373–417.

Green, C.G. & Swets, J.A. (1966) *Signal Detection Theory and Psychophysics.* J. Wiley and Sons, New York.

Hassani, O.K., Cromwell, H.C. & Schultz, W. (2001) Influence of expectation of different rewards on behavior-related neuronal activity in the striatum. *J. Neurophysiol.*, **85**, 2477–2489.

Hikosaka, O., Sakamoto, M. & Usui, S. (1989) Functional properties of monkey caudate neurons. III. Activities related to expectation of target and reward. *J. Neurophysiol.*, **61**, 814–832.

Kawasaki, H., Kaufman, O., Damasio, H., Damasio, A.R., Granner, M., Bakken, H., Hori, T., Howard, M.A. & Adolphs, R. (2001) Single-neuron responses to emotional visual stimuli recorded in human ventral prefrontal cortex. *Nature Neurosci.*, **4**, 15–16.

Kobayashi, S., Lauwereyns, J., Koizumi, M., Sakagami, M. & Hikosaka, O. (2002) Influence of reward expectation on visuospatial processing in macaque lateral prefrontal cortex. *J. Neurophysiol.*, **87**, 1488–1498.

Komura, Y., Tamura, R., Uwano, T., Nishijo, H., Kaga, K. & Ono, T. (2001) Retrospective and prospective coding for predicted reward in the sensory thalamus. *Nature*, **412**, 546–549.

Leon, M.I. & Shadlen, M.N. (1999) Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. *Neuron*, **24**, 415–425.

Liu, Z. & Richmond, B.J. (2000) Response differences in monkey TE and perirhinal cortex: stimulus association related to reward schedules. *J. Neurophysiol.*, **83**, 1677–1692.

Matsumura, M., Kojima, J., Gardiner, T.W. & Hikosaka, O. (1992) Visual and oculomotor functions of monkey subthalamic nucleus. *J. Neurophysiol.*, **67**, 1615–1632.

Mesulam, M.M. & Mufson, E.J. (1982) Insula of the old world monkey. III: Efferent cortical output and comments on function. *J. Comp. Neurol.*, **212**, 38–52.

Miller, E.K. (2000) The prefrontal cortex and cognitive control. *Nature Rev. Neurosci.*, **1**, 59–65.

Miller, E.K. & Cohen, J.D. (2001) An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.*, **24**, 167–202.

Morecraft, R.J., Geula, C. & Mesulam, M.M. (1992) Cytoarchitecture and neural afferents of orbitofrontal cortex in the brain of the monkey. *J. Comp. Neurol.*, **323**, 341–358.

Newsome, W.T., Britten, K.H. & Movshon, J.A. (1989) Neuronal correlates of a perceptual decision. *Nature*, **341**, 52–54.

Niki, H. & Watanabe, M. (1976) Cingulate unit activity and delayed response. *Brain Res.*, **110**, 381–386.

Niki, H. & Watanabe, M. (1979) Prefrontal and cingulate unit activity during timing behavior in the monkey. *Brain Res.*, **171**, 213–224.

Nishijo, H., Ono, T. & Nishino, H. (1988) Single neuron responses in amygdala of alert monkey during complex sensory stimulation with affective significance. *J. Neurosci.*, **8**, 3570–3583.

Norman, D.A. & Shallice, T. (1986) Attention to action: willed and automatic control of behaviour. In Schwartz, G.E., (Ed.), *Consciousness and Self-Regulation*. Plenum Press, New York.

Ono, T., Nishino, H., Fukuda, M., Sasaki, K. & Nishijo, H. (1984) Single neuron activity in dorsolateral prefrontal cortex of monkey during operant behavior sustained by food reward. *Brain Res.*, **311**, 323–332.

Pandya, D.N. & Barnes, C.L. (1987) Architecture and connections of the frontal lobe. In Perecman, E. (Ed.), *The Frontal Lobes Revisited*. The IRBN Press, New York, pp. 41–72.

Passingham, R.E. (1993) *The Frontal Lobes and Voluntary Action*. Oxford University Press, Oxford.

Paus, T. (2001) Primate anterior cingulate cortex: where motor control, drive and cognition interface. *Nature Rev. Neurosci.*, **2**, 417–424.

Petrides, M. & Pandya, D.N. (1994) Comparative architectonic analysis of the human and macaque frontal cortex. In Boller, F. & Grafman, J. (Eds), *Handbook of Neuropsychology*. Elsevier, New York, pp. 17–57.

Petrides, M. & Pandya, D.N. (1999) Dorsolateral prefrontal cortex: comparative cytoarchitectonic analysis in the human and the macaque brain and corticocortical connection patterns. *Eur. J. Neurosci.*, **11**, 1011–1036.

Platt, M.L. & Glimcher, P.W. (1999) Neural correlates of decision variables in parietal cortex. *Nature*, **400**, 233–238.

Rainer, G. & Miller, E.K. (2000) Effects of visual experience on the representation of objects in the prefrontal cortex. *Neuron*, **27**, 179–189.

Roberts, A.C. & Wallis, J.D. (2000) Inhibitory control and affective processing in the prefrontal cortex: neuropsychological studies in the common marmoset. *Cereb. Cortex*, **10**, 252–262.

Rolls, E.T. (1996) The orbitofrontal cortex. *Philos. Trans. R. Soc. Lond. B Biol. Sci.*, **351**, 1433–1443. [Discussion, pp. 1443–1434.]

Romo, R., Brody, C.D., Hernandez, A. & Lemus, L. (1999) Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature*, **399**, 470–473.

Romo, R. & Salinas, E. (2003) Flutter discrimination: neural codes, perception, memory and decision making. *Nature Rev. Neurosci.*, **4**, 203–218.

Rosenkilde, C.E., Bauer, R.H. & Fuster, J.M. (1981) Single cell activity in ventral prefrontal cortex of behaving monkeys. *Brain Res.*, **209**, 375–394.

Sanghera, M.K., Rolls, E.T. & Roper-Hall, A. (1979) Visual responses of neurons in the dorsolateral amygdala of the alert monkey. *Exp. Neurol.*, **63**, 610–626.

Schoenbaum, G., Chiba, A.A. & Gallagher, M. (1998) Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nature Neurosci.*, **1**, 155–159.

Schoenbaum, G., Chiba, A.A. & Gallagher, M. (1999) Neural encoding in orbitofrontal cortex and basolateral amygdala during olfactory discrimination learning. *J. Neurosci.*, **19**, 1876–1884.

Schultz, W. (1986) Activity of pars reticulata neurons of monkey substantia nigra in relation to motor, sensory, and complex events. *J. Neurophysiol.*, **55**, 660–677.

Schultz, W. (2000) Multiple reward signals in the brain. *Nature Rev. Neurosci.*, **1**, 199–207.

Schultz, W., Apicella, P., Scarnati, E. & Ljungberg, T. (1992) Neuronal activity in monkey ventral striatum related to the expectation of reward. *J. Neurosci.*, **12**, 4595–4610.

Schultz, W. & Dickinson, A. (2000) Neuronal coding of prediction errors. *Annu. Rev. Neurosci.*, **23**, 473–500.

Shimamura, A.P. (2000) The role of the prefrontal cortex in dynamic filtering. *Psychobiology*, **28**, 156–167.

Thorpe, S.J., Rolls, E.T. & Maddison, S. (1983) The orbitofrontal cortex: neuronal activity in the behaving monkey. *Exp. Brain Res.*, **49**, 93–115.

Tremblay, L. & Schultz, W. (1999) Relative reward preference in primate orbitofrontal cortex. *Nature*, **398**, 704–708.

Waelti, P., Dickinson, A. & Schultz, W. (2001) Dopamine responses comply with basic assumptions of formal learning theory. *Nature*, **412**, 43–48.

Wallis, J.D., Dias, R., Robbins, T.W. & Roberts, A.C. (2001) Dissociable contributions of the orbitofrontal and lateral prefrontal cortex of the marmoset to performance on a detour reaching task. *Eur. J. Neurosci.*, **13**, 1797–1808.

Watanabe, M. (1996) Reward expectancy in primate prefrontal neurons. *Nature*, **382**, 629–632.

Watanabe, T. & Niki, H. (1985) Hippocampal unit activity and delayed response in the monkey. *Brain Res.*, **325**, 241–254.

Yeterian, E.H. & Pandya, D.N. (1991) Prefrontostriatal connections in relation to cortical architectonic organization in rhesus monkeys. *J. Comp. Neurol.*, **312**, 43–67.