Abstract reward and punishment representations in the human orbitofrontal cortex

J. O'Doherty[1,2], M. L. Kringelbach[1,2], E. T. Rolls[1], J. Hornak[1] and C. Andrews[2]

[1] *Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford OX1 3UD, UK*

[2] *Oxford Centre for Functional Magnetic Resonance Imaging (FMRIB), John Radcliffe Hospital, Headley Way, Oxford OX3 9DU, UK*

Correspondence should be addressed to E.T.R. (Edmund.Rolls@psy.ox.ac.uk)

The orbitofrontal cortex (OFC) is implicated in emotion and emotion-related learning. Using event-related functional magnetic resonance imaging (fMRI), we measured brain activation in human subjects doing an emotion-related visual reversal-learning task in which choice of the correct stimulus led to a probabilistically determined 'monetary' reward and choice of the incorrect stimulus led to a monetary loss. Distinct areas of the OFC were activated by monetary rewards and punishments. Moreover, in these areas, we found a correlation between the magnitude of the brain activation and the magnitude of the rewards and punishments received. These findings indicate that one emotional involvement of the human orbitofrontal cortex is its representation of the magnitudes of abstract rewards and punishments, such as receiving or losing money.

Damage to the OFC (along with other sectors of the ventral and medial prefrontal cortex) in humans is associated with impairments in emotional and social behavior characterized by disinhibition, social inappropriateness and irresponsibility[1–4]. This damage in humans is also associated with impairments in 'gambling' tasks. In these tasks, patients choose from a set of stimuli and, on the basis of the monetary reward and loss obtained following each selection, learn to adaptively choose those stimuli that maximize the overall monetary reward, and minimize the monetary loss obtained in the task[5,6]. A fundamental deficit shown by non-human primates and humans with OFC lesions is a difficulty in reversing behavioral responses to objects associated with rewards and punishments, following a reversal of the reinforcement contingencies[4,7–10]. Such impairments may be due to a representation in the OFC of rewards and punishments that are received, and involvement of the OFC in learning and updating associations between visual stimuli and rewarding and punishing outcomes[3]. Consistent with this hypothesis are findings from single-cell neurophysiological investigations in non-human primates that the reward value of taste, olfactory and visual stimuli is represented in the orbitofrontal cortex, and that some neurons respond only when reinforcement contingencies change[3,11]. However, less is known about the nature of the representations in the human brain for reward and punishment, especially for abstract types of reward and punishment (such as praise or losing money). Previous PET (positron emission tomography) imaging studies have found that the human OFC can be activated with monetary reward[12], and with written positive and negative feedback during performance of a guessing task[13]. However, very little is known about whether abstract punishment, such as losing money, activates the human OFC. If it does, then the question is raised of whether the representations for monetary reward and punishment in the OFC are distinct or overlapping. Further, it is not known how the magnitudes of monetary rewards and punishments received after a choice are represented in the brain; the amount won or lost on any trial clearly is important information for the computation of the overall reward value associated with a particular stimulus.

To address these issues, we used event-related fMRI to investigate the involvement of the OFC while subjects did a variant of a standard reversal learning task (Fig. 1a and b; Methods), in which symbolic monetary gains and losses were used as the rewards and punishments. Subjects were informed in advance that they would not receive any actual remuneration based on the total 'money' accumulated throughout the experiment, but were encouraged to try to do well in the task. (The average total 'money' gained by the subjects at the end of the task was £5061.) The event-related fMRI design used here offered considerable advantages over the previous PET studies described above, as fMRI offers better spatial and temporal resolution, and enables trial-by-trial measurement of the individual hemodynamic responses related to the receipt of each reward or punishment[14]. Our fMRI design therefore was capable of showing whether brain activation in particular regions followed each reward or punishment received. This is in contrast to PET studies, which, because of the long imaging period and block design, cannot show whether any effect is due to some ongoing general effect of being in the reward condition, to anticipation of reward, or to the reward actually received. In addition, our event-related design enabled us to correlate the BOLD signal with the individual rewards and punishments received, to determine where in the brain the magnitudes of the rewards and punishments are represented. To achieve this, the monetary gains and losses were distributed according to differing probabilistic reinforcement schedules and ranges of magnitude for the rewarding stimulus (S+) and punishing stimulus (S–). An important feature of this task was that 'money' could be won or lost on both the S+ and the S–, but choosing the S+ gave larger rewards and smaller punishments, whereas the converse

a

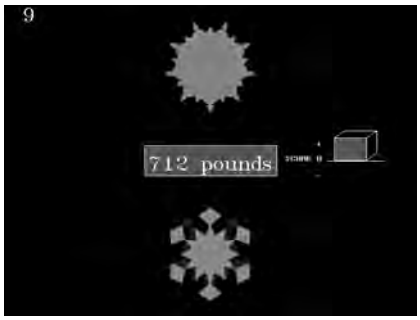

b

	S+	S–
Reward:punishment ratio	70:30	40:60
Reward ranges	£80–250	£30–60
Punishment ranges	£10–60	£250–600

c

		Reward	Punishment
Acquisition phase	S+	Reward acquisition [RewAcq]	Punishment acquisition [PunAcq]
Reversal phase	Old S+ [new S–]	Reward reversal [RewRev]	Punishment reversal [PunRev]

Fig. 1. Experimental design. (**a**) Two unfamiliar and easily discernable fractal patterns were displayed on the screen, with the total score displayed numerically and graphically as a bar chart. After a subject had selected a stimulus, a flashing text message was superimposed above the selected stimulus, indicating how much money the subject had won or lost. The two fractals were randomly assigned to the top or the bottom of the screen on each trial. At the beginning of the task, one of the stimuli was designated as the S+. Whether a trial was rewarded or punished was decided on the basis of a pseudo-random sequence according to the ratios shown in (**b**). The magnitude of the rewards and punishments were also varied using a uniform distribution between the ranges shown in the same figure. The specific ratios and ranges used were chosen on the basis of pilot experiments to give the task an appropriate level of difficulty. Consistent selection of the current S+ resulted in an overall monetary gain, whereas consistent selection of the S– resulted in an overall monetary loss. The criterion for successful acquisition of the S+ in the task was four of the previous five consecutive responses to the S+. After between one and three trials following criterion (randomly determined), the reward contingencies were reversed so that the old S+ became the new S– and the old S– became the new S+. The subject's task was, by trial and error, to determine which stimulus was the more profitable to choose, and to keep track of this when a reversal occurred. The different event types used in the statistical comparisons are shown in (**c**).

was true of the S–. The subjects' task was, by trial and error, to determine which stimulus was the more profitable to choose, to keep track of the profitability, and to reverse their choice when we reversed the reward and punishment contingencies. The specific ratios and ranges of rewards and punishments were chosen on the basis of pilot experiments; the task was given an appropriate level of difficulty, sensitive enough to detect impairments in patients with OFC damage (J.O., J.H., E.T.R., D. Wade and J. McGrath, unpublished data). The task was designed with probabilistic reinforcement contingencies to make it less likely that subjects would use alternative cognitive strategies, such as use of explicit verbal coding ("stimulus *x* gives reward"), to solve the task. During the same scanning session, subjects also participated in a control task in which they made similar responses without assurance of monetary reward or punishment.

Results

The principal events of interest (Fig. 1c) were as follows: first, a large reward obtained during the acquisition phase, when the subjects were consistently touching the S+ (RewAcq); second, a small punishment obtained during the acquisition phase (PunAcq); third, a large punishment obtained during the reversal phase when the contingencies have reversed (PunRev); fourth, a small reward obtained during the reversal phase (RewRev). Statistical comparisons of the BOLD activations related to these different events were based on a group random effects with a statistic threshold at $p < 0.005$, which followed *t*-tests done for each subject, for each comparison (Methods).

Comparisons considered first are those that show the brain areas more activated following a rewarding outcome than a punishing outcome. When the reward acquisition events were compared to the punish reversal events (RewAcq – PunRev), voxels were found to survive the statistical threshold in the medial OFC bilaterally (Fig. 2a; this was the case in all nine of the individual subjects). We also found significantly activated voxels in the medial OFC when comparing the reward acquisition events to the punishment acquisition events (RewAcq – PunAcq), demonstrating that the activation was a genuine response to the receipt of the reward, rather than a response to the differences between the acquisition and reversal phases of the task. The averaged percent change time course of significantly activated voxels in the medial OFC (Fig. 3a) showed a clear increase in the BOLD signal following a reward, with also some evidence of a decrease in the BOLD signal following a punishment. To determine whether the activation in the medial OFC was related to the magnitude of the reward obtained, we measured the correlation between the magnitude of the rewards and punishments on each trial and the corresponding BOLD signal on each trial in each voxel (see Methods). The group analysis (Fig. 4) showed that left medial OFC activation was correlated with the magnitude of the reward obtained. (Significantly activated clusters of voxels showed this effect in the individual subject analyses for six of the nine subjects, at a threshold of $p < 0.005$, and an extent threshold of $p < 0.05$.) This finding indicates that an increase in neural activity (as measured by the BOLD signal) in the medial OFC is particularly related to the receipt of a reward after selecting a stimulus, and that the magnitude of this activity increases with the magnitude of the reward. To examine the nature of this correlation further, we show in Fig. 5a the mean percent change of the BOLD signal for different magnitudes of reward or punishment in the most significant voxels in the subtraction analysis shown in Fig. 2. It is evident from Fig. 5a (for the peak RewAcq – PunRev voxels across subjects) that there was a graded increase in these voxels according to the magnitude of the reward, and a graded decrease according to the magnitude of the punishment. Indeed, the peak voxels from the subtraction analysis also showed a significant correlation between the peak BOLD signal on each trial and the magnitude of the reward obtained (significant at $p < 0.005$ in all 6 subjects that showed significant effects in the correlation analysis). Thus, the single-event correlation analysis showed that voxels in the medial OFC had BOLD changes related to the magnitude of the reward obtained (Fig. 4).

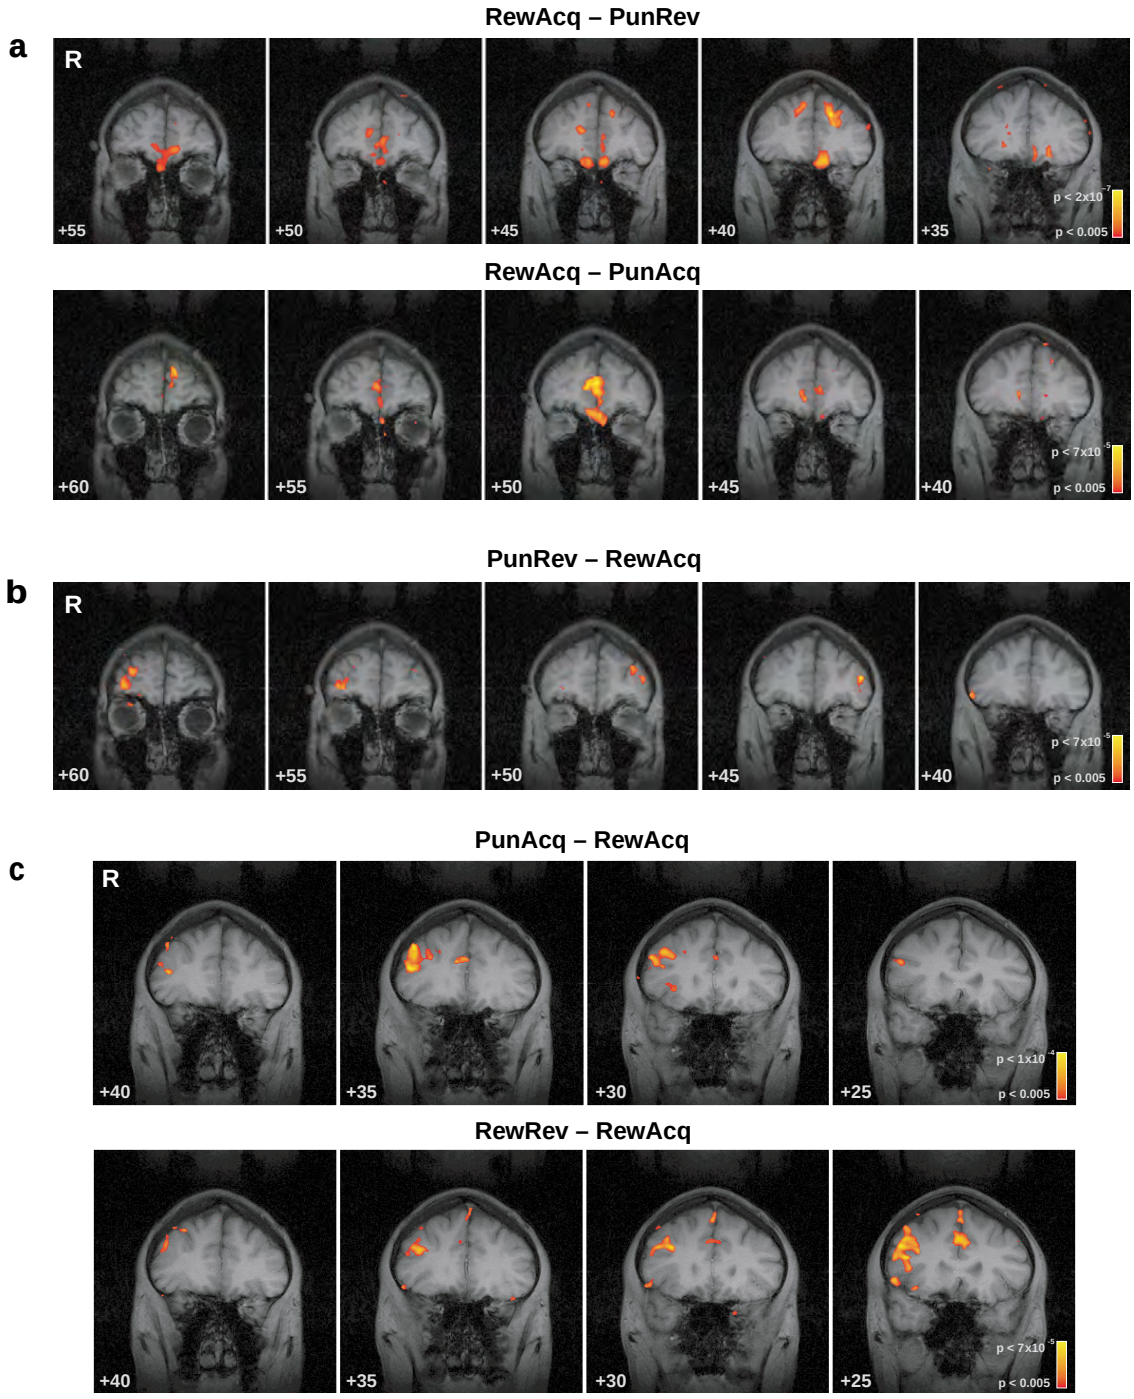

Fig. 2. The results from the group random effects analysis across the nine subjects with a threshold at $p < 0.005$ (uncorrected) superimposed on the structural MRI scan of a single subject. The equivalent *y*-values in Talairach space for each coronal section shown is provided on the bottom left of each frame. (**a**) Areas of bilateral medial OFC significantly activated by reward, as shown in the RewAcq – PunRev comparison. Talairach coordinates of activated clusters in the OFC, [*x*, *y*, *z*], [6, 42, –24], [–10, 40, –22], [–4, 28, –18] and the RewAcq – PunAcq comparison [–4, 48, –28]. (**b**) An area of right lateral OFC activated by punishment is shown in the PunRev – RewAcq comparison [36, 58, –12]. (**c**) Right inferior prefrontal sulcus (BA 44/45) significantly activated in the PunAcq – RewAcq comparison [48, 33, 23] and RewRev – RewAcq comparison [45, 23, 23]. The activation extends dorsally to include parts of areas 9/46.

We next describe comparisons that show the brain areas that are more activated following a punishing outcome than a rewarding outcome. We compared the punish reversal events to the reward acquisition events (PunRev – RewAcq), and saw activation in a lateral area of anterior OFC (BA 10/11) as well as in a region of nearby ventral prefrontal cortex. The area activated was in the right hemisphere in the group analysis (Fig. 2b), and on the right in eight of nine subjects in the individual subject analyses (at a threshold of $p < 0.005$, extent threshold of $p < 0.05$). In five of those eight subjects, activation was also

found on the left, and in the ninth subject, the activation was on the left only. The time course of significantly activated voxels in this region is shown in Fig. 3b. In addition, a correlation analysis, in which voxels were selected that were positively correlated with the magnitude of the punishments received, revealed an adjacent area of the lateral OFC (and ventral prefrontal cortex) where the BOLD signal was correlated with the magnitude of the punishment obtained (Fig. 4, group analysis; the same correlation existed in eight of nine subjects in the individual subject analyses at $p < 0.005$, with a $p < 0.05$ extent threshold). These results indicate that an increase in the neural activity of the lateral OFC is particularly related to the subject's receipt of punishment after selecting a stimulus, and that the magnitude of activation in this region is related to the magnitude of the punishments obtained. Confirmation of this finding comes from the data indicating that, in the peak voxels across subjects from the PunRev – RewAcq subtraction (Fig. 2b), there was a graded increase in these voxels according to the magnitudes of the punishment, and a graded decrease according to the magnitude of the reward (Fig. 5b). Indeed, the peak voxels from the subtraction analysis also showed a significant correlation between the peak BOLD signal in the lateral OFC area on each trial and the magnitude of the punishment obtained (significant at $p < 0.01$ in 5 subjects and at $p < 0.05$ in 3 other subjects). Thus, the single-event correlation analysis showed that voxels in the lateral OFC had BOLD changes that were related to the magnitude of the punishment obtained (Fig. 4).

We conducted a control condition in which the subjects had to choose one of two fractal images presented in the same format as the reversal task, but without receiving any monetary reward or punishment as feedback following the stimulus selection (see Methods). The aims of this task were to provide an affectively neutral baseline condition, with the same response selection and motor components as in the reversal task, but without the monetary feedback. In this analysis, we compared the activations during all of the trials in which rewards were received (RewAcq and RewRev trials) and the activations on the control trials, to reveal areas activated by reward relative to a neutral baseline (Reward – Control). Seven of the nine subjects participated in this condition. (Two subjects were unable to participate because of their time constraints.) Consistent with the previous results, an area of medial OFC was activated in this contrast (Fig. 6a). Similarly, when all of the trials in which the subjects received a monetary loss (PunAcq and PunRev trials) were subtracted from the neutral baseline (Punish–Control), an area of lateral orbitofrontal cortex was activated bilaterally (Fig. 6b). The results of the comparisons with the control task support the findings described above.

Other brain areas

Activation of some other brain areas was also found to be related to performance of the task. One such area included parts of the medial prefrontal cortex (BA 10/32), which were significantly activated in contrasts revealing brain activation to reward (Fig. 2a). A second region, the cortex in the right posterior inferior prefrontal sulcus (BA 44/45), showed most activation after trials in which the subject either received a small punishment on the currently correct stimulus (PunAcq), or a small reward for touching the currently incorrect stimulus (RewRev; Fig. 2c). A third activated region was the dorsal anterior cingulate cortex (BA 24/32), where effects were found in contrasts that revealed activation to punishment as well as in the RewRev – RewAcq contrast.

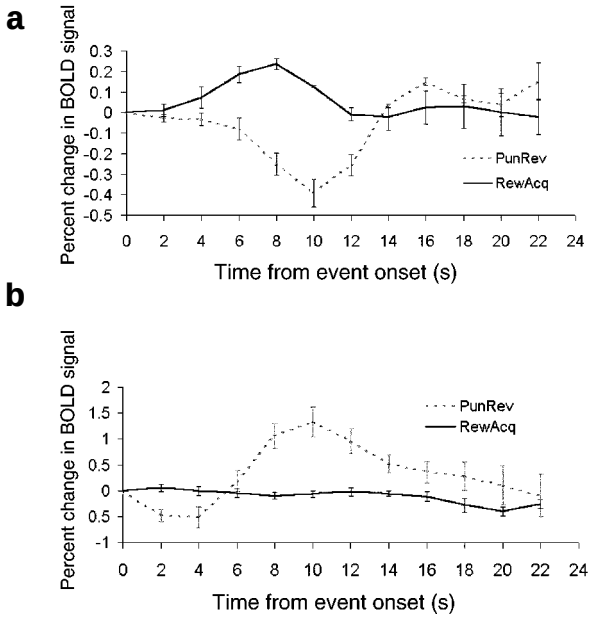

Fig. 3. Typical averaged time course relative to the time of the behavioral response made to select one of the stimuli from a single subject of significantly activated voxels from (**a**) the RewAcq – PunRev contrast in the medial OFC and (**b**) the PunRev – RewAcq contrast in the lateral OFC.

Discussion

The results from this study indicate that the lateral and medial orbitofrontal cortical areas are together responsive to the rewarding and punishing outcomes produced by selecting a stimulus. The lateral area of the OFC is activated following a punishing outcome, and the medial OFC is activated following a rewarding outcome. The results advance earlier understanding of the orbitofrontal cortex obtained, for example, by neurophysiology in the macaque, or by PET neuroimaging in humans. The results show the following. Abstract rewards and punishments can activate the orbitofrontal cortex. Different areas are activated by these rewards and punishers. The activations are produced by the rewards and punishments given (in that they follow the delivery of each reward and punishment, a finding made possible by the fMRI design). The magnitude of the activations reflects the magnitude of the reward or punishment delivered (a finding made possible by the event-related aspect of the fMRI design, which enabled correlations to be measured). The medial region that showed increased activation to reward also showed decreased BOLD signal when punishment was delivered, and *vice versa* for the lateral orbitofrontal cortex region.

The dissociation found in this study between the lateral and medial orbitofrontal cortex may be related to anatomical and neuropsychological differences between these regions in macaques. The lateral orbitofrontal cortex receives most of the sensory afferents[15–17], whereas the medial orbitofrontal cortex has the strongest connections with the rostral cingulate cortex. The intrinsic connectivity of the macaque lateral orbitofrontal cortex is such that areas within this region project strongly to each other, but very weakly to areas in the medial orbitofrontal cortex, whereas areas in the medial orbitofrontal cortex connect strongly with each other and weakly to lateral orbital areas[18]. On these grounds, those authors draw a firm distinction between the medial and lateral

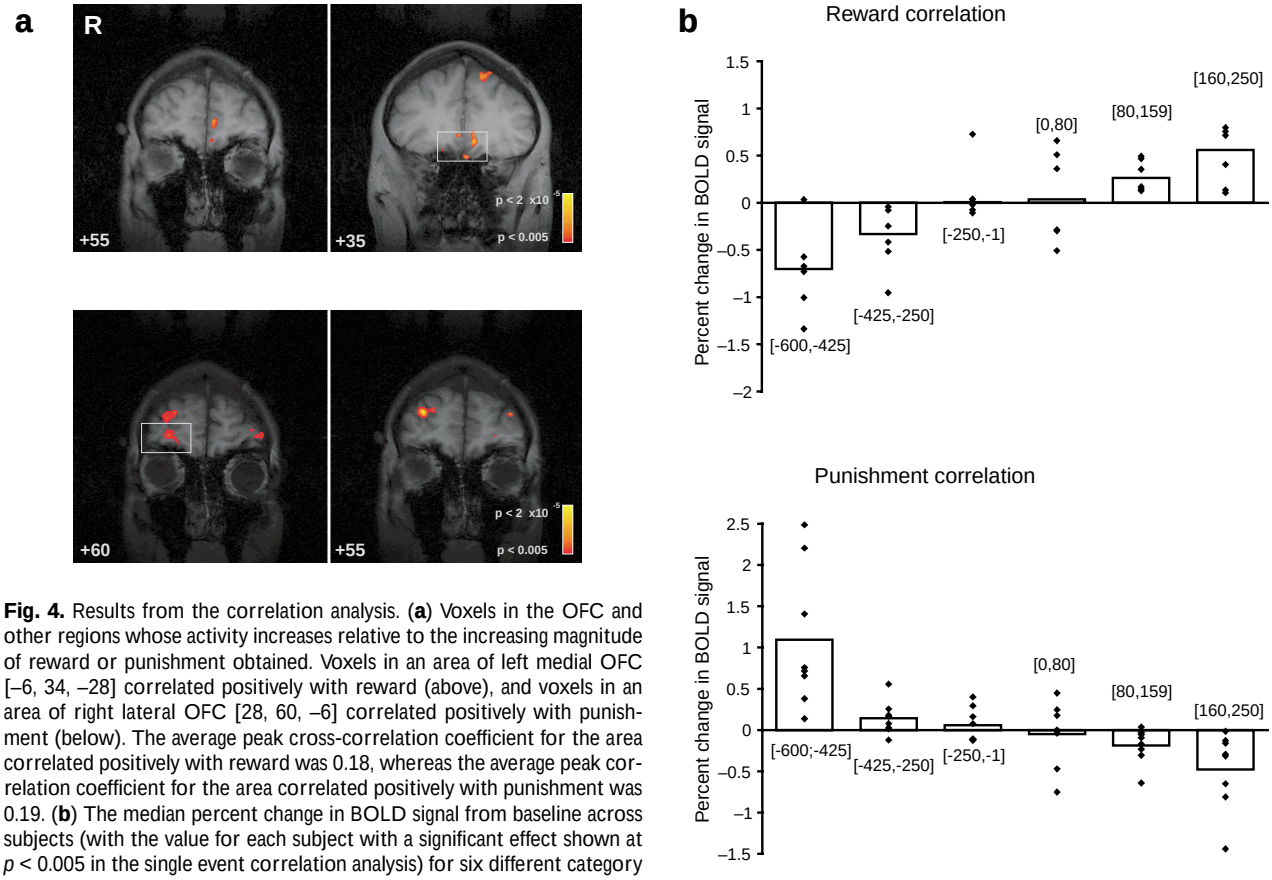

Fig. 4. Results from the correlation analysis. (**a**) Voxels in the OFC and other regions whose activity increases relative to the increasing magnitude of reward or punishment obtained. Voxels in an area of left medial OFC [–6, 34, –28] correlated positively with reward (above), and voxels in an area of right lateral OFC [28, 60, –6] correlated positively with punishment (below). The average peak cross-correlation coefficient for the area correlated positively with reward was 0.18, whereas the average peak correlation coefficient for the area correlated positively with punishment was 0.19. (**b**) The median percent change in BOLD signal from baseline across subjects (with the value for each subject with a significant effect shown at $p < 0.005$ in the single event correlation analysis) for six different category ranges of reward and punishment. The signal was averaged across a category range within each subject and then the average signal change from each category was averaged across subjects. This is plotted for the peak voxels in the medial OFC (above), which significantly correlated with reward, and in the lateral OFC (below), which significantly correlated with punishment. The ranges of monetary reward and punishment in each category are shown on the chart, and were determined by their relative frequencies, which follow from the experimental design. For the time series cross-correlation data shown here, the average of the Pearson correlation coefficient from each subject between the binned magnitudes of reward and the corresponding BOLD signals was 0.66, and the average of the Pearson correlation coefficient from each subject between the binned magnitudes of punishment and the corresponding BOLD signals was 0.79.

orbitofrontal cortex and suggest that the two areas are separate networks. In a neuropsychological investigation in macaques, differences were reported between the effects of lesions in different parts of the orbitofrontal cortex during emotion-related visual discrimination reversal learning[7]. Following a reversal, macaques with lesions in the lateral orbitofrontal cortex and inferior convexity continue responding to the previously rewarded but now unrewarded stimulus. In contrast, monkeys with medial OFC lesions make more errors during acquisition before criterion, and do not show the same perseverative responding following reversal. This dissociation in lesion effects could be explained by the results from the current study; if the lateral orbitofrontal cortex is more involved in representing the punishing consequences of continuing to select the previously rewarded stimulus following reversal, then this could account for a difficulty in reversing responses to the old S+ following reversal. If, alternatively, the medial orbitofrontal cortex is involved in representing the rewarding consequences of choosing the currently rewarded S+ during acquisition, then damage to this region could result in a failure to consistently choose the currently rewarded stimulus evident, for example, during acquisition.

The finding of a correlation between the size of the reward or punishment obtained on an individual trial and the magnitude of the BOLD response in the medial and lateral orbitofrontal cortical areas, respectively, is evidence that the magnitude of the rewards and punishments received are represented in those regions. This finding may account in part for the deficits shown by patients with OFC lesions during gambling tasks, in which the patients have to decide to choose advantageously from a number of choices, to maximize monetary gain[5]. An inability to represent the magnitude of the rewards and punishments received may then result in a difficulty in judging the degree to which a particular stimulus choice is advantageous on the basis of cumulative monetary gain.

The finding of inferior prefrontal sulcus activation during RewRev and PunAcq trials suggests that this area may be more activated during trials that are more ambiguous in their potential for signaling a change in the reinforcement contingencies. The activated area was very similar to the activated area seen in a number of imaging studies featuring tasks associated with prefrontal function, such as in the cognitive set-shifting component of the Wisconsin card sorting task[19], and in the no-go trials of a go/no-go task[20]. Taken together, the findings suggest that the cortex in the inferior prefrontal sulcus may be involved in inhibiting inappropriate behavioral strategies, such as to switch behavior after a small punishment or to stay after a small reward.

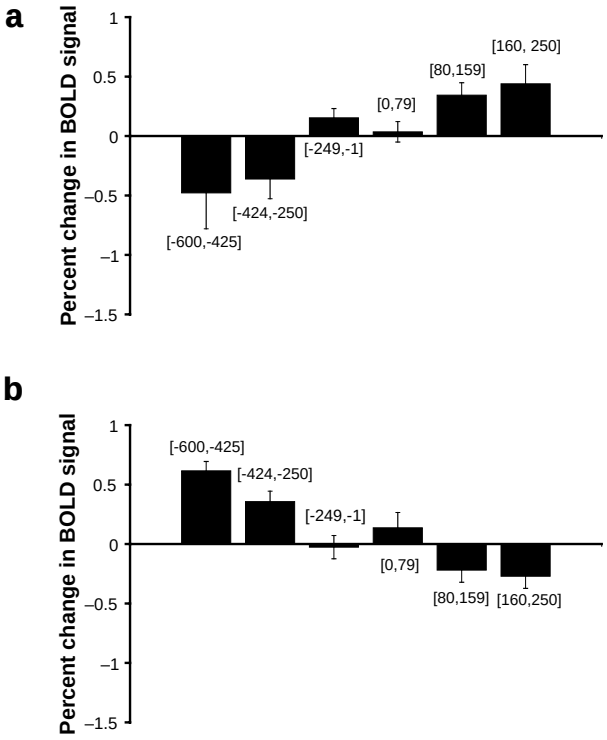

Fig. 5. The mean percent change in BOLD signal from baseline across subjects for six different category ranges of reward and punishment, shown for the peak voxels from the subtraction analysis. (**a**) Percent change in BOLD signal for the RewAcq – PunRev contrast. (**b**) Percent change in BOLD signal for the PunRev – RewAcq contrast. The signal was averaged across a category range within each subject, and the average signal change from each category was then averaged across subjects. The ranges of monetary reward and punishment in each category are shown on the chart, and were determined by their relative frequencies, which follow from the experimental design. The means and standard errors of the means are shown. For the subtraction analysis shown here, the average of the Pearson correlation coefficient from each subject, between the binned magnitudes of reward and the corresponding BOLD signals, was 0.61, and the average of the Pearson correlation coefficient from each subject, between the binned magnitudes of punishment and the corresponding BOLD signals, was 0.75.

Elliot and colleagues[21] reviewed results of functional imaging experiments conducted in their laboratory in which they also report a dissociation between the medial and lateral orbitofrontal cortex. They suggest, on the basis of findings in diverse tasks including sentence completion, story comprehension, guessing and delayed match to sample, that the medial orbitofrontal cortex is involved in monitoring and holding in mind reward values. Our direct evidence on this issue indicates that the medial orbitofrontal cortex was activated by monetary reward, and showed less BOLD signal relative to baseline following the punishment of monetary loss. On the basis of findings in tasks including delayed non-match to sample, sentence completion, guessing and hypothesis testing, these authors suggest that the lateral orbitofrontal cortex is likely to be activated when a response previously associated with reward has to be suppressed. Our evidence showed that the lateral orbitofrontal cortex was activated by punishment, and that it indeed reflected the magnitude of the punishment, with some additional evidence for a decrease in activation following a reward. A region sensitive to the magnitude of punishment could provide a useful signal for response suppression. However, in macaques, neurophysiological evidence[22] suggests that sensory stimuli that produce rewards and punishments, rather than the motor responses being made, are represented in the orbitofrontal cortex, and we therefore suggest that the response suppression itself is likely to be mediated outside the orbitofrontal cortex.

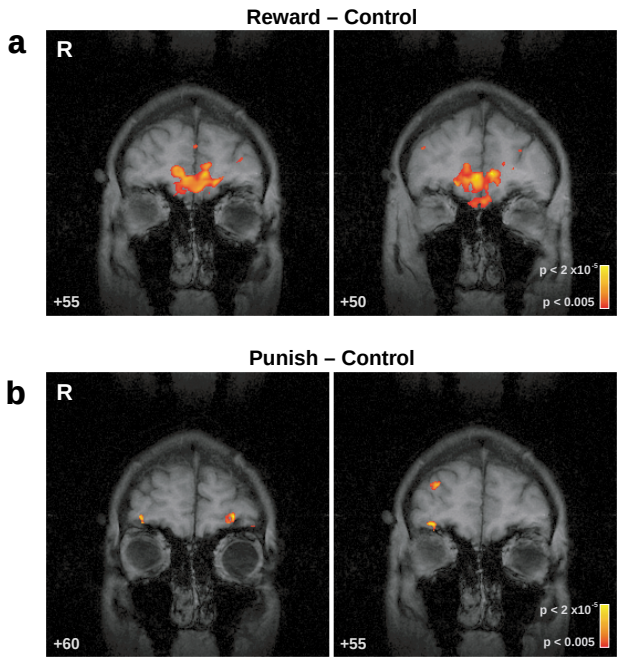

The time to peak of the activation in the orbitofrontal cortex (approximately 10 seconds) was later than the time to peak often found by fMRI investigations in other brain areas such as the primary visual or somatosensory cortex, in which a period between four to six seconds is commonly reported[23]. The relatively long latency of the BOLD response in the orbitofrontal cortex observed here has been noted previously. In particular, using a somatosensory stimulation protocol, we observed that the activation in the orbitofrontal cortex peaked much later than the activation in the primary somatosensory cortex[24]. A similar effect has also been observed in the anterior inferior prefrontal cortex[25]. Possible reasons for the slow hemodynamic response in the OFC include[24] the relatively low peak firing rates of neurons in the OFC (frequently in the range of 10–15 spikes/s[26,27], which contrast with frequent peak responses of 60–120 spikes/s in the temporal lobe cortical visual areas[28]), and the sparseness of the representations found in the OFC[26,29]. A full understanding of this effect may only be possible when the relationship between the underlying neural firing rate and the characteristics of the BOLD response are more completely understood[30].

Fig. 6. Comparison of rewards and punishments with control condition. (**a**) A region of bilateral medial OFC [0, 44, 26] and medial prefrontal cortex significantly activated relative to a neutral baseline in the Reward – Control contrast (**b**) A region of bilateral lateral OFC significantly activated relative to a neutral baseline in the Punish – Control contrast. Right OFC, [34, 52, –12]; left OFC, [–28, 64, –8].

A much discussed issue is the extent to which there is hemispheric specialization in the processing of rewards and punishments[31]. In the group analysis, the voxels that correlated with reward were predominantly left-sided, whereas the voxels correlated with punishment were predominantly on the right. However, bilateral activation in the reward and punishment correlation analyses was commonly seen in the individual subjects (five subjects in reward analysis and four subjects in punishment analysis). Furthermore, when the reward and punishment trials were subtracted from a neutral baseline, activation in the group analysis was found bilaterally in the OFC for both reward and punishment. Consequently, the most consistent separation between the brain areas responding to reward and punishment in this study were between lateral and medial areas of the OFC, not the left and right OFC.

Other areas activated in the task include parts of the medial prefrontal cortex (area 10) and the dorsal anterior cingulate. These results indicate that several prefrontal regions are activated during performance of the task. However, the evidence from lesion studies suggests that the OFC, not other regions of prefrontal cortex, is essential for performance on an emotion-related reversal learning task[9,32]. Furthermore, single-cell neurophysiological recordings in non-human primates emphasize the importance of OFC neurons in the responses to rewards and punishments[11,29,33].

These results are consistent with the hypothesis that the deficits in emotion-related learning tasks following OFC lesions are due to an impairment in using information from rewarding and punishing outcomes to guide behavioral choice. The finding that the magnitude of the activations depended on the size of the reward or punishment received is consistent with neurophysiological evidence from primates that rewarding and punishing stimuli are represented in the OFC, but this finding does not exclude a possible representation in other brain areas of rewards and punishments[3,29]. Our evidence that there is a representation for rewards and punishments in these regions is also consistent with the hypothesis that the changed emotional and social behavior demonstrated by humans with this damage is due to a fundamental impairment in the processing of rewards and punishments[3].

METHODS

Subjects and task. Nine healthy, right-handed subjects participated in this study (6 female and 3 male; 22–30 years old; average age, 25.4). Before scanning, the subjects were trained on a version of the task designed for use outside the scanner, as part of a neuropsychological study of patients with OFC lesions (J.O., J.H., E.T.R., D. Wade and J. McGrath, unpublished data). This task consisted of two parts: a simple acquisition task, where the subject had to find which of two stimuli was the S+ and touch that stimulus consistently, and a reversal task similar to the one used in the scanner. Previously unseen fractal images were used in the scanner, so that learning was required during the imaging. The subjects were trained on 50 trials of the reversal task outside the scanner (or until they had attained at least two reversals). In the scanner, stimuli were generated on a PC, and appeared on a back-projection screen mounted outside the scanner bore, which subjects viewed using mirror glasses. Subjects used their right hand to press one of two buttons corresponding to the top or bottom of the screen, to select one of the two fractals. The subjects did the task (described fully in Fig. 1) for 30 min, during which they completed an average of 136 trials and 15 reversals. After selection of one of the stimuli, the amount won or lost on a stimulus was displayed for 3 s; then the screen was cleared, and a fixation cross was presented for an additional 10 s, before the presentation of the next trial. Seven subjects also participated in a control condition in the same scanning session, in which two novel fractals were presented in an identical fashion to the reversal task, and the subjects were requested to consistently select one of the two fractals. (They were told in advance which stimulus to select.) After selection of one of the stimuli, the subjects were not provided with any monetary reward or punishment but instead were given neutral feedback, which consisted of the words "pattern selected." The screen was then cleared, and a fixation cross was shown before the presentation of the next trial. The subjects did the control task for 5 min, during which they completed an average of 18 trials.

Scanning procedures. Images were acquired with a 3.0-T VARIAN/SIEMENS (Palo Alto, California) whole-body scanner at Oxford Centre for functional magnetic resonance imaging (FMRIB, Oxford, UK). Fourteen T2* weighted EPI slices were acquired every 2 s (TR = 2). The following parameters were carefully selected in order to minimize susceptibility and distortion artifact in the orbitofrontal cortex. First, the data were acquired in a coronal rather than axial slicing direction, as this aligned the slices to be perpendicular to the predominant direction of the intrinsic susceptibility-induced field gradients, and helped to minimize through-plane dephasing. Second, the voxel resolution was minimized by using 3-mm in-plane resolution and a 5-mm slice thickness, which results in less phase cancellation than would be produced by lower voxel resolutions. Third, a relatively low TE of 25 ms was selected to decrease the signal dropout, as less phase dispersion is created across the voxels. Fourth, each subject was individually shimmed using both linear and second-order shimming to minimize static field inhomogeneities in the orbitofrontal cortex. Finally, geometric distortion was minimized by using a specialist head insert gradient coil (Magnex SGRAD III, Abingdon, UK) with a relatively high gradient switching frequency of 960 Hz.

The matrix size was 64×64, and the field of view was 192×192 mm. Continuous coverage was obtained from +60 (A/P) to –10 (A/P). Acquisition was done during task performance, which lasted a total of 30 min, yielding 900 volumes in total. A whole-brain T2* weighted EPI volume of the above dimensions and an anatomical T1 volume with slice thickness of 5 mm and in-plane resolution of 0.75×0.75 mm were also acquired. The acquisition protocol used in this study is similar to that used in previous studies from our laboratory, in which we used fMRI to image the orbitofrontal cortex[24,34].

Image analysis. The image analysis was carried out in MEDx (Sensor Systems, Sterling, Virginia). The datasets were motion corrected using a three-dimensional automated image registration algorithm, spatially smoothed using a Gaussian filter (full width at half maximum, 5 mm), intensity normalized, and temporally filtered. (None of the nine subjects were excluded from the analysis because of excessive motion.) The motion across all nine subjects ranged from 0.31–0.73 mm in the *x*-direction, 0.26–1.84 mm in the *y*-direction and 0.15–1.61 mm in the *z*-direction, and we verified that the three-dimensional automated image registration algorithm corrected for this small degree of motion, which was, in any case, small in relation to the voxel size. Further, the pattern of results could not be due to stimulus-correlated motion artifact, as the time course of the OFC activations found shows a peak at around 8–14 s, which is consistent with a normal hemodynamic response. If the results were related to any stimulus-induced motion artifact, it would be anticipated that the peak of the response would occur at the trial onset, and not 8–14 s afterward[35]).

The trials were designated as one of four event types (Fig. 1c). Events from the acquisition phase were selected from the second trial, after subjects had changed their choice of stimulus to the new S+, provided that the subjects consistently chose that S+ until the next reversal (denoting stable acquisition of the S+). Events from the reversal phase were selected after a reversal of the reinforcement contingencies, when the subject touched the old S+ and before the subjects had successfully switched their choice of stimulus to the new S+. Three volumes were selected from each trial corresponding to 8–14 seconds after the trial onset, as this represented the peak of the hemodynamic response in the OFC. The selected volumes from each trial were then used in the *t*-tests between each of the events of interest for each individual subject. A statistical threshold of $p < 0.005$ with an extent threshold of $p < 0.05$ was applied to the individual *z*-maps resulting from the statistical contrasts using the MEDx cluster detection function. For the analysis involving the control condition, the volumes from the control condition were registered to the first volume in the reversal task, and intensity normalized to the same glob-

al mean as the volumes from the reversal task. Further, we applied spatial and temporal filtering using the same parameters as those described for the reversal task above. For the statistical comparisons, three volumes were selected from each trial in the control condition corresponding to 8–14 s after the trial onset, and *t*-tests were done between the selected volumes from the reward and punishment events and the control events. The reward events used in the Reward – Control contrast consisted of all the events during the reversal task in which a reward was received, irrespective of the trial type (that is, the trials were pooled across both RewAcq and RewRev events). Similarly, the punishment events used in the Punish – Control contrast consisted of all trials during the reversal task in which a punishment was received (pooled across both PunRev and PunAcq events). The individual *z*-maps for each subject were registered to that subject's corresponding anatomical volume. The mean difference images for each subject from every comparison were registered to the anatomical volume from a single subject and random effects statistics were calculated for each of the contrasts (using a locally developed FMRIB tool). A threshold for statistical significance of $p < 0.005$ (uncorrected) was applied to the random effects *z*-maps. The statistical criteria incorporated in the random effects methods and used in this investigation are intrinsically robust, and within this context, the use of uncorrected statistics is reasonable, but is in any case further justified on the basis that the main area of interest is the orbitofrontal cortex, a region in which activation was predicted on the basis of our *a priori* hypothesis. Talairach coordinates for each of the activated clusters were obtained following registration of the single subject anatomical to a standard brain in the coordinate space of Talairach and Tournoux[36].

For the correlation analysis, a time-series waveform was constructed, which represented the magnitude of the rewards and punishments received during the experiment. The waveform was constructed starting with a simple ON-OFF square block waveform in which the magnitude in a 6-s ON period after the trial onset (allowing 8 s for hemodynamic lag) was set to the amount of monetary gain or loss obtained on each individual trial (positive for reward, negative for punishment), and the OFF period magnitude was set to zero. The resulting square waveform was convolved using a Poisson distribution (with $\lambda = 1.0$) to give a closer approximation to a real hemodynamic response. This waveform was then correlated on a voxel-by-voxel basis with the actual hemodynamic responses across the experiment for each of the subjects using the MEDx correlation analysis function. A group random effects statistic was calculated from a combination of the individual correlation maps, and a statistical threshold of $p < 0.005$ (uncorrected) was applied.

Note: supplementary information is available on the Nature Neuroscience web site (http://neurosci.nature.com/web_specials/).

Acknowledgements

This research is funded by the MRC (to FMRIB), the MRC (to E.T.R.), the MRC IRC for Cognitive Neuroscience, the MRC, Oxford McDonnell- Pew Centre for Cognitive Neuroscience and Firmenich (to J.O.D.) and the Danish Research Agency, Aarhus, Denmark (to M.L.K.). We thank P. Matthews, S. Smith, P. Jezzard, S. Clare and I. Tracey for their help.

Received 5 June; Accepted 20 October 2000

1. Damasio, A. R. *Descartes' Error* (Putnam, New York, 1994).
2. Anderson, S. W., Bechara, A., Damasio, H., Tranel, D. & Damasio, A. R. Impairment of social and moral behavior related to early damage in human prefrontal cortex. *Nat. Neurosci.* 2, 1032–1037 (1999).
3. Rolls, E. T. *The Brain and Emotion* (Oxford Univ. Press, Oxford, 1999).
4. Rolls, E. T., Hornak, J., Wade, D. & McGrath, J. Emotion-related learning in patients with social and emotional changes associated with frontal lobe damage. *J. Neurol. Neurosurg. Psychiatry* 57, 1518–1524 (1994).
5. Bechara, A., Damasio, A. R., Damasio, H. & Anderson, S. W. Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition* 50, 7–15 (1994).
6. Bechara, A., Damasio, H., Tranel, D. & Damasio, A. R. Deciding advantageously before knowing the advantageous strategy. *Science* 275, 1293–1295 (1997).
7. Iversen, S. D. & Mishkin, M. Perseverative interference in monkeys following selective lesions of the inferior prefrontal convexity. *Exp. Brain Res.* 11, 376–386 (1970).
8. Jones, B. & Mishkin, M. Limbic lesions and the problem of stimulus–reinforcement associations. *Exp. Neurol.* 36, 362–377 (1972).
9. Meunier, M., Bachevalier, J. & Mishkin, M. Effects of orbital frontal and anterior cingulate lesions on object and spatial memory in rhesus monkeys. *Neuropsychologia* 35, 999–1015 (1997).
10. Freedman, M., Black, S., Ebert, P. & Binns, M. Orbitofrontal function, object alternation and perseveration. *Cereb. Cortex* 8, 18–27 (1998).
11. Rolls, E. T., Critchley, H. D., Mason, R. & Wakeman, E. A. Orbitofrontal cortex neurons: role in olfactory and visual association learning. *J. Neurophysiol.* 75, 1970–1981 (1996).
12. Thut, G. *et al.* Activation of the human brain by monetary reward. *Neuroreport* 8, 1225–1228 (1997).
13. Elliott, R., Frith, C. D. & Dolan, R. J. Differential neural response to positive and negative feedback in planning and guessing tasks. *Neuropsychologia* 35, 1395–1404 (1997).
14. Buckner, R. L. Event-related fMRI and the hemodynamic response. *Hum. Brain Mapp.* 6, 373–377 (1998).
15. Cavada, C., Company, T., Tejedor, J., Cruz Rizzolo, R. J. & Reinoso Suarez, F. The anatomical connections of the macaque monkey orbitofrontal cortex. A review. *Cereb. Cortex* 10, 220–242 (2000).
16. Carmichael, S. T. & Price, J. L. Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *J. Comp. Neurol.* 363, 615–641 (1995).
17. Morecraft, R. J., Geula, C. & Mesulam, M. -M. Cytoarchitecture and neural afferents of orbitofrontal cortex in the brain of the monkey. *J. Comp. Neurol.* 232, 341–358 (1992).
18. Carmichael, S. T. & Price, J. L. Connectional networks within the orbital and medial prefrontal cortex of macaque monkeys. *J. Comp. Neurol.* 371, 179–207 (1996).
19. Konishi, S. *et al.* Transient activation of inferior prefrontal cortex during cognitive set shifting. *Nat. Neurosci.* 1, 80–84 (1998).
20. Konishi, S., Nakajima, K., Uchida, I., Sekihara, K. & Miyashita, Y. No-go dominant brain activity in human inferior prefrontal cortex revealed by functional magnetic resonance imaging. *Eur. J. Neurosci.* 10, 1209–1213 (1998).
21. Elliott, R., Dolan, R. J. & Frith, C. D. Dissociable functions in the medial and lateral orbitofrontal cortex: evidence from human neuroimaging studies. *Cereb. Cortex* 10, 308–317 (2000).
22. Rolls, E. T. in *The Frontal Lobes* (eds. Stuss, D. T. & Knight, R. T.) (Oxford Univ. Press, Oxford, in press).
23. Miezin, F. M., Maccotta, L., Ollinger, J. M., Petersen, S. E. & Buckner, R. L. Characterizing the hemodynamic response: effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. *Neuroimage* 11, 735–759 (2000).
24. Francis, S. *et al.* The representation of the pleasantness of touch in the human brain, and its relation to taste and olfactory areas. *Neuroreport* 10, 453–459 (1999).
25. Buckner, R. L. *et al.* Functional-anatomic study of episodic retrieval: II. Selective averaging of event-related fMRI trials to test the retrieval success hypothesis. *Neuroimage* 7, 163–175 (1998).
26. Rolls, E. T. & Baylis, L. L. Gustatory, olfactory, and visual convergence within the primate orbitofrontal cortex. *J. Neurosci.* 14, 5437–5452 (1994).
27. Rolls, E. T., Yaxley, S. & Sienkiewicz, Z. J. Gustatory responses of single neurons in the caudolateral orbitofrontal cortex of the macaque monkey. *J. Neurophysiol.* 64, 1055–1066 (1990).
28. Rolls, E. T. & Tovee, M. J. Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *J. Neurophysiol.* 73, 713–726 (1995).
29. Rolls, E. T. The orbitofrontal cortex and reward. *Cereb. Cortex* 10, 284–294 (2000).
30. Rees, G., Friston, K. & Koch, C. A direct quantitative relationship between the functional properties of human and macaque V5. *Nat. Neurosci.* 3, 716–723 (2000).
31. Davidson, R. J. Anterior cerebral asymmetry and the nature of emotion. *Brain Cogn.* 6, 245–268 (1992).
32. Dias, R., Robbins, T. & Roberts, A. Dissociation in prefrontal cortex of affective and attentional shifts. *Nature* 380, 69–72 (1996).
33. Tremblay, L. & Schultz, W. Relative reward preference in primate orbitofrontal cortex. *Nature* 398, 704–708 (1999).
34. O'Doherty, J. *et al.* Sensory-specific satiety-related olfactory activation of the human orbitofrontal cortex. *Neuroreport* 11, 893–897 (2000).
35. Birn, R. M., Bandettini, P. A., Cox, R. W. & Shaker, R. Event-related fMRI of tasks involving brief motion. *Hum. Brain Mapp.* 7, 106–114 (1999).
36. Talairach, J. & Tournoux, P. *Co-Planar Stereotaxic Atlas of the Human Brain* (Thieme, Stuttgart, 1988).