

Choosing the Lesser of Two Evils, the Better of Two Goods: Specifying the Roles of Ventromedial Prefrontal Cortex and Dorsal Anterior Cingulate in Object Choice

Karina Blair,¹ Abigail A. Marsh,¹ John Morton,² Meena Vythilingam,¹ Matthew Jones,¹ Krystal Mondillo,¹ Daniel C. Pine,¹ Wayne C. Drevets,¹ and James R. Blair¹

¹Mood and Anxiety Disorders Program, National Institute of Mental Health, National Institutes of Health, Department of Health and Human Services, Bethesda, Maryland 20892-2670, and ²Institute of Cognitive Neuroscience, University College London, London WC1X 8EE, United Kingdom

The ventromedial prefrontal cortex (vmPFC) and dorsal anterior cingulate cortices (ACd) are considered important for reward-based decision making. However, work distinguishing their individual functional contributions has only begun. One aspect of decision making that has received little attention is that making the right choice often translates to making the better choice. Thus, response choice often occurs in situations where both options are desirable (e.g., choosing between mousse au chocolat or crème caramel cheesecake from a menu) or, alternatively, in situations where both options are undesirable. Moreover, response choice is easier when the reinforcements associated with the objects are far apart, rather than close together, in value. We used functional magnetic resonance imaging to delineate the functional roles of the vmPFC and ACd by investigating these two aspects of decision making: (1) decision form (i.e., choosing between two objects to gain the greater reward or the lesser punishment), and (2) between-object reinforcement distance (i.e., the difference in reinforcements associated with the two objects). Blood oxygen level-dependent (BOLD) responses within the ACd and vmPFC were both related to decision form but differentially. Whereas ACd showed greater responses when deciding between objects to gain the lesser punishment, vmPFC showed greater responses when deciding between objects to gain the greater reward. Moreover, vmPFC was sensitive to reinforcement expectations associated with both the chosen and the forgone choice. In contrast, BOLD responses within ACd, but not vmPFC, related to between-object reinforcement distance, increasing as the distance between the reinforcements of the two objects decreased. These data are interpreted with reference to models of ACd and vmPFC functioning.

Key words: decision; fMRI; reward; punishment; vmPFC; ACd

Introduction

Making choices on the basis of expected rewards and punishments is crucial for survival. Several studies involving reward-based decision making have shown activation in the ventromedial prefrontal cortex (vmPFC) and the dorsal/supracallosal parts of the anterior cingulate (ACd) (Bush et al., 2002; Williams et al., 2004; Rushworth et al., 2005). However, work distinguishing their functional roles has only begun (Walton et al., 2004).

One aspect of decision making that has received little attention is that making the right choice often translates to making the better choice. That is, response choice often occurs in situations in which both options are desirable (e.g., choosing between mousse au chocolat or crème caramel cheesecake from a menu) or both are undesirable. In addition, the impact of between-object reinforcement distance has received little attention. Yet, it is easier to choose between objects far apart in desirability (e.g.,

mousse au chocolat vs buttered toast) than objects close together (e.g., mousse au chocolat vs crème caramel cheesecake).

Previous work might suggest that both the vmPFC and ACd respond to reward-based decisions. Thus, several studies have reported vmPFC activation to reward and ventrolateral prefrontal cortex (vlPFC) activation to punishment (O'Doherty et al., 2001; Anderson et al., 2003; Kringelbach, 2005). This could predict that deciding between two options associated with reward will recruit vmPFC whereas deciding between two options associated with punishment will recruit vlPFC. Recent work has also stressed the role of the ACd in reward-based decision making (Bush et al., 2002; Richmond et al., 2003; Williams et al., 2004). This could predict that ACd too will show greater activity when deciding between two options associated with reward rather than two options associated with punishment.

Previous work on decision making might also suggest that the vmPFC and ACd will be impacted by between-object reinforcement distance. Thus, on the basis of significant correlation between vmPFC activity and subjective reports of choice difficulty, it has been suggested that the vmPFC is involved in the comparison of reward values (Arana et al., 2003), predicting increased vmPFC activity as reinforcement differences decrease (cf. Blair, 2004). Alternatively, if outcome expectancies are translated as

Received April 17, 2006; revised Aug. 28, 2006; accepted Sept. 19, 2006.

This work was supported by the Intramural Research Program of the National Institutes of Health, National Institute of Mental Health. We thank Madeline Jacobs for help with preparation of this paper.

Correspondence should be addressed to Karina Blair, Mood and Anxiety Disorders Program, National Institute of Mental Health, 15K North Drive, MSC 2670, Bethesda, MD 20892. E-mail: blairka@mail.nih.gov.

DOI:10.1523/JNEUROSCI.1640-06.2006

Copyright © 2006 Society for Neuroscience 0270-6474/06/2611379-08\$15.00/0

approach/avoidance tendencies, then as distance between the reinforcements associated with the two objects decreases, there should be greater similarity in approach/avoidance strengths, greater response competition, and greater ACd activation, given its suggested role in response conflict resolution (Carter et al., 2000; Botvinick et al., 2004).

In short, the above literature suggests a similar responsiveness of the vmPFC and ACd to decision form and between-object reinforcement level distance. However, initial evidence of functional specificity exists (Walton et al., 2004). Moreover, whereas the data suggesting that vmPFC signals reinforcement expectancies has been relatively clear (Elliott et al., 2003; Cox et al., 2005; Schoenbaum and Roesch, 2005), the role of ACd with respect to reward is less certain. Similarly, data on the role of the vmPFC in object selection is tentative. Thus, an alternative hypothesis is that the vmPFC will show greater responsiveness related to decision form whereas ACd will show greater responsiveness to between-object reinforcement distance. The current study tests these hypotheses.

Materials and Methods

Subjects. Twenty-one right-handed subjects (11 males, 10 females; aged 21–42; mean age, 27.85) volunteered for the study and were paid for their participation. All subjects gave written informed assent/consent to participate in the study, which was approved by the National Institute of Mental Health Institutional Review Board. Subjects were in good health, with no history of psychiatric or neurological disease. The data from one subject whose behavioral responses did not meet criterion levels for successful performance on the task (criterion was set at ≥ 20 correct responses in every condition) were collected but excluded from the analyses.

MRI data acquisition. Whole-brain blood oxygen level-dependent (BOLD) functional magnetic resonance imaging (fMRI) data were acquired using a 1.5 tesla Siemens (Erlangen, Germany) MRI scanner. After sagittal localization, functional T2* weighted images were acquired using an echo-planar single-shot gradient echo pulse sequence with a matrix of 64×64 mm, repetition time (TR) of 3000 ms, echo time (TE) of 30 ms, field of view (FOV) of 240 mm, and voxels of $3.75 \times 3.75 \times 4$ mm. Images were acquired in 31 contiguous 4 mm axial slices per brain volume. The functional data were acquired over four runs, each lasting 8 min 45 s. In the same session, a high-resolution T1-weighted anatomical image was acquired to aid with spatial normalization (three-dimensional spoiled grass; TR, 8.1 ms; TE, 3.2 ms, flip angle, 20°; FOV, 240 mm; 124 axial slices; thickness, 1.0 mm; 256×256 acquisition matrix).

Differential reward/punishment task and experimental procedure. The stimuli were a set of 10 line drawings from the Snodgrass and Vanderwart (1980) picture set. Each stimulus depicted a common object: a house, cup, fork, duck, pineapple, necklace, raccoon, door, flashlight, or shoe. Before the study, each stimulus was randomly assigned a differential associated reinforcement value (−900, −700, −500, −300, −100, 100, 300, 500, 700, or 900 points) that would be uniquely associated with that particular stimulus throughout the experiment.

On each trial, objects were presented together in pairs for 1.5 s, appearing in two of four (left-hand top, left-hand bottom, right-hand top, right-hand bottom) screen locations. Feedback would then appear for 1 s (e.g., “You have won 900 points”). If no response was recorded during the presentation of the two objects (<2% of the trials), subjects were displayed the feedback: “Respond faster.” After feedback, a fixation point would appear on the screen for 500 ms. Subjects were told that on each trial one of the two objects must be chosen, and that choosing some objects would mean losing points and that choosing some objects would mean winning points. They were told that their goal was to win as many points as possible.

Subjects received pretraining to acquaint them with the task procedure. The training stimuli were not included in the main paradigm. Thus, subjects did not know the contingencies associated with particular stimuli at the beginning of the imaging study.

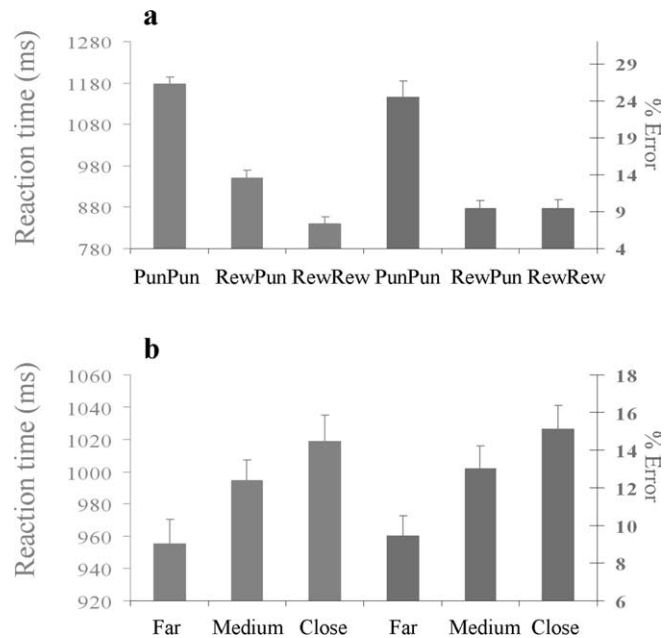


Figure 1. RTs and error rates to decision form and between-object reinforcement distance. **a**, RTs and error rates were significantly affected by decision form. Subjects were significantly slower and made significantly more errors when they responded to PunPun relative to RewPun and RewRew trials. **b**, In addition, RTs and error rates were significantly affected by between-object reinforcement distance. Subjects were significantly slower and made significantly more errors when they responded to close relative to medium trials, and to medium relative to far trials. Error bars indicate SE.

The study involved a three (decision form: PunPun, RewRew, RewPun) by three [between-object reinforcement level distance (distance): close, medium, far] design. RewRew trials involved two objects both associated with a reward (e.g., 100 vs 300; 100 vs 500; 300 vs 700). On these trials, response choice of either object would result in a point gain; however, one of the objects would result in the greater point gain (see Fig. 2). PunPun trials involved two objects both associated with a punishment (e.g., −100 vs −300; −100 vs −500; −300 vs −700). On these trials, response choice of either object would result in a point loss; however, one of the objects would result in the greater point loss. Thus, on these trials the subjects’ strategy was to select the object associated with the smaller point loss. RewPun trials involved one object associated with reward and one object associated with punishment (e.g., 100 vs −100; 100 vs −300; 100 vs −500). On these trials, response choice of one of the objects would result in a point gain whereas the response choice of the competing object would result in a punishment (for example trials, see Fig. 2). RewRew and RewPun trials both involved the point rewards 300, 500, 700, and 900 and the study was designed so that subjects should win a comparable number of points on these two different type trials. However, subjects won significantly more on RewRew relative to RewPun trials (858 and 560 points per trial type respectively). Unsurprisingly, they also won very considerably more on both RewRew and RewPun trials than on PunPun trials (−391 points).

The “close” between-object reinforcement distance trials involved two objects associated with values that were close together in value (e.g., −900 vs −700; 900 vs 700; 300 vs −100). The “far” between-object reinforcement distance trials involved two objects associated with values that were far apart in magnitude (e.g., −900 vs −100; 900 vs 100; 300 vs −900) (for example trials, see Fig. 3). Distance conditions were matched for total points won.

On any trial, selecting the superior choice over the inferior choice was scored as “correct”. Thus, on PunPun trials, where both objects represented a point loss (e.g., −100 vs −300), selecting the object representing the smaller loss of −100 was scored as correct, and on RewRew trials, where both objects represented a gain (e.g., 100 vs 300), selecting the object representing the greater gain of 300 was scored as correct. On

Table 1. Significant main effect areas of activation for decision form and between-object reinforcement level difference*

Region	BA	mm ³	X	Y	Z
Moderated by decision form					
R vmPFC ^{a,b,c}	10	9404	0	51	0
R amygdala ^{a,b,c,*}		98	26	−6	−14
L superior temporal gyrus ^{a,b,c}	42	3605	−60	−29	17
L middle temporal gyrus ^{a,b,c}	19	2941	−54	−74	15
R middle temporal gyrus ^{a,b}	37	2474	62	−62	4
L precuneus ^{a,b,c}	19	7839	−29	−63	40
R ACd ^{d,e}	32	9454	3	16	46
R middle frontal gyrus ^{d,e}	8	3605	49	22	44
L insula ^{d,e}	13	5395	−39	15	5
R insula ^{d,e}	13	2125	36	19	4
R inferior parietal lobule ^{d,e}	40	9812	43	−56	50
L middle occipital gyrus ^{d,e}	18	6943	−47	−78	−11
Moderated by between-object reinforcement level difference					
R ACd	32	4753	1	12	24
L middle frontal gyrus	6	1443	−35	9	32

*All activations are effects observed in whole-brain analyses corrected for multiple comparisons (significant at $p < 0.05$) except ** $p < 0.001$, uncorrected for multiple comparisons. BA, Brodmann's area; L, left; R, right.

^{a,b,c,d,e}The BOLD response was significantly greater to ^aRewRew than PunPun trials, ^bRewRew than RewPun trials, ^cRewPun than PunPun trials, ^dPunPun than RewRew trials, and ^ePunPun than RewPun trials.

RewPun trials, where one object represented a gain and one object represented a loss (e.g., −100 vs 100), selecting the rewarding object was scored as correct.

The paradigm was programmed in E-Studio. Stimuli were presented on a computer display that was projected onto a mirror in the MRI scanner. Subjects were placed in a light head restraint within the scanner to limit movement during acquisition.

The fMRI scan acquisition followed an event-related design, and consisted of four runs, each containing 150 experimental trials and 25 fixation point trials. Only correct responses were analyzed (incorrect responses were modeled within one separate error regressor).

fMRI analysis. Data were analyzed within the framework of a random effects general linear model using Analysis of Functional Neuroimages (Cox, 1996). Both individual and group-level analyses were conducted. The first four volumes in each scan series, collected before equilibrium magnetization was reached, were discarded. Motion correction was performed by registering all volumes in the echo-planar imaging (EPI) dataset to a volume that was collected shortly before acquisition of the high-resolution anatomical dataset.

The EPI datasets for each subject were spatially smoothed (using an isotropic 6 mm Gaussian kernel) to reduce the influence of anatomical variability among the individual maps in generating group maps. Next, the time series data were normalized by dividing the signal intensity of a voxel at each time point by the mean signal intensity of that voxel for each run and multiplying the result by 100. Resultant regression coefficients represented a percent signal change from the mean. After this, regressors depicting each of the response types were created by convolving the train of stimulus events with a gamma-variate hemodynamic response function to account for the slow hemodynamic response. This involved 10 regressors (RewRew close, RewRew medium, RewRew far, PunPun close, PunPun medium, PunPun far, RewPun close, RewPun medium, RewPun far, error/missed responses) with fixation point baseline trials. The regressors were modeled at time of trial onset. Linear regression modeling was then performed using the regressors described above plus regressors to model a first order baseline drift function. This produced for each voxel and each regressor, a β coefficient and its associated t statistic.

Voxel-wise group analyses involved transforming single subject β coefficients into the standard coordinate space of Talairach and Tournoux (1988). Subsequently, a three-way ANOVA involving a 3 (decision form: RewRew, PunPun, RewPun) by 3 (distance: close, medium, far) by 20 (subject 1–20) design was performed to produce statistical parametric maps of the main effect of decision form (stimuli associated with reward or punishment values), and between-object reinforcement distance (associated values close, medium, or far apart in value). The result was three

whole-brain group maps of areas of differential activation ($p < 0.001$). To correct for multiple comparisons, a spatial clustering operation was performed using AlphaSim with 1,000 Monte Carlo simulations taking into account the entire EPI matrix ($p < 0.001$).

Although the purpose of the study was to test our a priori hypotheses, whole-brain analyses were conducted to ensure identification of the most statistically significant regions involved in task performance.

Results

Behavioral data

Mean RTs and error rates for each trial were computed for each subject. Separate three (decision form: PunPun, RewRew, RewPun) by three (distance: close, medium, far) repeated-measures ANOVAs were conducted on the RT and error rate data.^a These revealed main effects for decision form ($F_{(2,38)} = 163.81$ and 47.30 for RT and error rate, respectively; $p < 0.001$); subjects were slower and less accurate at selecting the correct object on PunPun trials relative to RewRew trials and RewPun trials ($[M(\text{PunPun RT})] = 1176.94$, SE = 16.29; $[M(\text{RewRew RT})] = 840.66$, SE = 16.24; $[M(\text{RewPun RT})] = 951.28$, SE = 18.59; $[M(\text{PunPun error rate})] = 24.50$, SE = 2.17; $[M(\text{RewRew error rate})] = 9.40$, SE = 1.24; $[M(\text{RewPun error rate})] = 9.45$, SE = 1.03). They were also slower at selecting the correct object on RewPun trials relative to RewRew trials (Fig. 1). In addition, there were main effects for distance ($F_{(2,38)} = 38.00$ and 22.41 for RT and error rate, respectively; $p < 0.001$); subjects were slower and less accurate at selecting the correct object as between-object reinforcement distance decreased. There was a significant decision form by distance interaction for the RT data ($F_{(4,76)} = 12.97$; $p < 0.001$); the increase in RTs across between-object reinforcement distance was significantly greater for RewPun relative to both RewRew and PunPun trials ($p < 0.001$ and 0.005, respectively).

fMRI data

A three (decision form: PunPun, RewRew, RewPun) by three (distance: close, medium, far) ANOVA was performed on the data. This revealed significant main effects for both decision form and between-object reinforcement distance but no significant interaction.

Main effect of decision form

The first main effect identified regions that showed a differential BOLD response for decision form. These included the right vmPFC, right amygdala, and bilateral temporal regions, which all showed significantly greater responses for RewRew rather than RewPun trials and RewPun rather than PunPun trials. They also included the right ACd, right middle frontal gyrus, right inferior parietal lobule, and bilateral insula which, in contrast, showed significantly greater responses for PunPun rather than RewPun or RewRew trials (Table 1, Fig. 2*d,e*).

^aTo examine time course effects, we also applied a four (run: 1, 2, 3, 4) by three (decision form: PunPun, RewRew, RewPun) by three (distance: close, medium, far) ANOVA to the RT data. There was a main effect of run ($F_{(1,74,33,11)} = 5.84$; $p < 0.01$); performance did improve across runs, however, there was no significant run by decision form, run by distance, or run by decision form by distance interactions ($F = 2.51, 1.41, \text{ and } 0.84$, respectively). That is, although there was a general improvement in performance, it did not differentially impact on the three different decision forms or distances.

Main effect of between-object reinforcement distance

The main effect of reinforcement distance revealed BOLD responses in the ACd and left frontal gyrus, which increased as between-object reinforcement distance decreased (Table 1, Fig. 3*d*). In contrast, there was no main effect activation in the mvPFC (the percentage signal change within the previously identified area of mvPFC to between-object reinforcement distance is depicted in Fig. 3*e*).

Analysis excluding RewPun trials

The distances used for the distance contingencies for the RewRew and PunPun trials were identical. However, the distances for the distances for the RewPun trials differed from those used in the RewRew and PunPun trials. To be sure that data from the RewPun trials was not weighting the results, we reanalyzed the data only from the RewRew and PunPun trials. The results from that analysis were almost identical to those reported above and confirm the study's main findings (supplemental information, supplemental Figs. 1, 2, available at www.jneurosci.org as supplemental material).

Effect of total available value

Subjects won more for RewRew than RewPun trials and considerably more for both than PunPun trials (858, 560, and –391 points, respectively). Previous literature showing that the vmPFC is preferentially involved in the processing of reward should therefore predict greater vmPFC activity for RewRew relative to RewPun trials and considerably greater activity for both than for PunPun trials. However, the data indicated that vmPFC showed considerably greater activity for RewRew than either RewPun ($F_{(1,19)} = 27.82$; $p < 0.001$) or PunPun ($F_{(1,19)} = 29.02$; $p < 0.001$) trials and a relatively small difference between RewPun and PunPun trials ($F_{(1,19)} = 7.50$; $p < 0.05$). These data suggested the interesting possibility that vmPFC might not only be representing the expected reinforcement associated with the chosen option but also the forgone option. If the vmPFC represents multiple reinforcements attached to multiple stimulus options, then activity within this region should be influenced by total available reinforcement (i.e., the combined reward/punishment associated with both options).

We examined this possibility in a secondary analysis. Initially, the total available reinforcement for each trial was calculated. So, for example, the total available reinforcement for a trial where one object was associated with a 100 point gain and the other object was associated with a 500 point gain would be 600, as would the total available reinforcement for a trial involving one object associated with respectively a 100 point loss and another a 700 point gain. This resulted in 17 different total available reinforcement conditions: –1600, –1400, –1200, –1000, –800, –600, –400, –200, 0, 200, 400, 600, 800, 100, 1200, 1400, and 1600. Regressors for these 17 different total available reinforcement conditions were created using the method described above.

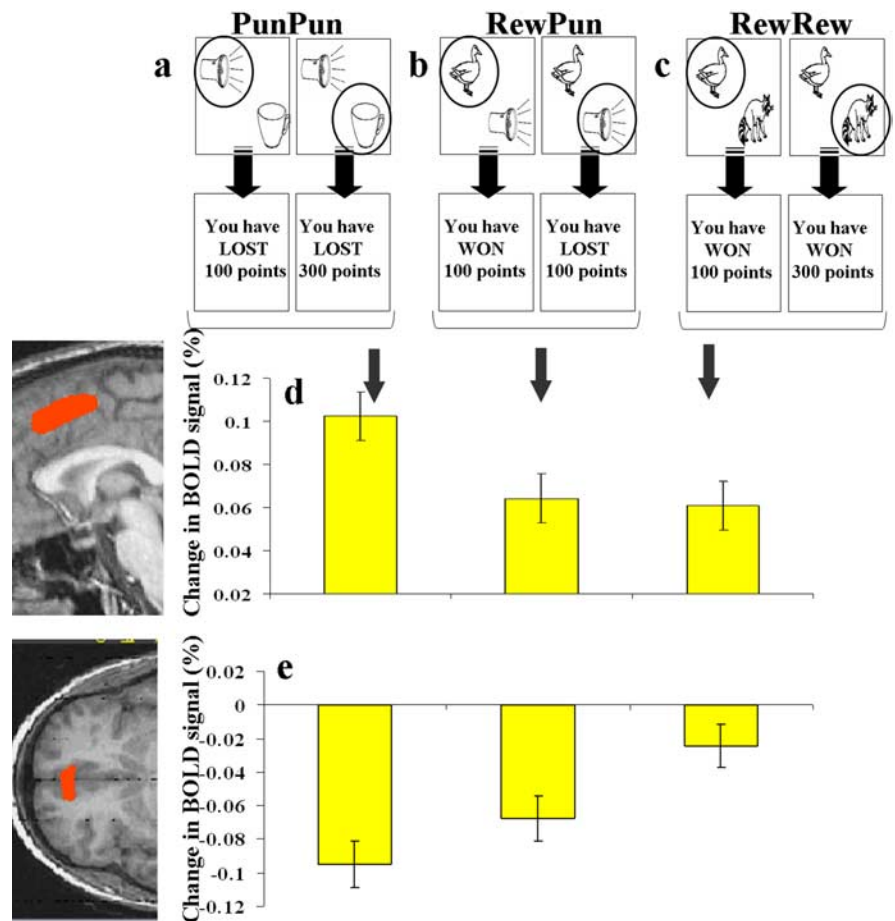


Figure 2. Decision form and Acd/vmPFC responses. *a–c*, Sample stimuli and results of (ringed) responses for PunPun (*a*), RewPun (*b*), and RewRew (*c*) trials. *d*, The BOLD response within the ACd (3, 16, 46) was significantly greater to PunPun relative to RewPun and RewRew trials. *e*, The BOLD response within the vmPFC (0, 51, 0) was significantly greater to RewRew relative to RewPun and PunPun trials. Error bars indicate SE.

Given the origins of this second analysis, we decided to examine activity associated with these regressors within the vmPFC region identified by the decision form main effect (i.e., the region which showed significantly greater BOLD response to RewRew relative to RewPun and PunPun trials). This analysis showed a highly significant linear distance effect for total available reinforcement; activity within this region increased as the total available reinforcement value increased ($F_{(1,19)} = 30.30$; $p < 0.001$) (Fig. 4*a*).

We also examined the relationship between activity within the vmPFC and the value of the chosen option (900, 700, 500, 300, 100, –100, –300, –500, –700). We found a significant relationship between vmPFC activity and the value of the chosen option; activity within vmPFC increased as the value of the chosen option increased ($F_{(1,19)} = 8.17$; $p < 0.05$) (Fig. 4*b*).

To examine whether vmPFC activity was best predicted by total available reinforcement or value of the chosen option, we examined the correlations of vmPFC activity with these two variables for each subject. This revealed that subjects showed significantly stronger correlations between vmPFC activity and total available reinforcement than vmPFC and the value of the chosen option ($F_{(1,19)} = 8.05$; $p < 0.01$); [*M* *r*Pearson's (total available reinforcement)] = 0.470; [*M* *r*Pearson's (value of option chosen)] = 0.271.

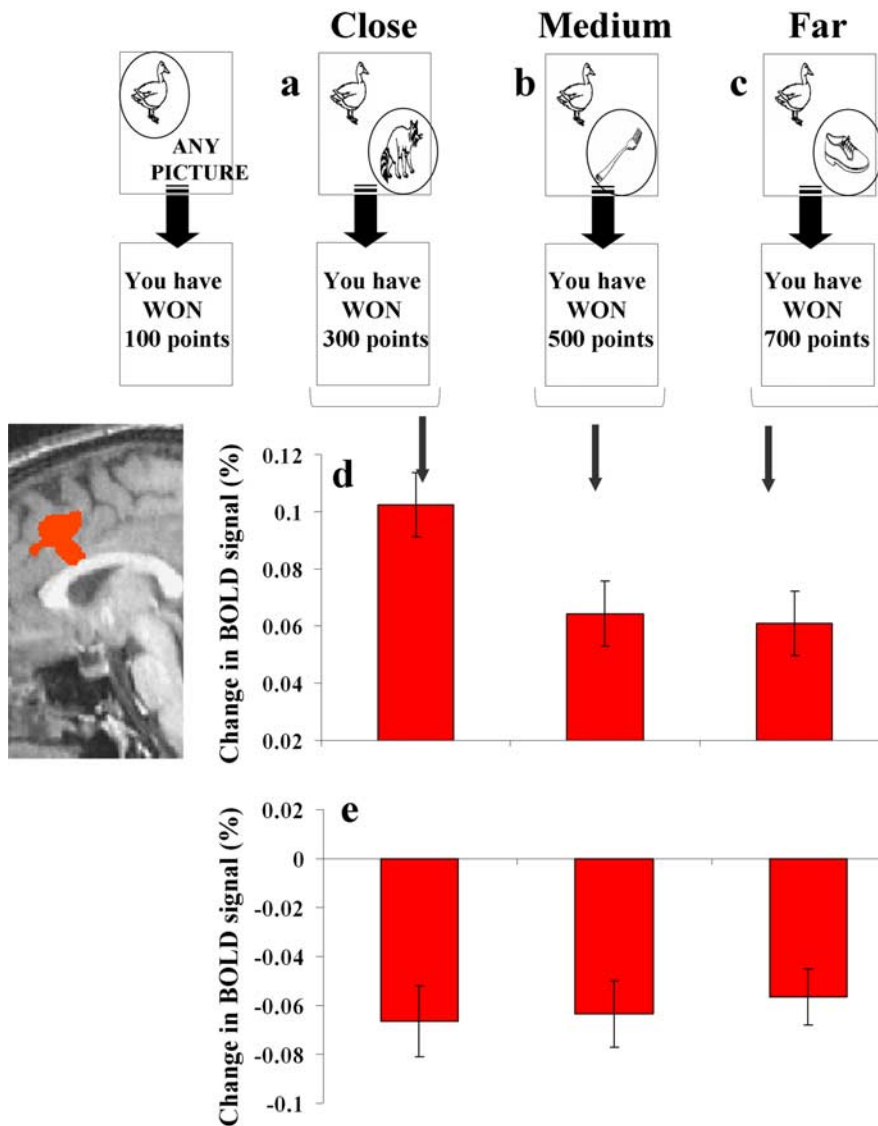


Figure 3. Between-object reinforcement distance and ACd/vmPFC responses. *a–c*, Sample stimuli and results of (ringed) responses for close (*a*), medium (*b*), or far (*c*) between-object reinforcement distance. *d*, The BOLD responses within the ACd (1, 12, 24) increased as between-object reinforcement distance decreased. *e*, In contrast, the vmPFC (0, 51, 0) did not show any differential activation for close, medium, or far trials. Error bars indicate SE.




Figure 4. vmPFC (0, 51, 0) response to total available reinforcement. *a, b*, The BOLD response within vmPFC increased as a function of increased total available reinforcement value (*a*) and value of option chosen (*b*). Error bars indicate SE.

ACd: overlay analysis

Our analysis showed that there was a differential BOLD response in ACd for both decision form and distance. To examine whether these two ACd activations represented distinct areas with functional specificity, we conducted an overlay analysis and obtained statistical maps for the area of ACd (1) exclusively activated by

decision form (ACd decision form), (2) exclusively activated by distance (ACd distance), and (3) activated by both decision form and distance (ACd decision form and distance). There was a significant area activated by both decision form and distance (2200 mm³); however, there were also two areas of ACd associated with decision form and distance (7247 and 2709 mm³, respectively) (Fig. 5). Thus, we applied a three (decision form: PunPun, RewRew, RewPun) by three (distance: close, medium, far) ANOVA to the percentage signal change within these three different functional regions of interest. There was a significant main effect for decision form in both ACd (decision form) and ACd (both decision form and distance) ($F = 22.29$ and 18.19 , respectively; $p < 0.001$); i.e., the two areas that the overlay analysis demonstrated to be associated with decision form. There was a trend toward a significant effect for decision form in ACd (distance) ($F = 2.30$; $p = 0.08$). There was a significant main effect of distance for all three ACd areas [$F = 7.68, 21.47, \text{ and } 16.17$, respectively; $p < 0.005$ for ACd (decision form), ACd (both decision form and distance), and ACd (distance)]. In short, the ACd activations appear to show significant functional overlap.

Discussion

The main goal of the current study was to determine whether manipulation of two parameters, decision form and between-object reinforcement distance, would help to distinguish the roles of the ACd and vmPFC in decision making. With respect to decision form, we found that ACd showed a greater signal to choices between “bad” options whereas the vmPFC showed the greatest signal to choices between “good” options. With respect to reinforcement difference, ACd showed the greatest signal as the difference between the reinforcement levels associated with the two options decreased. In contrast, the vmPFC showed no significant modulation by this parameter.

Recent work has suggested that the anterior cingulate cortex/ACd plays a role in reward-based decision making (Bush et al., 2002; Richmond et al., 2003; Rogers et al., 2004). Other studies have shown ACd activity during task conditions where an alteration in reward level suggests a behavioral change (Bush et al., 2002; Williams et al., 2004) or when freely choosing a new rule (Walton et al., 2004). On the basis of such data, it has been suggested that the ACd uses reward and error outcome information to guide voluntary response selection (Walton et al., 2004; Rushworth et al., 2005). The current data can be considered compatible with this

view. Certainly, the ACD activation identified by the decision form main effect showed significant activity, relative to baseline, when choosing between two options associated with different levels of reward. However, this region showed significantly greater BOLD response when choosing between two options associated with different levels of punishment. Thus, at the very least, ACD must be using reward and punishment information to guide decision making. However, it remains unclear on the basis of this position why the BOLD response within the ACD should be greater when choosing between bad options than when choosing between good options.

An alternative conceptualization of ACD functioning is that it is involved in the monitoring of response conflict (Carter et al., 2000; Cohen et al., 2000; Botvinick et al., 2004; Kerns et al., 2004). On the basis of this position, we argued that if outcome expectancies are translated as approach/avoidance tendencies, then response options that are close in reinforcement value should be associated with similar strength approach/avoidance tendencies. Increased similarity in approach/avoidance strength should mean increased response conflict and, on the basis of this model, increased ACD activation. In contrast, response options distant in reinforcement value from one another should be associated with approach/avoidance tendencies of notably different strengths. As one will more easily “win” over the other, there will be less response conflict and less recruitment of ACD. This prediction was confirmed. A significant increase in ACD signal [centered in the paracingulate cortex signal and extending into presupplementary motor area, with focal coordinates within the region identified in the review by Botvinick et al. (2004) as most involved in the response to conflict] was seen as reinforcement level differences between the objects to be chosen decreased.

Increased response competition may also be the explanation for the increased signal within the ACD when choosing between bad rather than good options. As can be seen in Figure 1, deciding between bad options was more difficult than deciding between good options (indexed both in terms of RT and error rate). Although the reason for this increased difficulty is uncertain, it is likely to result in increased competition between the two choice options which the models (Carter et al., 2000; Cohen et al., 2000; Botvinick et al., 2004; Kerns et al., 2004) suggest would lead to increased ACD activity. In line with the suggestion, the overlay analysis revealed considerable functional overlap between the ACD activations associated with decision form and between-object reinforcement distance.

With respect to the vmPFC, some, although not all (Elliott et al., 2003), studies have reported greater vmPFC/medial orbitofrontal cortex (mOFC) responses to reward whereas punishment activates more lateral orbitofrontal cortex (lOFC)/vlPFC (O’Doherty et al., 2001; Anderson et al., 2003; Kringelbach, 2005). This position would suggest that deciding between two

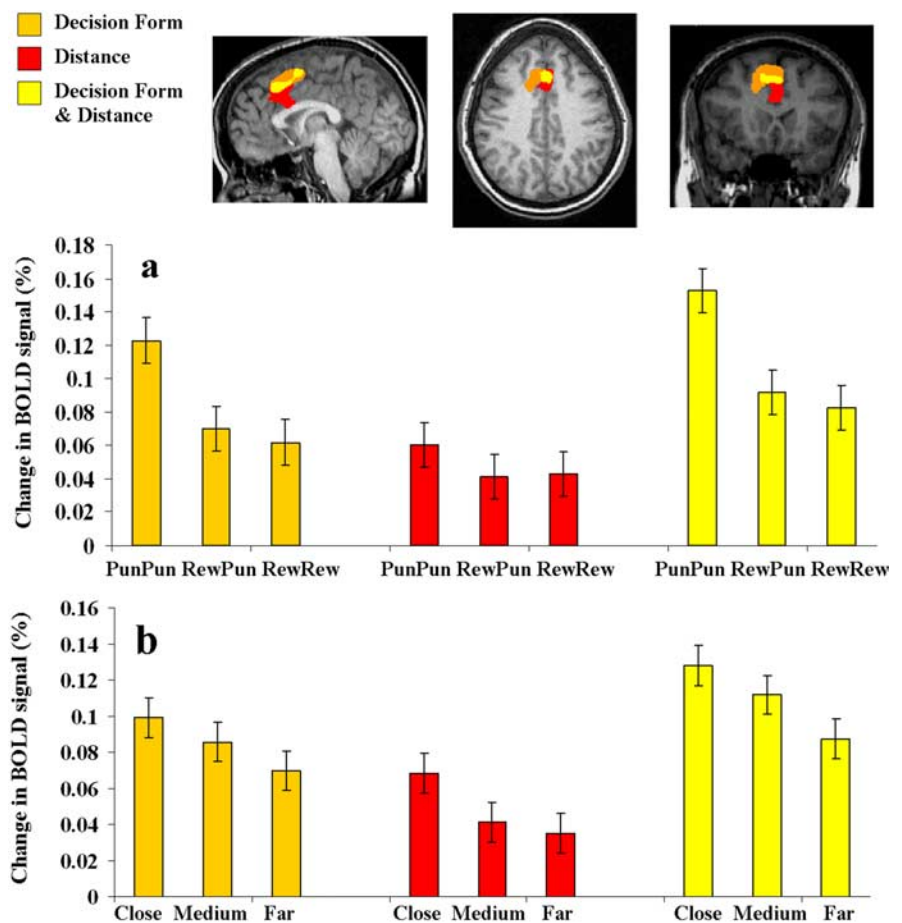


Figure 5. Overlay analysis for ACD. *a, b*, The BOLD response within the three identified ACD areas differentiated for type of decision form (*a*) and between-object reinforcement distance (*b*). Error bars indicate SE.

options associated with reward will involve the vmPFC whereas deciding between two options associated with punishment will involve the lOFC/vlPFC. The current data are in line with the suggestion that the vmPFC responds to reward expectations (Schoenbaum et al., 1998; Montague and Berns, 2002; Schoenbaum and Roesch, 2005). However, we saw no indication of increased BOLD responses within the lOFC/vlPFC either when choosing between response options that were both associated with punishment or as a function of expected punishment.

Of course, previous findings of responses within the lOFC/vlPFC to aversive stimulation may partly reflect the response demands of this stimulation. In response reversal studies, where lOFC/vlPFC responses to punishment information have been seen, this information was a cue for response control/change (O’Doherty et al., 2001, 2003). Findings of vlPFC activation to unpleasant olfactory cues may also reflect the modulation of behavior, control over the withdrawal response prompted by the cue (Anderson et al., 2003). In contrast, in the current task, on PunPun trials the subject had to decide between two objects which were both associated with punishment. Such trials required response selection but not the overruling/modulation of a pre-existing motor response. We cannot rule out the possibility that dropout susceptibility within the lOFC/vlPFC could have decreased the possibility of finding punishment-related activity within this area. However, it should be noted that the scanning parameters used in this study have successfully demonstrated

IOFC/vIPFC activity in other work (Budhani et al. 2006; Finger et al., 2006; Luo et al. 2006; Mitchell et al., 2006). This suggests that activity within these regions was not significantly related to the task parameters.

In contrast to BOLD response within the ACD, activity in the vmPFC was significantly greater when choosing between two good options rather than one good and one bad option or two bad options. This is in line with previous suggestions that the OFC has a role in signaling outcome expectancies (Montague and Berns, 2002; Schoenbaum and Roesch, 2005) and previous results indicating its activity reflects the values of anticipated rewards (Schoenbaum et al., 1998; Tremblay and Schultz, 1999; Elliott et al., 2000, 2003; Cox et al., 2005; Kosson et al., 2006). However, as noted above, average reward received by subjects when choosing between two good options rather than one good and one bad option did not significantly differ. This suggested that the vmPFC activity was not solely determined by the reinforcement expectancy associated with the selected object; if this were the case, there would have been no difference in BOLD response for RewRew and RewPun trials. We therefore conducted a follow-up analysis to determine whether BOLD response in the vmPFC reflected the reinforcement value attached to the foregone option. We found that BOLD response within the vmPFC region identified by the decision form main effect, was a function of total available reinforcement (i.e., the combined reward/punishment of both options) on any given trial.

This result is of interest with respect to recent work on “regret” (Coricelli et al., 2005). Coricelli et al. (2005) found that vmPFC activity was modulated by information on reinforcement associated with both the selected and the foregone choice (activity increased if the foregone choice involved reward and decreased if it involved punishment). In the study by Coricelli et al. (2005), BOLD response within the vmPFC changed as a result of the revelation of the reinforcement associated with the foregone option; i.e., not as a result of an expectancy. Our data extend this result by indicating that the vmPFC response does not require the reinforcement associated with the foregone choice to be revealed. If previous learning is sufficient to generate a reinforcement expectancy associated with the foregone conclusion, this will be reflected in vmPFC activation.

Previous work has suggested that the vmPFC is involved in the comparison of goal values (Arana et al., 2003; Blair, 2004). Data in support of this suggestion include the finding by Arana et al. (2003) that vmPFC activity correlated with subjective reports of choice difficulty. We predicted that if the vmPFC is involved in the comparison of goal values, then signal within this region should be sensitive to between-object reinforcement distance. Choice difficulty, as indexed by both RT and error rate, did increase with decreasing distance between the reinforcement values associated with the response options. However, in contrast to the BOLD response within ACD, there was no significant change in BOLD response within vmPFC as a function of between-object reinforcement distance. We suggest therefore, following Schoenbaum and Roesch (2005), that the vmFC/mOFC codes reinforcement expectancies, perhaps normalizing the value of competing outcomes (cf. Montague and Berns, 2002). However, this region itself does not directly select between responses but rather allows the representation of information crucial for selection, information that ACD appears to be operating on.

In short, in this study we demonstrated the complementary yet dissociable roles that the ACD and vmPFC play with respect to decision making. Both were related to decision form but differentially. The ACD showed greater responses when both choices

were undesirable, whereas the vmPFC showed greater responses when both choices were desirable. In addition, BOLD responses within the ACD, but not the vmPFC, related to the distance between the desirability of the choices on display. Finally, the vmPFC was sensitive to the reinforcement expectations associated with both the chosen and the foregone option. These results may help resolve how the ACD and vmPFC select objects in situations where making the right choice is a matter of making the better choice.

References

- Anderson AK, Christoff K, Stappen I, Panitz D, Ghahremani DG, Glover G, Gabrieli JD, Sobel N (2003) Dissociated neural representations of intensity and valence in human olfaction. *Nat Neurosci* 6:196–202.
- Arana FS, JA Parkinson, et al (2003) Dissociable contributions of the human amygdala and orbitofrontal cortex to incentive motivation and goal selection. *J Neurosci* 23:9632–9638.
- Blair RJ (2004) The roles of orbital frontal cortex in the modulation of antisocial behavior. *Brain Cogn* 55:198–208.
- Botvinick MM, Cohen JD, Carter CS (2004) Conflict monitoring and anterior cingulate cortex: an update. *Trends Cogn Sci* 8:539–546.
- Budhani S, Pine DS, Blair RJR (2006) Ventrolateral prefrontal cortex and response reversal: inhibiting prepotent responses, representing punishment or activating alternative responses? *NeuroImage*, in press.
- Bush G, Vogt BA, Holmes J, Dale AM, Greve D, Jenike MA, Rosen BR (2002) Dorsal anterior cingulate cortex: a role in reward-based decision making. *Proc Natl Acad Sci USA* 99:523–528.
- Carter CS, Macdonald AM, Botvinick M, Ross LL, Stenger VA, Noll D, Cohen JD (2000) Parsing executive processes: strategic versus evaluative functions of the anterior cingulate cortex. *Proc Natl Acad Sci USA* 97:1944–1948.
- Cohen JD, Botvinick M, Carter CS (2000) Anterior cingulate and prefrontal cortex: who's in control? *Nat Neurosci* 3:421–423.
- Coricelli GG, Critchley HD, Joffly MM, O'Doherty JP, Sirigu A, Dolan RJ (2005) Regret and its avoidance: a neuroimaging study of choice behavior. *Nat Neurosci* 8:1255.
- Cox RW (1996) AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput Biomed Res* 29:162–173.
- Cox SM, Andrade A, Johnsrude IS (2005) Learning to like: a role for human orbital frontal cortex in conditioned reward. *J Neurosci* 25:2733–2740.
- Elliott R, Friston KJ, Dolan RJ (2000) Dissociable neural responses in human reward systems. *J Neurosci* 20:6159–6165.
- Elliott R, Newman JL, Longe OA, Deakin JF (2003) Differential response patterns in the striatum and orbitofrontal cortex to financial reward in humans: a parametric functional magnetic resonance imaging study. *J Neurosci* 23:303–307.
- Finger EC, Marsh AA, Kamel N, Mitchell DG, Blair JR (2006) Caught in the act: the impact of audience on the neural response to morally and socially inappropriate behavior. *NeuroImage* 33:414–421.
- Kerns JG, Cohen JD, MacDonald AW, Cho RY, Stenger VA, Carter CS (2004) Anterior cingulate conflict monitoring and adjustments in control. *Science* 303:1023–1026.
- Kosson DS, Budhani S, Nakic M, Chen G, Saad ZS, Vythilingam M, Pine DS, Blair RJ (2006) The role of the amygdala and rostral anterior cingulate in encoding expected outcomes during learning. *NeuroImage* 29:1161–1172.
- Kringelbach ML (2005) The human orbitofrontal cortex: linking reward to hedonic experience. *Nat Rev Neurosci* 6:691–702.
- Luo Q, Nakic M, Wheatley T, Richell R, Martin A, Blair RJ (2006) The neural basis of implicit moral attitude: an IAT study using event related fMRI. *NeuroImage* 30:1449–1457.
- Mitchell DGV, Nakic M, Fridberg D, Pine DS, Blair RJR (2006) The impact of processing load on emotion. *NeuroImage*, in press.
- Montague PR, Berns GS (2002) Neural economics and the biological substrates of valuation. *Neuron* 36:265–284.
- O'Doherty J, Kringelbach ML, Rolls ET, Hornak J, Andrews C (2001) Abstract reward and punishment representations in the human orbitofrontal cortex. *Nat Neurosci* 4:95–102.
- O'Doherty J, Critchley H, Deichmann R, Dolan RJ (2003) Dissociating valence of outcome from behavioral control in human orbital and ventral prefrontal cortices. *J Neurosci* 23:7931–7939.

- Richmond BJ, Liu Z, Shidara M (2003) Neuroscience. Predicting future rewards. *Science* 301:179–180.
- Rogers RD, Ramnani N, Mackay C, Wilson JL, Jezzard P, Carter CS, Smith SM (2004) Distinct portions of anterior cingulate cortex and medial prefrontal cortex are activated by reward processing in separable phases of decision-making cognition. *Biol Psychiatry* 55:594–602.
- Rushworth MF, Kennerley SW, Walton ME (2005) Cognitive neuroscience: resolving conflict in and over the medial frontal cortex. *Curr Opin Biol* 15:R54–R56.
- Schoenbaum G, Roesch M (2005) Orbitofrontal cortex, associative learning, and expectancies. *Neuron* 47:633–636.
- Schoenbaum G, Chiba AA, Gallagher M (1998) Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nat Neurosci* 1:155–159.
- Snodgrass JG, Vanderwart M (1980) A standardized set of 260 pictures: norms for name agreement, familiarity, and visual complexity. *J Exp Psychol [Hum Learn]* 6:174–215.
- Talairach J, Tournoux P (1988) Co-planar stereotaxic atlas of the human brain. New York: Thieme.
- Tremblay L, Schultz W (1999) Relative reward preference in primate orbitofrontal cortex. *Nature* 398:704–708.
- Walton ME, Devlin JT, Rushworth MF (2004) Interactions between decision making and performance monitoring within prefrontal cortex. *Nat Neurosci* 7:1259.
- Williams ZM, Bush G, Rauch SL, Cosgrove GR, Eskandar EN (2004) Human anterior cingulate neurons and the integration of monetary reward with motor responses. *Nat Neurosci* 7:1370.