

Single-Neuron Mechanisms Underlying Cost-Benefit Analysis in Frontal Cortex

Takayuki Hosokawa,^{1*} Steven W. Kennerley,^{2*} Jennifer Sloan,¹ and Jonathan D. Wallis¹

¹Helen Wills Neuroscience Institute and Department of Psychology, University of California at Berkeley, Berkeley, California 94720-3190, and ²Institute of Neurology, University College London, London, WC1N 3BG, England, United Kingdom

Effective decision-making requires consideration of costs and benefits. Previous studies have implicated orbitofrontal cortex (OFC), dorsolateral prefrontal cortex (DLPFC), and anterior cingulate cortex (ACC) in cost-benefit decision-making. Yet controversy remains about whether different decision costs are encoded by different brain areas, and whether single neurons integrate costs and benefits to derive a subjective value estimate for each choice alternative. To address these issues, we trained four subjects to perform delay- and effort-based cost-benefit decisions and recorded neuronal activity in OFC, ACC, DLPFC, and the cingulate motor area (CMA). Although some neurons, mainly in ACC, did exhibit integrated value signals as if performing cost-benefit computations, they were relatively few in number. Instead, the majority of neurons in all areas encoded the decision type; that is whether the subject was required to perform a delay- or effort-based decision. OFC and DLPFC neurons tended to show the largest changes in firing rate for delay- but not effort-based decisions; whereas, the reverse was true for CMA neurons. Only ACC contained neurons modulated by both effort- and delay-based decisions. These findings challenge the idea that OFC calculates an abstract value signal to guide decision-making. Instead, our results suggest that an important function of single PFC neurons is to categorize sensory stimuli based on the consequences predicted by those stimuli.

Introduction

Damage to frontal cortex can produce dramatic deficits in value-based decision-making (Kennerley et al., 2006; Rudebeck et al., 2008; Noonan et al., 2010; Walton et al., 2010; Camille et al., 2011) and neurons in anterior cingulate cortex (ACC), orbitofrontal cortex (OFC) and dorsolateral prefrontal cortex (DLPFC) encode many decision-related attributes, such as the potential costs or benefits of a decision (Amiez et al., 2006; Kim et al., 2008; Hayden et al., 2009, 2011; Kennerley et al., 2009; Hillman and Bilkey, 2010; Amemori and Graybiel, 2012; Kennerley, 2012). Theoretical models of decision-making emphasize the importance of integrating these attributes to form a single “abstract” value estimate for each decision alternative to simplify comparison of disparate outcomes (Glimcher et al., 2005; Rangel and Hare, 2010; Padoa-Schioppa, 2011; Wallis and Rich, 2011) and abstract value signals are evident in neuroimaging studies, commonly in the vicinity of OFC including adjacent ventromedial PFC (Levy and Glimcher, 2012; Clithero and Rangel, 2013).

At the single neuron level, neurons in both OFC and ACC integrate different attributes of reward, such as its size and taste preference (Padoa-Schioppa and Assad, 2006; Cai and Padoa-Schioppa, 2012), consistent with encoding subjective reward value. However, it remains less clear whether single neurons integrate decision costs and benefits as if encoding the discounted value of a reward. Some neurons in ACC and LPFC exhibit signals that reflect cost-benefit integration during effort- and delay-based decision-making, respectively (Kim et al., 2008; Hillman and Bilkey, 2010). Yet, most OFC neurons do not exhibit cost-benefit integration for delays to reward (Roesch et al., 2006) or risk of obtaining reward (O’Neill and Schultz, 2010). Cost-benefit integration may not be necessary at the single neuron level; optimal choice might arise from populations of neurons that evaluate costs and benefits independently, with each conveying this information to the motor system to bias which action is selected.

Clouding this issue is that decision costs can be associated with the action required to obtain the outcome (e.g., physical effort) or associated directly with the outcome (e.g., delay to outcome; Rangel and Hare, 2010). These costs are typically confounded. For example, manipulating effort by varying the number of actions required to obtain reward (Crosson et al., 2009; Kennerley et al., 2009; Gan et al., 2010; Toda et al., 2012) inherently introduces a delay to obtain reward. When effort and delay are independent, both rodent lesion (Rudebeck et al., 2006) and human neuroimaging studies (Prévost et al., 2010) suggest effort- and delay-based decisions may be supported by ACC and OFC, respectively. Yet effort-based MRI activations are often near the cingulate motor area (CMA; Crosson et al., 2009; Prévost et al., 2010; Kurniawan et al., 2013). CMA neurons are implicated in action valuation and execution (Shima and Tanji, 1998; Akkal et al., 2002), and are modulated as animals work to

Received May 24, 2013; revised Sept. 19, 2013; accepted Sept. 26, 2013.

Author contributions: S.W.K. and J.D.W. designed research; T.H., S.W.K., and J.S. performed research; T.H., S.W.K., and J.D.W. analyzed data; T.H., S.W.K., and J.D.W. wrote the paper.

The project was funded by NIDA Grant R01DA19028 and NINDS Grant P01NS040813 to J.D.W., NIMH training Grant F32MH081521 to S.W.K., and JSPS Research Fellowships for Young Scientists 09J07528 to T.H. We thank Drs. Erin Rich and Timothy Behrens for their comments on an earlier draft of this paper.

The authors declare no competing financial interests.

*T.H. and S.W.K. contributed equally to this work.

Correspondence should be addressed to Dr Jonathan D. Wallis, 132, Barker Hall, MC#3190, Berkeley, CA 94720-3190. E-mail: wallis@berkeley.edu.

DOI:10.1523/JNEUROSCI.2221-13.2013

Copyright © 2013 the authors 0270-6474/13/3317385-13\$15.00/0

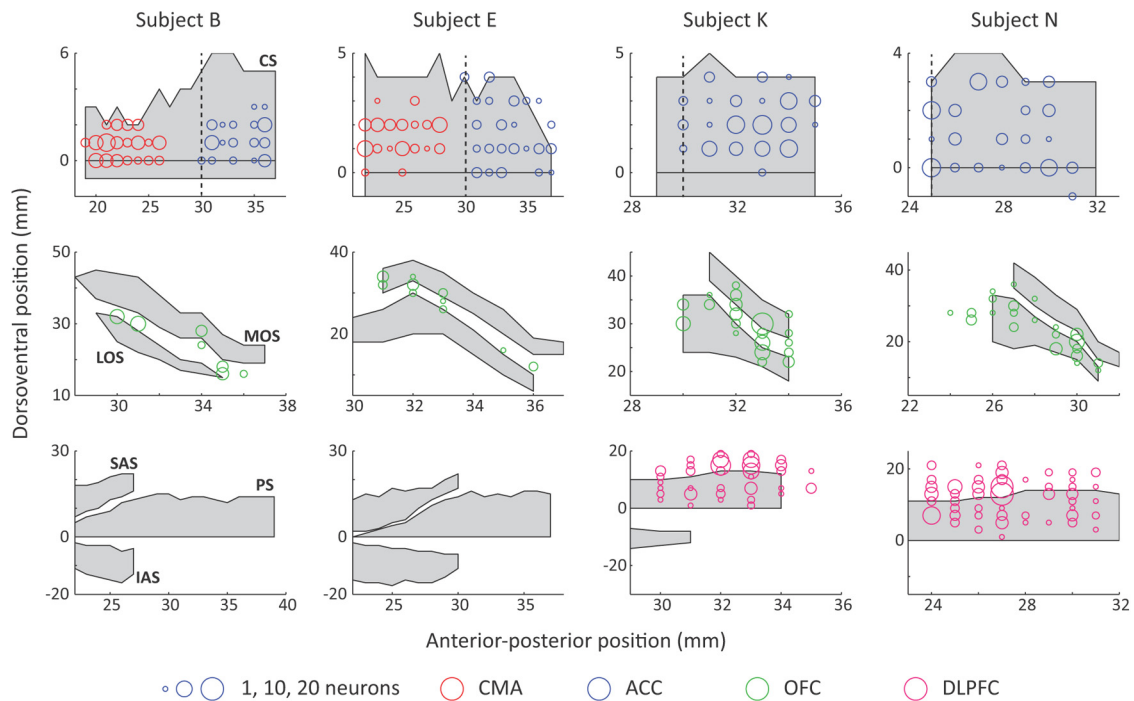


Figure 1. Flattened cortical representations illustrating recording locations from the four subjects. Gray shading indicates the position of a sulcus. The anterior–posterior (AP) position is measured relative to the interaural line. ACC and CMA recording locations were in the dorsal bank of the cingulate sulcus and the lateral–medial position was calculated relative to the fundus. We defined ACC recording locations as anterior to the genu of the corpus callosum, whereas CMA recording locations were defined as at least 2 mm posterior to the genu. The position of the genu is indicated by a dashed vertical line. In three subjects (B, E, and K) it was at AP 30 mm, whereas in the fourth subject (N) it was at AP 25 mm. OFC recording locations were largely located within and between the lateral and medial orbital sulci. They are plotted relative to the ventral bank of the principal sulcus, which is a more consistent landmark across subjects than the orbital sulci. DLPFC recording locations were located within and dorsal to the principal sulcus. CS, Cingulate sulcus; MOS, medial orbital sulcus; LOS, lateral orbital sulcus; SAS, superior arcuate sulcus; PS, principal sulcus; IAS, inferior arcuate sulcus.

obtain reward (Shidara and Richmond, 2002). These findings raise the possibility that any specialization for effort-based decision making might be better ascribed to CMA than ACC. To address these issues, we trained four monkeys to perform a novel decision-making task that enabled us to fully dissociate effort and delay costs and recorded simultaneously from primate ACC, OFC, DLPFC, and CMA neurons.

Materials and Methods

Subjects and neurophysiological procedures

Four male rhesus monkeys (*Macaca mulatta*) served as subjects (B, E, K, and N). Subjects were 5–6 years of age and weighed 8–15 kg at the time of recording. We regulated daily fluid intake to maintain motivation. Our methods for neurophysiological recording were reported in detail previously (Lara et al., 2009). We used a 1.5 T magnetic resonance imaging (MRI) scanner to ensure accurate positioning of our electrodes. Neuronal waveforms were digitized and analyzed offline (Plexon Instruments). We recorded from ACC, DLPFC, OFC, and CMA (Fig. 1) using arrays of 10–24 tungsten microelectrodes (FHC Instruments). In most sessions we recorded from at least three of the four brain areas simultaneously. We determined the approximate distance to lower electrodes from the MRI images and advanced the electrodes using custom-built, manual microdrives. We randomly sampled neurons; we did not attempt to select neurons based on responsiveness. This procedure aimed to reduce bias in our estimate of neuronal activity thereby allowing a fairer comparison of neuronal properties between the different brain regions. All procedures were in accord with the National Institute of Health guidelines and the recommendations of the University of California Berkeley Animal Care and Use Committee.

Behavioral task

Subjects were trained to make choices between pairs of stimuli (Fig. 2). In one set of trials, subjects were offered choices between stimuli that indicated both the amount of reward, as well as the amount of effort required to lift a lever to obtain reward. We held the time between effort onset and

reward delivery constant, allowing us to examine effort costs independent of delay to reward. In the other set of trials, subjects made choices between stimuli that indicated both the amount of reward as well as the delay before reward delivery. As there was no required effort in these trials, we were able to examine delay costs independent of effort. Subjects completed an average of 465 trials during a single recording session.

For subjects B and E, we used NIMH Cortex (<http://dally.nimh.nih.gov>) to control the presentation of stimuli and task contingencies. For subjects K and N, we used MonkeyLogic (Asaad and Eskandar, 2008). We monitored eye position and pupil dilation using an infrared system (ISCAN). For subjects B and E, the sampling rate was 120 Hz, whereas for subjects K and N it was 60 Hz. Each trial began with the subject fixating a central square 0.3° in width (Fig. 2A). If the subject maintained fixation within 1.8° of the cue for 1000 ms (fixation epoch), two pictures (2.5° in size) appeared at 5.0° to the left and right of fixation, drawn from one of two sets of 16 pictures (Fig. 2B). Subjects were required to maintain fixation throughout the choice epoch. However, one subject (K) was unable to do this and was allowed to freely view the pictures. After 1 s the fixation cue changed color, which instructed the subject to choose one of the pictures by making an eye movement toward it. The first set of pictures tested effort-based decisions. On selecting one of these pictures, the subject had to lift a lever against resistance (varied using a custom-built pneumatic pressure regulator) and hold the lever in the up position for 1.5 s. A specific volume of juice reward was then delivered. The second set of pictures tested delay-based decisions. On selecting one of these pictures, the subject simply had to sit and wait a specific delay until the juice reward was delivered. We varied the volume of juice delivered by adjusting its flow rate. The subject could detect this as soon as juice delivery began, thereby enabling us to detect neuronal responses related to reward onset that encoded the reward amount. A 1000 ms intertrial interval (ITI) separated each trial.

We tailored the precise costs and benefits for each subject to ensure that they received their daily fluid allotment over the course of the recording session and to ensure that they were sufficiently motivated to perform the task. Occasionally we would adjust these amounts from one recording ses-

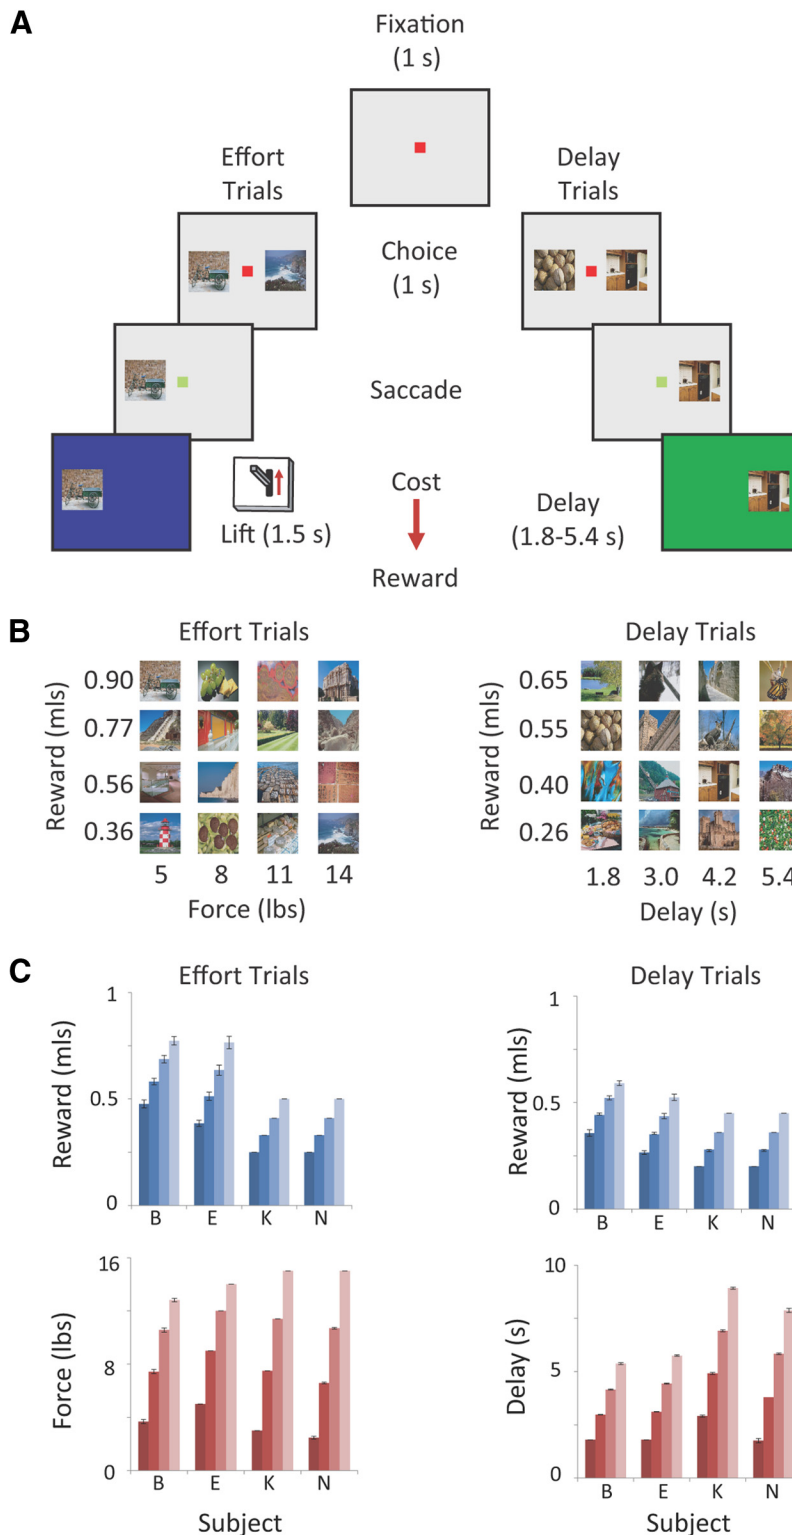


Figure 2. The behavioral task. **A**, Sequence of events in the effort and delay cost-benefit task. Subjects made choices between pairs of presented pictures. On effort trials, they had to lift a lever and hold it for 1.5 s to earn the reward. On delay trials, they simply had to wait for a specific delay before they received the reward. The amount of reward they received, the force to lift the lever and the length of delay varied depending on the picture they chose. Effort and delay trials were intermingled in the same session. **B**, The sets of pictures used on effort (left) and delay (right) trials. There were four levels of increasing benefit and four levels of increasing cost (effort or delay), comprising 16 pictures for each type of trial. We tailored the precise amounts of cost and benefit from subject to subject and even across recording sessions so as to ensure that the subjects worked for sufficient trials and received their daily aliquot of fluid within a single recording session. **C**, Mean (\pm SEM) parameters used for each subject.

sion to another, for example, increasing the reward amounts if it appeared that a subject required additional motivation. Figure 2C shows the mean parameters used for each subject. We note that for the purposes of interpreting our data, the precise costs and benefits that we used are not important because we calculated discount functions for each subject for each individual recording session (Fig. 3B).

We used the same pictures throughout recording. Within a set (effort or delay), we presented all possible picture pairings. However, not all pairings required a cost-benefit analysis. For example, if the pictures were equivalent in cost or benefit, or if one picture was better in terms of both cost and benefit, then a cost-benefit analysis was not necessary. To ensure that we had sufficient trials for our analyses, we presented those choices requiring a cost-benefit analysis 20% more frequently than would be expected by pairing pictures randomly. In all other respects the ordering of presented choices was pseudorandom and effort- and delay-based choices were randomly intermingled. If the subject broke fixation during the trial, failed to lift the lever for 1.5 s, or moved the joystick before the go cue, the subject was given a timeout penalty of 5 s and the trial was aborted. Only successfully performed trials are included in our analysis.

Statistical analysis

Behavior. To examine the influence of our experimental parameters on choice behavior, for each recording session and each set of pictures, we calculated the discounted value of the left and right picture (DV_L and DV_R respectively). We tested whether our data were best fit by either a linear, hyperbolic or exponential discount function.

Linear:

$$DV_L = B_L - w_1 C_L \quad (1)$$

$$DV_R = B_L - w_2 C_R \quad (2)$$

Hyperbolic:

$$DV_L = \frac{B_L}{1 + w_1 C_L} \quad (3)$$

$$DV_R = \frac{B_R}{1 + w_2 C_R} \quad (4)$$

Exponential:

$$DV_L = B_L \times e^{w_1 C_L} \quad (5)$$

$$DV_R = B_R \times e^{w_2 C_R} \quad (6)$$

C and B represent the costs and benefits associated with the left and right picture. We then fit a logistic regression model using the difference between the discounted values ($DV_L - DV_R$) to predict P_L , the probability that the subject will choose the left picture. We included a bias term, b , which accounted for any tendency of the subject to select the leftward picture that was independent of the pictures' values.

$$P_L = \frac{1}{1 + e^{w_3(DV_L - DV_R) - w_4 b}} \quad (7)$$

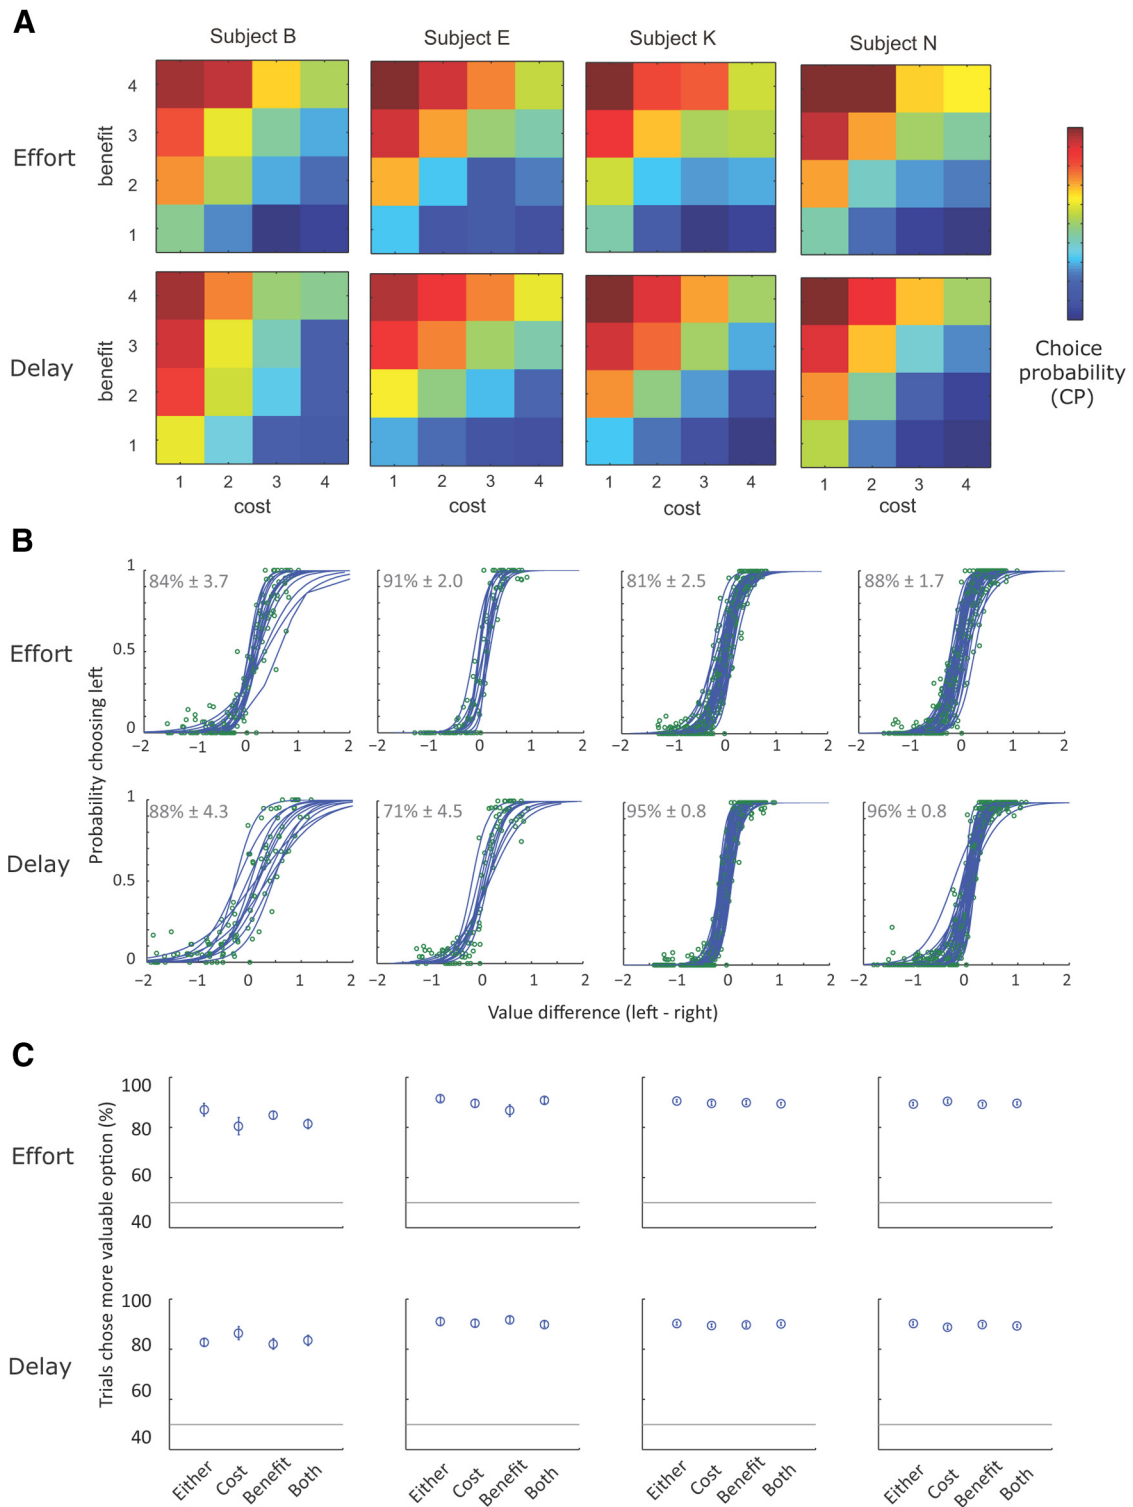


Figure 3. Behavioral performance. **A**, The probability that the subject would choose a picture increased with increasing benefit and decreased with increasing cost. The choice probabilities for 16 pictures collapsed across all recording sessions and all possible picture pairings are shown. Each cell of the matrix corresponds to a picture in Figure 2B. **B**, The ability of linear discount functions to predict subjects' choice. The x-axis indicates the difference in value between the left and right pictures as determined by Equations 1 and 2. Green circles indicate actual data, whereas blue lines indicate the results predicted by the model for each individual session. The values in the top left of each plot indicate the mean percentage of variance (\pm SEM) explained by the linear discount function as opposed to a null model in which we assumed no discounting i.e., that costs have no effect on the pictures' value. **C**, Behavioral performance as assessed by the percentage of trials on which the subject chose the more valuable option. Trials are grouped according to whether the subjects needed to attend to both the costs and benefits to choose optimally. The gray line indicates chance level performance (50%).

We estimated the free parameters in the model by determining the values that minimized the log likelihood of the model.

Neuronal encoding. We began our neuronal analysis by visualizing neuronal selectivity through the construction of spike density histo-

grams. We averaged neuronal activity across the appropriate conditions using a sliding window of 100 ms. Although it was evident in the spike density histograms that neurons were encoding some of the parameters related to the decisions, many of these parameters were highly correlated

Table 1. Parameters included in the stepwise regression

1	Discounted value	Chosen
2		Unchosen
3		Chosen + unchosen
4		Chosen–unchosen
5		Ipsilateral
6		Contralateral
7		Ipsilateral–contralateral
8	Benefit	Chosen
9		Unchosen
10		Chosen + unchosen
11		Chosen–unchosen
12		Ipsilateral
13		Contralateral
14		Ipsilateral–contralateral
15	Cost	Chosen
16		Unchosen
17		Chosen + unchosen
18		Chosen–unchosen
19		Ipsilateral
20		Contralateral
21		Ipsilateral–contralateral
22	Behavioral response	
23	Decision type	

Discounted value was calculated for each session by integrating the costs and benefits using the best fitting model from the behavioral analysis. For each parameter, we examined different ways in which the parameter could be related to the pictures, sorted according to either the picture that the subject chose or did not choose, or the spatial position of the picture relative to the recorded neuron. Note that the sum of the ipsilateral and contralateral picture is identical to the sum of the chosen and unchosen picture, and so it was not included in the model as a separate predictor. We also included the interaction between decision type and the other 22 predictors so that a total of 45 regressors were tested.

with one another. Therefore to identify what information neurons were encoding, we performed stepwise regression. For each neuron, we calculated its average firing rate in two epochs, corresponding to the first half (100–600 ms) and second half (600–1100 ms) of the choice epoch. (The 100 ms offset allowed time for visual information to reach the frontal cortex). We contrasted this activity with that recorded in the fixation period, defined as the 500 ms immediately preceding the onset of the pictures.

The first four parameters tested whether the neuron encoded the benefit associated with either the picture that the subject did or did not choose, or the spatial position of the picture relative to the recorded neuron. We also included three additional parameters that looked at how these parameters might be combined: chosen + unchosen, chosen – unchosen and ipsilateral – contralateral. Note that the sum of the ipsilateral and contralateral picture is identical to the sum of the chosen and unchosen picture, and so it was not included in the model as a separate predictor. Thus, there were a total of seven parameters related to potential ways of encoding benefits. We added analogous parameters for costs, bringing the number of parameters to 14. We also calculated the discounted value for each parameter using the best fitting model from the behavioral analysis to determine how the subject was integrating costs and benefits, thereby bringing the number of parameters to 21. We added a parameter encoding the subjects’ behavioral response (whether he chose left and right) and a dummy coded decision type parameter that indicated whether the subject was performing an effort or delay-based decision. Finally, we included the interaction between decision type and the previous 22 parameters. This enabled us to detect neurons that encoded one of the parameters but only for one type of decision. Thus, we tested 45 parameters in total (Table 1).

The first step in the stepwise regression was to regress each parameter individually against the neuron’s firing rate and determine whether any parameter significantly predicted firing rate (assessed at $p < 0.05$). If more than one parameter predicted firing rate, we determined which parameter was the best predictor (X) and the amount of variance explained by this predictor alone. For each of the remaining predictors, we computed the amount of variance explained if they were included in the

model in addition to X . To establish whether the inclusion of an additional predictor explained significantly more variance than X alone, we calculated an F value using the following equation:

$$F = \frac{(N - f - 1) \times (R_f^2 - R_r^2)}{(f - r) \times (1 - R_f^2)} \tag{8}$$

R_r^2 is the variance explained by the reduced model (in this case, X alone) and R_f^2 is the variance explained by the full model (i.e., including the additional predictor). In addition, N is the number of data points, f is the number of predictors in the full model, and r is the number of predictors in the reduced model. The significance of F was then calculated using $f - r$ and $N - f - 1$ degree of freedom. If the addition of the extra predictor variable significantly improved the amount of explained variance (assessed at $p < 0.05$) it was included in the model. We repeated this process until the addition of further variables produced no significant improvement in explained variance. In addition, at each step, before we added a new variable, we tested whether any variable included at an earlier step could be removed on the grounds that it no longer made a significant contribution (assessed at $p > 0.1$). To determine whether the proportion of neurons encoding a given parameter exceeded the proportion we would expect by chance, we performed the same stepwise analysis on the data from the fixation period. We then performed a binomial test to examine whether the proportion of selective neurons in the choice period exceed the proportion of selective neurons observed in the fixation period. We used a Bonferroni correction to correct for the multiple comparisons (the 45 parameters) involved in this procedure.

To examine the time course of neural selectivity related to the encoding of decision type we calculated the coefficient of partial determination (CPD) for this parameter. This is the amount of variance in the neuron’s firing rate that can be explained by decision type over and above the variance explained by other predictors included in the model. It is defined by:

$$CPD = \frac{(SSE_r - SSE_f)}{SSE_r} \tag{9}$$

SSE_r is the sum of squared errors in a regression model that includes all of the relevant predictor variables (as indicated by the stepwise regression) except decision type, whereas SSE_f is the sum of squared errors in a regression model that includes decision type. We calculated CPD for each neuron using a 200 ms window of neural activity at every time point in the trial. To calculate the latency at which neurons first encoded information about decision type, we determined the maximum CPD value for each neuron during the fixation period. We then used the 95th percentile of this distribution of values as our criterion for encoding the decision type during the choice epoch.

Results

Analysis of behavioral data

Choice behavior

Subjects chose a picture more often as benefit increased and cost decreased (Fig. 3A), demonstrating that they were able to weigh the costs and benefits associated with different pictures. To examine how subjects integrated costs and benefits, we estimated the extent to which the costs associated with the picture discounted the value of the benefit, and then examined how well the discounted values of the left and right pictures predicted subjects’ choice. In addition to a linear discount function, we also examined the fit of hyperbolic and exponential discount functions, which have frequently been found to better predict choice behavior (Kable and Glimcher, 2007; Hwang et al., 2009). A linear discount function provided a clearly superior fit for all subjects for both effort- and delay-based decisions (Table 2). Figure 3B illustrates the ability of the fitted discount functions to predict the subjects’ choices. We calculated the pseudo r -squared values of our models by comparing the predicted choice probabilities to a “null” model in which the value of the pictures was not dis-

Table 2. Formal model comparison of discount functions used to predict subjects' choice behavior

	B	E	K	N
Effort				
Linear	809	442	2558	2252
Hyperbolic	1049	538	2882	2912
Exponential	1016	524	2911	2843
Delay				
Linear	848	739	1596	1798
Hyperbolic	1039	817	2431	3136
Exponential	992	799	2213	2993

The values shown are Akaike Information Criterion (AIC) values, with smaller values indicating a better fit. The best fit was provided by a linear discount function for all subjects and for both effort- and delay-based decisions. Indeed, a comparison of AIC weights showed that in all cases a linear fit was at least 1×10^6 times more likely than a hyperbolic or exponential fit.

counted by the cost. Calculating the pseudo r -squared values of the logistic regression model revealed that the model explained a mean of $85 \pm 1\%$ of the variance in the subjects' choices for effort-based decisions and $92 \pm 1\%$ of the variance for delay-based decisions.

The fact that a linear discount function predicted subjects' choices better than a nonlinear discount function appears counter to the behavioral evidence that nonlinear discount functions typically provide a better fit of choice behavior (Kable and Glimcher, 2007; Hwang et al., 2009). However, our experiment used a relatively limited range of costs compared with previous studies. For example, in behavioral studies of temporal discounting in humans, delays of several months or even years are used (Thaler, 1981). It is possible that had we used a larger range of costs we would have seen a better fit for nonlinear discount functions. However, given the limited number of trials that we can collect in a single recording session, we wanted to focus on the part of the decision space where cost is maximally affecting the value of the benefit, rather than the part of the decision space where the effect of cost is reaching an asymptote. This may have biased us toward seeing a linear discount function.

Finally, we note that to motivate the subjects to complete a sufficient number of effort trials, we needed to provide a slightly larger reward on effort trials relative to delay trials (Fig. 2C). However, the difference in the mean value of the pictures across effort and delay trials was relatively small compared with the range of values for pictures within trial type. The mean difference in the value of the effort and delay pictures was equivalent to 0.03 ml of juice. In comparison, the mean range of values within effort pictures was 0.48 ml and within delay pictures was 0.47 ml.

We next grouped trials according to whether the subject could make his choice by focusing on either the cost or the benefit ("Either" trials: one option was both lower cost and higher benefit than the other option), or by focusing on solely the cost or benefit ("Cost" or "Benefit" trials: e.g., one option was the same cost but a different benefit to the other option) or whether the choice required the consideration of both the costs and benefits ("Both" trials: one option was higher cost but also higher benefit than the other option). Only this last group of trials required a cost-benefit analysis. We then calculated the percentage of trials on which the subject chose the more valuable picture (Fig. 3C). Subjects' performance was equally good regardless of whether or not a cost-benefit analysis was required: for every subject there were no significant differences in performance of the different trial types (one-way ANOVA, $p > 0.1$ in all cases). Thus, subjects appeared to attend to both the costs and benefits associated with each picture on every trial.

Table 3. Number of neurons recorded from each area in each subject

Subject	DLPPFC	OFC	ACC	CMA
B	0	28	37	41
E	0	19	54	67
K	125	98	117	0
N	185	61	112	0
Total	310	206	320	108

Movement times

One problem with many studies of effort-based decision-making is that the manipulation to increase the effort associated with a choice option inadvertently increases the delay until the subject receives a reward. For example, clambering over a higher barrier (Walton et al., 2002), not only takes more effort, but also takes more time. Such manipulations could potentially confound delay and effort discounting. Our behavioral task was explicitly designed to avoid such confounds but they could conceivably arise if it took longer for the subject to lift the lever when there was a greater load. There was little evidence that this was the case. We measured mean movement time for each level of effort from the time from that a choice target was fixated (causing the other target to disappear and indicating that the subject was free to lift the lever) until the time at which the completion of the movement was registered (which initiated the 1.5 s period that the subject had to hold the lever in position). In all subjects, there were significant differences in the time to perform movements requiring different degrees of effort (one-way ANOVA, $F > 5.1$, $p < 0.005$ in all cases). However, the pattern of effects differed across subjects. In subject N, more effortful movements took longer to complete than less effortful movements. However, in two subjects (E and K) more effortful movements were completed more quickly than less effortful movements, and in the final subject (B) there was no consistent relationship between effort and movement times. Furthermore, in all subjects, the movement times varied by only ~ 200 ms, a relatively trivial delay compared with the delays used in the delay discounting trials. Thus, differences in movement times on effort trials are unlikely to explain the subjects' valuation of the different pictures.

Analysis of neuronal data

We recorded from 944 neurons distributed across four brain areas (Fig. 1; Table 3), and examined whether neurons had a firing rate that correlated with the costs and benefits associated with the two options on each trial. We focused our analysis on the choice epoch (the 1 s period between picture onset and go cue) as this was uncontaminated by any movement or exerted effort. Neurons often encoded the benefit and/or cost associated with the presented pictures, particularly the picture that the animal eventually chose. Figure 4 illustrates three examples. To quantify the information each neuron was encoding, we performed stepwise regression analysis on neuronal data during the first and second half of the choice epoch (see Materials and Methods; Table 1).

We tested a total of 45 parameters. During the first half of the choice epoch 843/944 or 89% of neurons encoded at least one parameter, and this figure rose to 872/944 or 92% during the second half of the choice epoch. The majority of the neurons (79%) encoded no more than three variables. The most prevalent encoding in all areas indicated the type of decision the subject was performing—that is, whether the subject was currently performing an effort- or delay-based choice (Fig. 5). We will first describe encoding of value- and response-related results as these were the

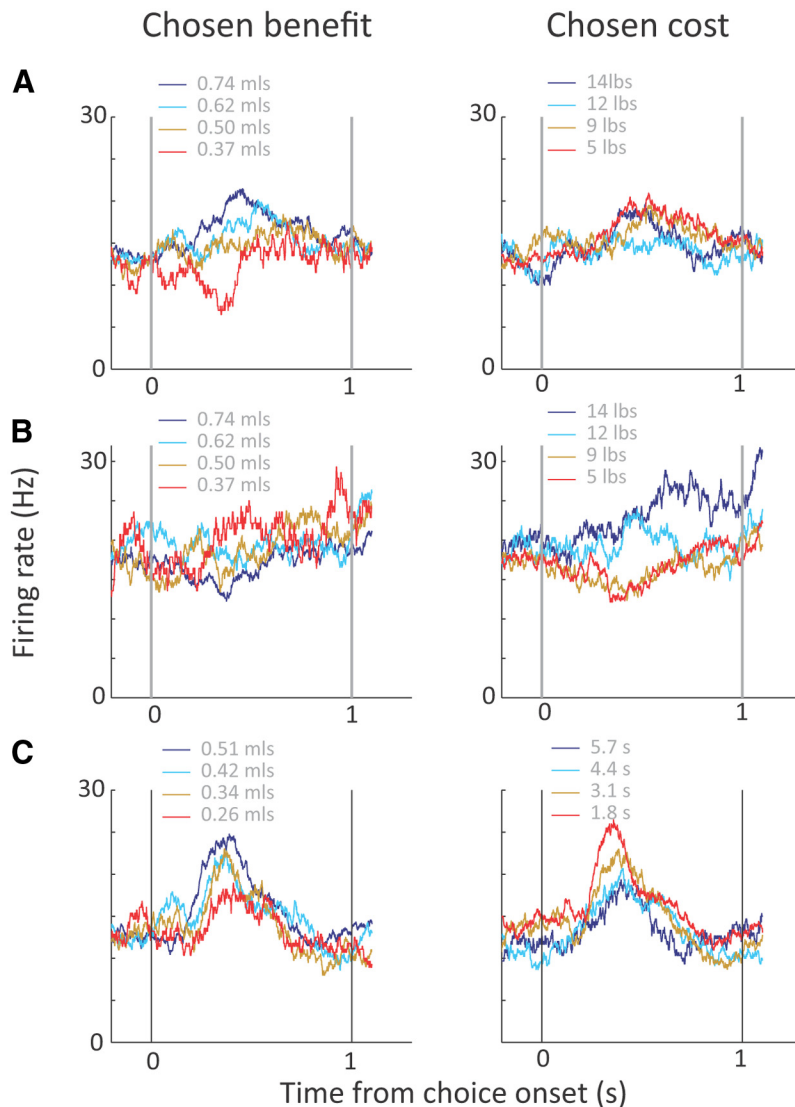


Figure 4. Spike density histograms from neurons encoding costs or benefits. The vertical gray lines indicate the onset and offset of the choice epoch. **A**, A neuron that increased its firing rate as the benefit associated with the chosen picture increased, but did not encode the effort cost. **B**, A neuron that increased its firing rate as cost increased, but was not affected by the benefit. **C**, A neuron that performed a cost-benefit analysis for delay decisions. It increased its firing rate as the benefit increased or the delay decreased. All three of these examples were recorded in ACC, the area in which we typically observed the most robust encoding of costs and benefits.

initial focus of our investigation, and we will then return to describe the encoding of decision type in more detail.

Encoding of costs, benefits, and responses

With respect to value coding, ACC differed from the other areas in that a sizeable proportion of neurons integrated information about costs and benefits to encode the value of the chosen picture. Across both epochs, 74/320 or 23% of ACC neurons encoded chosen value, significantly more than the proportion in OFC (18/206 or 9%, $\chi^2 = 17, p < 5 \times 10^{-5}$) or DLPFC (44/310 or 14%, $\chi^2 = 7.7, p < 0.01$) and a trend toward significance compared with the proportion in CMA (15/108 or 14%, $\chi^2 = 3.6, p = 0.056$). A significant population of neurons in both ACC and OFC encoded chosen cost—that is, they encoded the cost regardless of cost decision type. This suggests that the neural representation of decision costs is in an abstract or general form. However, we did not find any significant population of neurons in any of the four brain areas that encoded chosen benefit. Thus, the neural

representation of reward, at least in frontal cortex, is only encoded as it relates to the decision costs, i.e., as a signal that reflects the discounted value of the reward based on the integration of costs and benefits.

Other variables were also encoded, albeit rather modestly. Neurons tended to be biased toward encoding the variables associated with either the picture the subject intended to choose or the picture located contralaterally to the position of the recorded neuron. DLPFC tended to be biased toward encoding the behavioral response, either on its own or in an interaction with decision type (e.g., the neuron encoded the behavioral response but only for effort-based decisions). In total, across both epochs, 140/310 or 45% of DLPFC neurons encoded something about the behavioral response, which was significantly more than the proportion in either ACC (85/320 or 27%, $\chi^2 = 23, p < 5 \times 10^{-6}$), CMA (24/108 or 22%, $\chi^2 = 17, p < 5 \times 10^{-5}$), or OFC (44/206 or 21%, $\chi^2 = 30, p < 1 \times 10^{-7}$).

Neuronal dynamics of decision type encoding

Across all four brain areas, 638/944 (68%) of neurons encoded the decision type in at least one of the choice epochs. Although the prevalence of such neurons was significantly higher in ACC (250/320 or 78%) and CMA (83/108 or 77%), there was still a sizeable majority of neurons encoding decision type in DLPFC (212/310 or 68%) and OFC (117/206 or 57%, $\chi^2 > 6.7, p < 0.01$ for all relevant pairwise comparisons). Neurons exhibited a variety of dynamics in these signals. Some showed tonic encoding of decision type throughout the choice epoch (Fig. 6A,B), whereas others showed relatively phasic encoding (Fig. 6C,D). In DLPFC, OFC, and ACC there was an approximately even proportion of neurons

that showed a higher firing rate on either effort or delay trials. During the first half of the choice epoch, of those neurons that encoded the decision type 97/177 or 55% of ACC neurons responded more strongly on effort trials, compared with 84/154 or 54% of DLPFC neurons and 37/80 or 46% of OFC neurons. None of these proportions were significantly different from what would be expected by chance (binomial test, $p > 0.1$). We saw similar results in the second half of the choice epoch (binomial test, $p > 0.1$ in all cases). However, neurons in CMA were biased toward encoding effort. During the first half of the choice epoch, 28/41 or 68% of the decision type neurons in CMA responded more strongly on effort trials compared with delay trials (binomial test, $p < 0.05$) and this bias become even more pronounced in the second half of the choice epoch (53/74 or 72% of neurons responded more strongly on effort trials, binomial test, $p < 0.0005$).

We next examined the strength and time course of these responses at the population level. For those neurons that encoded

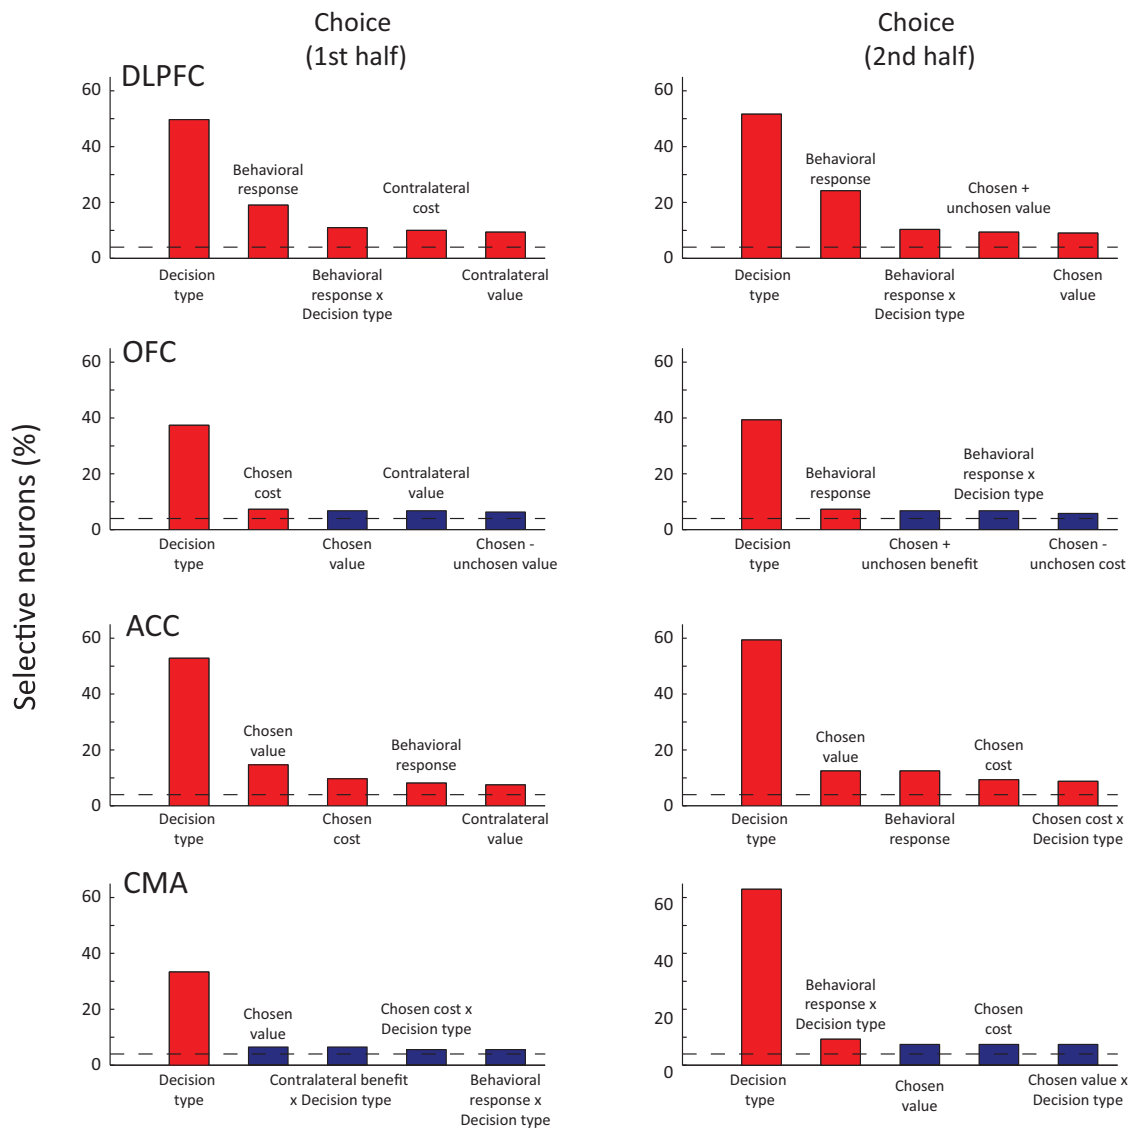


Figure 5. Percentage of neurons that were identified, using stepwise regression, as encoding specific variables during the first or second half of the choice epoch. For clarity, we only plotted the five variables that were most frequently encoded, because beyond that there were $<10\%$ of the neurons encoding a given variable. The horizontal dotted line represents chance levels of neuronal selectivity which we determined by calculating the mean percentage of neurons that were quantified as selective using the same analysis method during the fixation period i.e., before any decision-related information was revealed to the subject. Red bars indicate that the proportion of selective neurons is significantly higher than the average proportion of selective neurons in the fixation period, assessed using a binomial test at $p < 0.05$, with a Bonferroni correction for multiple comparisons (the 45 parameters).

decision type (as defined by the sliding CPD analysis; see Fig. 8), we divided them into two groups based on whether they showed a higher firing rate on effort or delay trials. We then plotted their mean firing rate on effort and delay trials relative to their mean firing rate during the fixation period (Fig. 7). For each neuron we calculated its mean firing rate on effort and delay trials in a 500 ms epoch that was centered on the brain area's mean latency to encode decision type (see Materials and Methods). We then performed a two-way ANOVA, with each neuron's mean firing rate as the dependent variable and two independent variables: decision (effort or delay) and preference (whether the neuron showed a higher firing rate on effort or delay trials).

In each area there was a significant decision \times preference interaction, but this was to be expected, given how we had grouped the data. More interesting was the pattern of simple effects, particularly how the encoding of the Decision varied between effort- and delay-preferring neurons. In LPFC and OFC, there was a strong significant difference in neuronal firing on

delay trials between the delay- and effort-preferring neurons (LPFC: $F_{(1,160)} = 18.5, p < 5 \times 10^{-5}$; OFC: $F_{(1,90)} = 19, p < 5 \times 10^{-5}$) but only a weak difference between these two groups of neurons on effort trials (LPFC: $F_{(1,160)} = 5.5, p < 0.05$, OFC: $F_{(1,90)} = 3.5, p < 0.1$). In other words, in both LPFC and OFC, the encoding of decision type was driven almost entirely by modulation on the delay trials, with one group of neurons showing inhibition on delay trials and another group of neurons showing excitation. The reverse was true in CMA: there was a significant difference in neuronal firing on effort trials between the delay- and effort-preferring neurons ($F_{(1,56)} = 33, p < 5 \times 10^{-7}$), but no difference on delay trials ($F_{(1,56)} < 1, p > 0.1$). In CMA, the encoding of decision type was driven entirely by modulation on effort trials, with one group of neurons showing excitation and the other group showing inhibition. Finally, both contrasts were equally significant in ACC (effect of neuronal Preference on delay trials: $F_{(1,197)} = 17, p < 0.0005$, effort trials: $F_{(1,197)} = 15, p < 0.0005$). Thus, ACC was the only area that contained populations

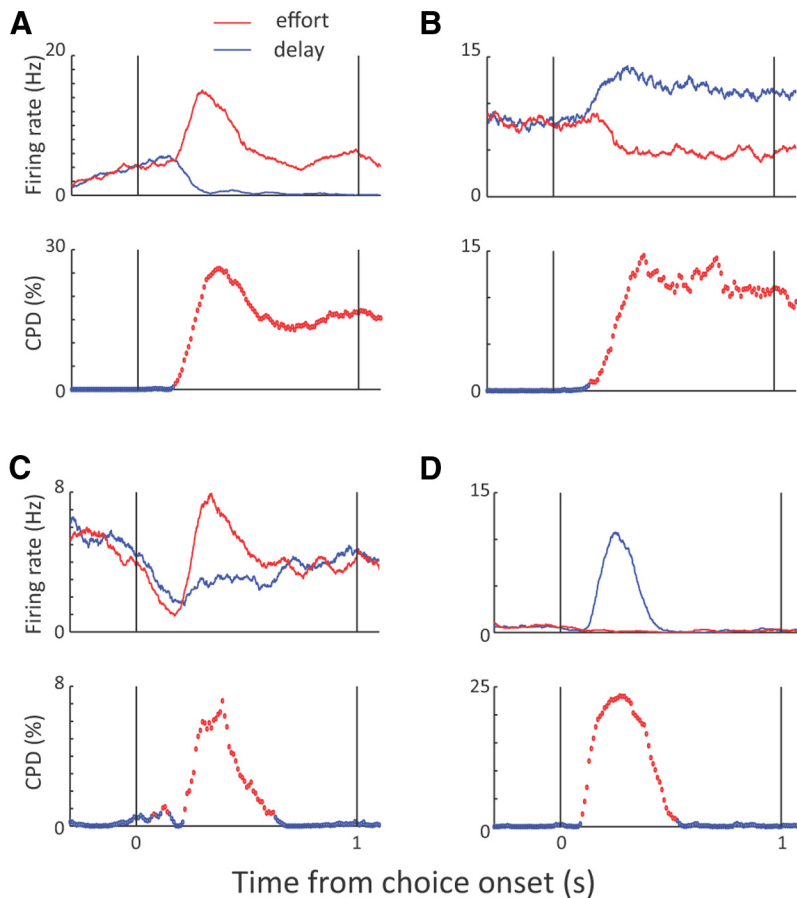


Figure 6. Examples of neuronal firing during the choice epoch for effort- or delay-based decisions. In each example, the top plot shows the spike density histogram, with the red line indicating the neuron’s mean firing rate on effort trials and the blue line indicating the neuron’s mean firing rate on delay trials. The bottom plot shows the coefficient of partial determination for the decision type parameter. This measure indicates the amount of variance in the neuron’s firing rate that is accounted for by the decision type and cannot be explained by any of the other parameters in the model (see Materials and Methods). Red data points indicate that the decision type significantly predicts neuronal firing rate ($p < 0.05$ corrected for multiple comparisons). Neuron **A** was recorded from OFC, whereas neurons **B–D** were recorded from ACC.

of neurons that were strongly modulated by both types of decision costs.

Although we found much stronger encoding of decision type relative to the encoding of value, it might be possible that value coding was dependent on decision type. For example, neurons might encode the value of the chosen picture, but only on effort trials, which would have been evident as an interaction between chosen value and decision type. In general, there was little evidence that encoding of other variables depended on the type of decision. There were three exceptions to this. In DLPFC, 58/310 or 19% of neurons encoded the behavioral response for one decision type but not the other (behavioral response \times decision type interaction) in at least one of the choice epochs, whereas 14/108 or 13% of CMA neurons encoded this same information. In ACC, 45/283 or 16% of neurons encoded the chosen cost for one decision type but not the other (chosen cost \times decision type interaction). To examine whether these neurons favored either effort or delay decisions we performed two separate regressions. For each neuron, we included the variables identified by the stepwise regression, and performed the regression separately on effort and delay trials. We then compared the absolute value of the standardized β coefficient for the variable that interacted with decision type on effort and delay trials. There was no systematic difference in the size of the β coefficients for the interacting vari-

able across the two types of decision (t test, $p > 0.1$ in all three cases). In other words, although individual neurons frequently showed stronger modulation for a given variable on delay or effort trials, across the population, the stronger modulation was equally likely to occur for either type of decision.

Cingulate motor area

A major difference between effort and delay trials is that effort trials require the animal to plan an additional motor response (lifting the lever) whereas delay trials simply require the animal to wait. Thus, neural activity differentiating effort and delay trials may simply reflect motor preparation processes on effort trials. This explanation appears to account for the neuronal activity in CMA. As we already noted, for those neurons that encoded decision type, there was an even split between neurons encoding either effort or delay in DLPFC, OFC, and ACC, but in CMA, the majority of neurons encoded effort. In addition, in CMA encoding of any variable other than decision type was very weak (Fig. 5). The only proportion to exceed the amount of encoding present in the fixation period was the interaction between decision type and response during the second half of the choice epoch, and even this accounted for only a small minority of neurons (10/108 or 9%).

If decision type encoding in CMA actually reflected motor preparation processes we hypothesized that it would have a later onset than neural activity that was responsible for identifying the type of decision with which the subject was faced.

To examine whether there were differ-

ences between the areas in terms of the onset of decision type encoding, we calculated the CPD for decision type at each point in the trial. Figure 8 shows that decision type encoding did appear to occur later in CMA relative to DLPFC, OFC, and ACC. To calculate the latency at which neurons first encoded information about decision type, for each neuron we determined the time at which the CPD exceeded 95% of the CPD values during the fixation period. Because the latency to detect a signal is correlated with the strength of the signal, we included the maximum value of the CPD as a covariate in the analysis. The mean latency at which decision type was encoded was at least 120 ms later in CMA (596 ms \pm 26 ms) compared with the other three areas (DLPFC: 417 ms \pm 15 ms, OFC: 462 ms \pm 20 ms, ACC: 474 ms \pm 14 ms). A one-way ANCOVA revealed significant differences between the areas in terms of the mean latency to encode the decision type ($F_{(3,501)} = 10, p < 5 \times 10^{-6}$). Tukey’s Honestly Significant Difference *post hoc* tests revealed that CMA was significantly slower to encode the decision type than either DLPFC ($p < 0.001$) or OFC ($p < 0.01$) and a trend toward being significantly slower than ACC ($p < 0.1$). Decision type was also encoded significantly more quickly in DLPFC relative to ACC ($p < 0.05$).

Thus, the data we collected from CMA, a predominately motor area, differed markedly from the data we collected from

ACC, OFC, and DLPFC. Although CMA neurons did encode decision type, neural activity was consistently higher on effort trials, it occurred later in the choice epoch and it did not reflect the costs and benefits associated with the choice. In turn, this highlights the aspects of the neuronal activity in ACC, OFC, and DLPFC that are not consistent with encoding the motor response on effort trials, namely that the selectivity in these frontal areas appears relatively quickly, it consists of neurons that prefer delay trials as well as those that prefer effort trials, and it is one of many variables about the choice that these neurons encode. Furthermore, we note that encoding of the more immediate motor response, the eye movement to indicate which picture the subject wanted to choose, was much weaker than encoding of decision type in all three areas, and indeed was completely absent in OFC. Thus, encoding of decision type in ACC, OFC, and DLPFC appears to signal whether the subject is performing an effort- or delay-based decision and does not reflect the differing motor demands of the two types of trial.

Discussion

We trained monkeys to make choices requiring the integration of reward with effort or delay costs and found sparse value coding across four frontal areas. This contrasts with previous studies showing OFC neurons representing an integrated value signal (Padoa-Schioppa and Assad, 2006, 2008; Padoa-Schioppa, 2009). Instead, neurons predominantly encoded the type of cost decision the subject faced. We previously found that OFC neurons encoded multiple decision attributes (reward probability, magnitude, and effort costs) when each attribute was tested independently, which suggested that OFC may integrate multiple different decision attributes to calculate an abstract value signal to guide subject's choice (Kennerley et al., 2009; Wallis and Kennerley, 2010); a view echoing others (Rangel and Hare, 2010; Padoa-Schioppa, 2011). However, not all data has been consistent with this view. Rodent studies have reported separate representations of delay costs and reward magnitude in OFC (Roesch et al., 2006), whereas in the monkey, OFC contains separate populations of neurons encoding reward risk and reward magnitude (O'Neill and Schultz, 2010). Even in our previous study, although we identified some OFC neurons that encoded value across three decision attributes, the majority of OFC neurons in fact encoded value for only one or two decision attributes (Kennerley et al., 2009).

Previous studies that have identified integrated value signals in OFC have typically manipulated the type and magnitude of a juice reward (Padoa-Schioppa and Assad, 2006, 2008; Padoa-Schioppa, 2009). OFC processes taste and smell information to

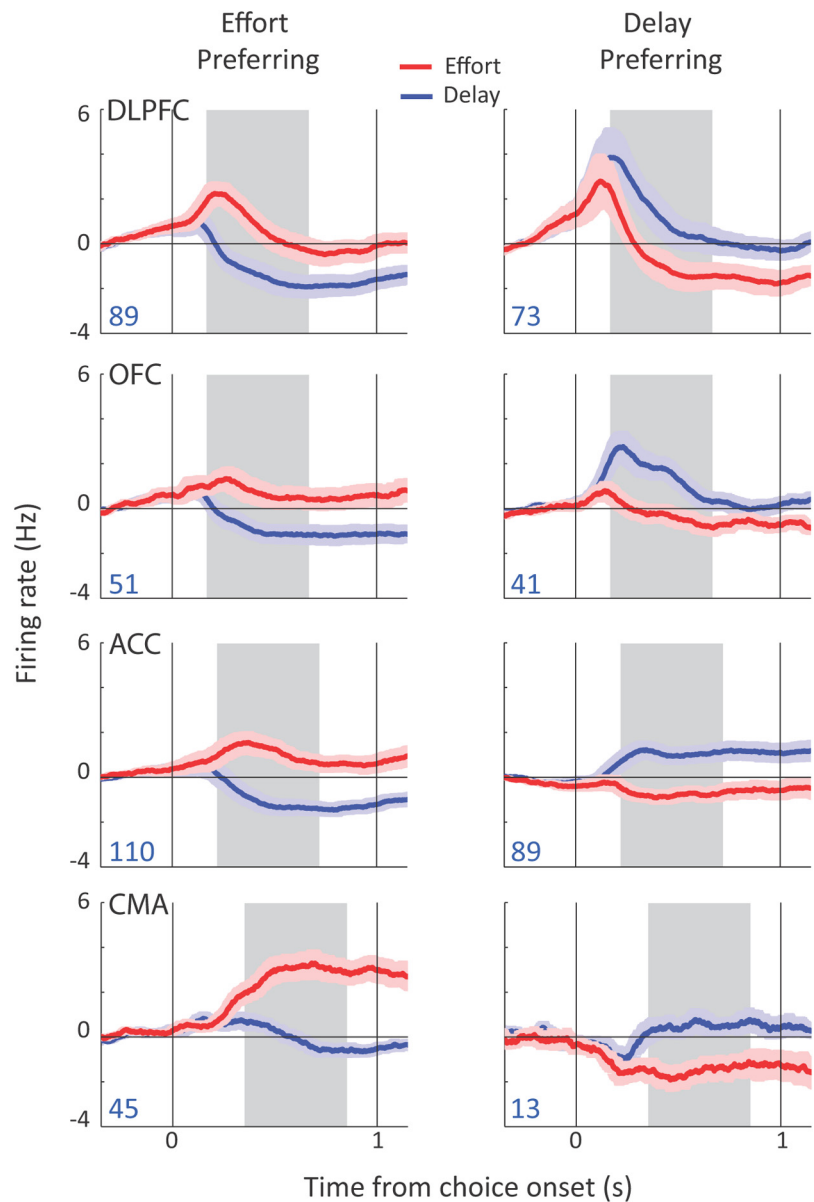


Figure 7. Time course of activity encoding the type of decision for effort-preferring and delay-preferring neuronal populations. Each plot shows the mean firing rate relative to baseline \pm SEM. The gray shading illustrates the 500 ms epoch used for analysis which is centered on the mean latency at which the information is encoded in each area. The blue numbers indicate the number of neurons included in each plot. Neurons were defined as encoding decision type based on reaching criterion on the sliding coefficient of partial determination analysis (see Fig. 8).

construct flavor representations (Small et al., 2007; Grabenhorst et al., 2008; Wu et al., 2012). It is possible that OFC is responsible for integrating only a subset of decision attributes, perhaps those related to the sensory characteristics of reinforcers (Burke et al., 2008; Mainen and Kepecs, 2009). Therefore, type and magnitude of juice may have been ideal parameters for detecting OFC value signals. More generally, decision-making may be a relatively distributed process with different brain areas responsible for encoding and/or integrating different types of decision attributes.

We found stronger encoding of integrated value in ACC rather than OFC, consistent with several studies reporting value signals in ACC (Amiez et al., 2006; Hayden et al., 2009; Kennerley et al., 2011; Amemori and Graybiel, 2012; Cai and Padoa-Schioppa, 2012). In some models of decision-making, the process of choosing between alternatives occurs between different action

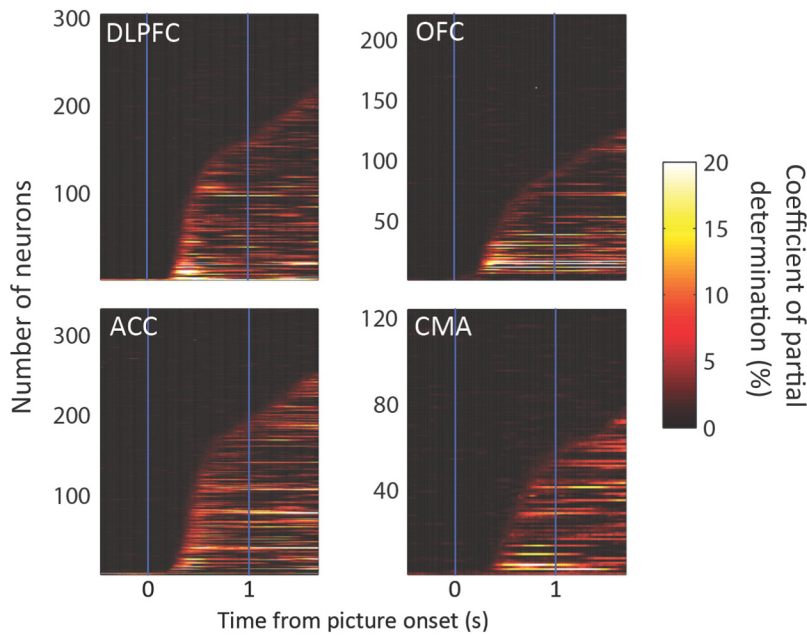


Figure 8. Encoding of decision type across the population. Each horizontal line on the plot indicates the selectivity of a single neuron as measured using the coefficient of partial determination (see Materials and Methods). Neurons have been sorted according to the latency at which they first show selectivity. The vertical blue lines indicate the onset and offset of the choice epoch.

(rather than stimulus) values (Cisek, 2012; Rushworth et al., 2012). This predicts that integrated value signals may be best represented in relatively downstream structures, such as ACC, which is more strongly connected with motor areas compared with OFC (Van Hoesen et al., 1993; Carmichael and Price, 1995). Interestingly, there were no value signals in CMA, suggesting that CMA is responsible for movement implementation rather than the value-based decision.

We also found that OFC and ACC neurons encoded a chosen cost signal, whereas DLPFC encoded the cost associated with the contralateral picture. These findings suggest that frontal neurons may encode decision costs in an abstract manner, independent of the specific way in which the cost is manipulated. Surprisingly, we did not find any population of neurons that encoded chosen reward independent of the costs. Thus, unlike decision costs, the neural representation of reward in frontal cortex appears to depend on the decision costs or decision context.

Functional dissociation of effort- and delay-based decisions

Based on lesion studies in rats and neuroimaging studies in humans, we hypothesized that ACC and OFC may be specialized for effort- and delay-based decisions, respectively. The strongest evidence for specialization of decision cost coding would be to see neurons encoding an interaction of the costs and/or benefits with decision type. There were no such neurons in OFC and, in ACC; although there were neurons that showed an interaction between chosen cost and decision type, there was no evidence that this population of neurons favored encoding costs on effort trials. Nevertheless, the different brain areas did show different degrees of modulation in overall firing rate depending on the type of decision. OFC and DLPFC showed stronger modulation on delay trials, consistent with evidence suggesting a preferential involvement of these areas in delay-based decisions (McClure et al., 2004; Roesch and Olson, 2005; Kable and Glimcher, 2007; McClure et al., 2007; Kim et al., 2008; Prévost et al., 2010). However, we found the strongest modulation of firing rate on effort trials in

CMA rather than ACC. We interpreted this result as related to the increasing motor demands required as effort increases, consistent with the relatively late onset of neuronal selectivity in CMA.

Our results suggest reappraisal of previous effort-related results that were attributed to ACC. First, although the homology between human and monkey cingulate areas is strong, the homology between primate and rodent is not straightforward (Rushworth et al., 2004; Wise, 2008). Thus, studies in rodents that implicated ACC in effort-based decision making (Walton et al., 2002; Rudebeck et al., 2006; Hillman and Bilkey, 2010) may have explored regions of rodent ACC that are more homologous to CMA in primates. Second, effort-based activations in neuroimaging studies are caudal to the genu of the corpus callosum in and around the location of human CMA (Crosson et al., 2009; Prévost et al., 2010; Kurniawan et al., 2013). In the present study, all the value coding ascribed to ACC was rostral to the genu of the callosum, whereas the firing rate modulations

in CMA noted on effort trials were caudal to the genu. Further, ACC/CMA BOLD activations correlate with effort exerted, rather than value (Prévost et al., 2010; Kurniawan et al., 2013), consistent with an explanation in motor terms rather than decision-making. It is also noteworthy that ACC neurons in rodent do not encode reward value when no effort is required (Hillman and Bilkey, 2010). Together, we propose that more rostral parts of primate ACC have a prominent role in decision making and appear to encode decision value both abstractly (Kennerley et al., 2009, 2011) and as an integrated value signal (current results), whereas the caudal CMA region encodes much less information about the value of choices (Kennerley et al., 2009) but may be critical for motivating and executing a course of action.

Representation of decision type

Despite relatively sparse value coding across all four frontal areas, all areas exhibited prominent decision type selectivity. The question remains as to why frontal neurons encode decision type. More specifically, given the surprising lack of value signals in OFC, what function is OFC performing if it is not calculating value? OFC is argued to be responsible for encoding associations between stimuli and the outcomes they predict (Ostlund and Balleine, 2007; Balleine and O’Doherty, 2010; Camille et al., 2011; McDannald et al., 2011; Luk and Wallis, 2013). In addition, prefrontal cortex is thought to be responsible for the abstraction of stimuli and responses into high-level categories, rules and behavioral “sets” (Miller et al., 2003) and OFC and DLPFC neurons encode abstract rules (Wallis et al., 2001). Putting these findings together, a function of OFC may be a dimensionality reduction of stimulus space into important decision categories or sets based on predicted outcomes. Encoding of decision type was also the most prevalent encoding in DLPFC and ACC, suggesting that the task is parsed in a similar way across multiple frontal areas. Indeed, recent results have emphasized the role of DLPFC in parsing the structure of behavioral tasks (Schapiro et al., 2013).

These ideas might help to explain why value coding may or may not be observed in OFC. Value becomes one of many dimensions of an outcome that might help categorize the stimulus space. Furthermore, in the current task, it is not a particularly useful dimension for this purpose. The value of each of the pictures is uniquely specified by a combination of the costs and benefits associated with that picture and so value does not condense the stimulus set. In contrast, decision type categorizes the stimuli into two sets, based on whether they predict effort- or delay-based decisions. Moreover, we presented every possible combination of picture pairs within each decision cost, presenting subjects with 240 unique choice value contexts. In comparison, our previous experiment (Kennerley et al., 2009) grouped pictures according to decision type and varied their associated outcomes by a single attribute, leading to only 24 unique choice contexts. In this case, OFC neurons encoded both the value of the picture as well as the decision type. Thus, as the number of possible choice combinations increases, value coding may shift away from a single neuron representation in OFC. The brain may encode the value of specific pictures in other areas, such as the tail of the caudate (Kim and Hikosaka, 2013) or perirhinal cortex (Ohya et al., 2012), or the information may remain in OFC, but rely on higher dimensional representations at the population level (Rigotti et al., 2013; Shenoy et al., 2013). Finally, it is notable that those tasks that observed the clearest value signals in OFC (Padoa-Schioppa and Assad, 2006, 2008; Padoa-Schioppa, 2009) used visual stimuli that directly correlated with the type and value of the outcomes which may also bias OFC to categorizing the visual stimuli according to their value.

Recent computational accounts of OFC function reached similar conclusions (Takahashi et al., 2011). OFC lesions in rats are more consistent with a loss of “state” information, which corresponds to information about the task that enables more accurate value predictions, rather than value information per se. Thus, although the value of predicted outcomes is a component of the information that OFC represents about sensory stimuli, it is perhaps premature to conclude that this is its primary function. In straightforward tasks, where value is the only manipulation, it is the main feature encoded by OFC neurons. However, as tasks become more complex it appears that OFC encodes higher-level aspects of the task, such as the decision set that is guiding choice.

References

- Akkal D, Bioulac B, Audin J, Burbaud P (2002) Comparison of neuronal activity in the rostral supplementary and cingulate motor areas during a task with cognitive and motor demands. *Eur J Neurosci* 15:887–904. [CrossRef Medline](#)
- Amemori K, Graybiel AM (2012) Localized microstimulation of primate pregenual cingulate cortex induces negative decision-making. *Nat Neurosci* 15:776–785. [CrossRef Medline](#)
- Amiez C, Joseph JP, Procyk E (2006) Reward encoding in the monkey anterior cingulate cortex. *Cereb Cortex* 16:1040–1055. [CrossRef Medline](#)
- Asaad WF, Eskandar EN (2008) A flexible software tool for temporally-precise behavioral control in Matlab. *J Neurosci Methods* 174:245–258. [CrossRef Medline](#)
- Balleine BW, O’Doherty JP (2010) Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action. *Neuropsychopharmacology* 35:48–69. [CrossRef Medline](#)
- Burke KA, Franz TM, Miller DN, Schoenbaum G (2008) The role of the orbitofrontal cortex in the pursuit of happiness and more specific rewards. *Nature* 454:340–344. [CrossRef Medline](#)
- Cai X, Padoa-Schioppa C (2012) Neuronal encoding of subjective value in dorsal and ventral anterior cingulate cortex. *J Neurosci* 32:3791–3808. [CrossRef Medline](#)
- Camille N, Tsuchida A, Fellows LK (2011) Double dissociation of stimulus-value and action-value learning in humans with orbitofrontal or anterior cingulate cortex damage. *J Neurosci* 31:15048–15052. [CrossRef Medline](#)
- Carmichael ST, Price JL (1995) Sensory and premotor connections of the orbital and medial prefrontal cortex of macaque monkeys. *J Comp Neurol* 363:642–664. [CrossRef Medline](#)
- Cisek P (2012) Making decisions through a distributed consensus. *Curr Opin Neurobiol* 22:927–936. [CrossRef Medline](#)
- Clithero JA, Rangel A (2013) Informatic parcellation of the network involved in the computation of subjective value. *Social Cogn Affect Neurosci*, in press. [CrossRef Medline](#)
- Crosson PL, Walton ME, O’Reilly JX, Behrens TE, Rushworth MF (2009) Effort-based cost-benefit valuation and the human brain. *J Neurosci* 29:4531–4541. [CrossRef Medline](#)
- Gan JO, Walton ME, Phillips PE (2010) Dissociable cost and benefit encoding of future rewards by mesolimbic dopamine. *Nat Neurosci* 13:25–27. [CrossRef Medline](#)
- Glimcher PW, Dorris MC, Bayer HM (2005) Physiological utility theory and the neuroeconomics of choice. *Games Econ Behav* 52:213–256. [CrossRef Medline](#)
- Grabenhorst F, Rolls ET, Bilderbeck A (2008) How cognition modulates affective responses to taste and flavor: top-down influences on the orbitofrontal and pregenual cingulate cortices. *Cereb Cortex* 18:1549–1559. [CrossRef Medline](#)
- Hayden BY, Pearson JM, Platt ML (2009) Fictive reward signals in the anterior cingulate cortex. *Science* 324:948–950. [CrossRef Medline](#)
- Hayden BY, Pearson JM, Platt ML (2011) Neuronal basis of sequential foraging decisions in a patchy environment. *Nat Neurosci* 14:933–939. [CrossRef Medline](#)
- Hillman KL, Bilkey DK (2010) Neurons in the rat anterior cingulate cortex dynamically encode cost-benefit in a spatial decision-making task. *J Neurosci* 30:7705–7713. [CrossRef Medline](#)
- Hwang J, Kim S, Lee D (2009) Temporal discounting and inter-temporal choice in rhesus monkeys. *Front Behav Neurosci* 3:9. [CrossRef Medline](#)
- Kable JW, Glimcher PW (2007) The neural correlates of subjective value during intertemporal choice. *Nat Neurosci* 10:1625–1633. [CrossRef Medline](#)
- Kennerley SW (2012) Is the reward really worth it? *Nat Neurosci* 15:647–649. [CrossRef Medline](#)
- Kennerley SW, Walton ME, Behrens TE, Buckley MJ, Rushworth MF (2006) Optimal decision making and the anterior cingulate cortex. *Nat Neurosci* 9:940–947. [CrossRef Medline](#)
- Kennerley SW, Dahmubed AF, Lara AH, Wallis JD (2009) Neurons in the frontal lobe encode the value of multiple decision variables. *J Cogn Neurosci* 21:1162–1178. [CrossRef Medline](#)
- Kennerley SW, Behrens TE, Wallis JD (2011) Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nat Neurosci* 14:1581–1589. [CrossRef Medline](#)
- Kim HF, Hikosaka O (2013) Distinct basal ganglia circuits controlling behaviors guided by flexible and stable values. *Neuron* 79:1001–1010. [CrossRef Medline](#)
- Kim S, Hwang J, Lee D (2008) Prefrontal coding of temporally discounted values during intertemporal choice. *Neuron* 59:161–172. [CrossRef Medline](#)
- Kurniawan IT, Guitart-Masip M, Dayan P, Dolan RJ (2013) Effort and valuation in the brain: the effects of anticipation and execution. *J Neurosci* 33:6160–6169. [CrossRef Medline](#)
- Lara AH, Kennerley SW, Wallis JD (2009) Encoding of gustatory working memory by orbitofrontal neurons. *J Neurosci* 29:765–774. [CrossRef Medline](#)
- Levy DJ, Glimcher PW (2012) The root of all value: a neural common currency for choice. *Curr Opin Neurobiol* 22:1027–1038. [CrossRef Medline](#)
- Luk CH, Wallis JD (2013) Choice coding in frontal cortex during stimulus-guided or action-guided decision-making. *J Neurosci* 33:1864–1871. [CrossRef Medline](#)
- Mainen ZF, Kepecs A (2009) Neural representation of behavioral outcomes in the orbitofrontal cortex. *Curr Opin Neurobiol* 19:84–91. [CrossRef Medline](#)
- McClure SM, Laibson DI, Loewenstein G, Cohen JD (2004) Separate neural systems value immediate and delayed monetary rewards. *Science* 306:503–507. [CrossRef Medline](#)
- McClure SM, Ericson KM, Laibson DI, Loewenstein G, Cohen JD (2007)

- Time discounting for primary rewards. *J Neurosci* 27:5796–5804. [CrossRef Medline](#)
- McDannald MA, Lucantonio F, Burke KA, Niv Y, Schoenbaum G (2011) Ventral striatum and orbitofrontal cortex are both required for model-based, but not model-free, reinforcement learning. *J Neurosci* 31:2700–2705. [CrossRef Medline](#)
- Miller EK, Nieder A, Freedman DJ, Wallis JD (2003) Neural correlates of categories and concepts. *Curr Opin Neurobiol* 13:198–203. [CrossRef Medline](#)
- Noonan MP, Walton ME, Behrens TE, Sallet J, Buckley MJ, Rushworth MF (2010) Separate value comparison and learning mechanisms in macaque medial and lateral orbitofrontal cortex. *Proc Natl Acad Sci U S A* 107:20547–20552. [CrossRef Medline](#)
- Ohyama K, Sugase-Miyamoto Y, Matsumoto N, Shidara M, Sato C (2012) Stimulus-related activity during conditional associations in monkey perirhinal cortex neurons depends on upcoming reward outcome. *J Neurosci* 32:17407–17419. [CrossRef Medline](#)
- O'Neill M, Schultz W (2010) Coding of reward risk by orbitofrontal neurons is mostly distinct from coding of reward value. *Neuron* 68:789–800. [CrossRef Medline](#)
- Ostlund SB, Balleine BW (2007) Orbitofrontal cortex mediates outcome encoding in Pavlovian but not instrumental conditioning. *J Neurosci* 27:4819–4825. [CrossRef Medline](#)
- Padoa-Schioppa C (2009) Range-adapting representation of economic value in the orbitofrontal cortex. *J Neurosci* 29:14004–14014. [CrossRef Medline](#)
- Padoa-Schioppa C (2011) Neurobiology of economic choice: a good-based model. *Annu Rev Neurosci* 34:333–359. [CrossRef Medline](#)
- Padoa-Schioppa C, Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. *Nature* 441:223–226. [CrossRef Medline](#)
- Padoa-Schioppa C, Assad JA (2008) The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nat Neurosci* 11:95–102. [CrossRef Medline](#)
- Prévost C, Pessiglione M, Météreau E, Cléry-Melin ML, Dreher JC (2010) Separate valuation subsystems for delay and effort decision costs. *J Neurosci* 30:14080–14090. [CrossRef Medline](#)
- Rangel A, Hare T (2010) Neural computations associated with goal-directed choice. *Curr Opin Neurobiol* 20:262–270. [CrossRef Medline](#)
- Rigotti M, Barak O, Warden MR, Wang XJ, Daw ND, Miller EK, Fusi S (2013) The importance of mixed selectivity in complex cognitive tasks. *Nature* 497:585–590. [CrossRef Medline](#)
- Roesch MR, Olson CR (2005) Neuronal activity in primate orbitofrontal cortex reflects the value of time. *J Neurophysiol* 94:2457–2471. [CrossRef Medline](#)
- Roesch MR, Taylor AR, Schoenbaum G (2006) Encoding of time-discounted rewards in orbitofrontal cortex is independent of value representation. *Neuron* 51:509–520. [CrossRef Medline](#)
- Rudebeck PH, Walton ME, Smyth AN, Bannerman DM, Rushworth MF (2006) Separate neural pathways process different decision costs. *Nat Neurosci* 9:1161–1168. [CrossRef Medline](#)
- Rudebeck PH, Behrens TE, Kennerley SW, Baxter MG, Buckley MJ, Walton ME, Rushworth MF (2008) Frontal cortex subregions play distinct roles in choices between actions and stimuli. *J Neurosci* 28:13775–13785. [CrossRef Medline](#)
- Rushworth MF, Walton ME, Kennerley SW, Bannerman DM (2004) Action sets and decisions in the medial frontal cortex. *Trends Cogn Sci* 8:410–417. [CrossRef Medline](#)
- Rushworth MF, Kolling N, Sallet J, Mars RB (2012) Valuation and decision-making in frontal cortex: one or many serial or parallel systems? *Curr Opin Neurobiol* 22:946–955. [CrossRef Medline](#)
- Schapiro AC, Rogers TT, Cordova NI, Turk-Browne NB, Botvinick MM (2013) Neural representations of events arise from temporal community structure. *Nat Neurosci* 16:486–492. [CrossRef Medline](#)
- Shenoy KV, Sahani M, Churchland MM (2013) Cortical control of arm movements: a dynamical systems perspective. *Annu Rev Neurosci* 36:337–359. [CrossRef Medline](#)
- Shidara M, Richmond BJ (2002) Anterior cingulate: single neuronal signals related to degree of reward expectancy. *Science* 296:1709–1711. [CrossRef Medline](#)
- Shima K, Tanji J (1998) Role for cingulate motor area cells in voluntary movement selection based on reward. *Science* 282:1335–1338. [CrossRef Medline](#)
- Small DM, Bender G, Veldhuizen MG, Rudenga K, Nachtigal D, Felsted J (2007) The role of the human orbitofrontal cortex in taste and flavor processing. *Ann N Y Acad Sci* 1121:136–151. [CrossRef Medline](#)
- Takahashi YK, Roesch MR, Wilson RC, Toreson K, O'Donnell P, Niv Y, Schoenbaum G (2011) Expectancy-related changes in firing of dopamine neurons depend on orbitofrontal cortex. *Nat Neurosci* 14:1590–1597. [CrossRef Medline](#)
- Thaler RH (1981) Some empirical evidence on dynamic inconsistency. *Econ Lett* 8:201–207. [CrossRef](#)
- Toda K, Sugase-Miyamoto Y, Mizuhiki T, Inaba K, Richmond BJ, Shidara M (2012) Differential encoding of factors influencing predicted reward value in monkey rostral anterior cingulate cortex. *PLoS One* 7:e30190. [CrossRef Medline](#)
- Van Hoesen GW, Morecraft RJ, Vogt BA (1993) Connections of the monkey cingulate cortex. In: *Neurobiology of cingulate cortex and limbic thalamus: a comprehensive handbook* (Vogt BA, Gabriel M, eds), pp 249–284. Cambridge, MA: Birkhaeuser.
- Wallis JD, Kennerley SW (2010) Heterogeneous reward signals in prefrontal cortex. *Curr Opin Neurobiol* 20:191–198. [CrossRef Medline](#)
- Wallis JD, Rich EL (2011) Challenges of interpreting frontal neurons during value-based decision-making. *Front Neurosci* 5:124. [CrossRef Medline](#)
- Wallis JD, Anderson KC, Miller EK (2001) Single neurons in prefrontal cortex encode abstract rules. *Nature* 411:953–956. [CrossRef Medline](#)
- Walton ME, Bannerman DM, Rushworth MF (2002) The role of rat medial frontal cortex in effort-based decision making. *J Neurosci* 22:10996–11003. [Medline](#)
- Walton ME, Behrens TE, Buckley MJ, Rudebeck PH, Rushworth MF (2010) Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. *Neuron* 65:927–939. [CrossRef Medline](#)
- Wise SP (2008) Forward frontal fields: phylogeny and fundamental function. *Trends Neurosci* 31:599–608. [CrossRef Medline](#)
- Wu KN, Tan BK, Howard JD, Conley DB, Gottfried JA (2012) Olfactory input is critical for sustaining odor quality codes in human orbitofrontal cortex. *Nat Neurosci* 15:1313–1319. [CrossRef Medline](#)