

Choice Coding in Frontal Cortex during Stimulus-Guided or Action-Guided Decision-Making

Chung-Hay Luk and Jonathan D. Wallis

Helen Wills Neuroscience Institute and Department of Psychology, University of California at Berkeley, Berkeley, California 94720-3190

To optimally obtain desirable outcomes, organisms must track outcomes predicted by stimuli in the environment (stimulus–outcome or SO associations) and outcomes predicted by their own actions (action–outcome or AO associations). Anterior cingulate cortex (ACC) and orbitofrontal cortex (OFC) are implicated in tracking outcomes, but anatomical and functional studies suggest a dissociation, with ACC and OFC responsible for encoding AO and SO associations, respectively. To examine whether this dissociation held at the single neuron level, we trained two subjects to perform choice tasks that required using AO or SO associations. OFC and ACC neurons encoded the action that the subject used to indicate its choice, but this encoding was stronger in OFC during the SO task and stronger in ACC during the AO task. These results are consistent with a division of labor between the two areas in terms of using rewards associated with either stimuli or actions to guide decision-making.

Introduction

Orbitofrontal cortex (OFC) is known to make important contributions to decision-making, but its precise role remains unclear. One possibility is that OFC is important for associating stimuli with the rewarding outcomes they predict (stimulus–outcome or SO associations) while the anterior cingulate cortex (ACC) is important for associating actions with rewarding outcomes (action–outcome or AO associations). Although there is neuropsychological evidence to support this dissociation in humans (Camille et al., 2011), monkeys (Rudebeck et al., 2008) and rats (Balleine and Dickinson, 1998; Pickens et al., 2003; Ostlund and Balleine, 2007), there is little support for such a dissociation at the single-neuron level. While OFC neurons in monkeys typically encode the value of predicted outcomes rather than the motor response necessary to obtain the outcome (Tremblay and Schultz, 1999; Wallis and Miller, 2003; Padoa-Schioppa and Assad, 2006; Ichihara-Takeda and Funahashi, 2008; Abe and Lee, 2011), there have been some notable exceptions (Tsujiimoto et al., 2009). Furthermore, robust encoding of actions has been seen in rat OFC (Feierstein et al., 2006; Furuyashiki et al., 2008; Sul et al., 2010; van Wingerden et al., 2010). With regard to ACC, many studies have emphasized the role it plays in predicting the outcome associated with a given action (Ito et al., 2003; Matsumoto et al., 2003; Williams et al., 2004; Luk and Wallis, 2009; Hayden and Platt, 2010), but there have also been studies showing ACC

neurons encoding the rewards predicted by stimuli (Seo and Lee, 2007; Kennerley et al., 2009; Cai and Padoa-Schioppa, 2012).

One problem in interpreting these results is that tasks often do not allow the researcher to unambiguously determine whether the choice was driven by AO or SO associations. Often animals are presented with pairs of reward-predictive stimuli and must pick the stimulus associated with their preferred reward. The assumption is that they access SO associations to recall the reward associated with the stimulus, and then use this information to guide their choice. However, with repeated presentation of these choices, the animal may learn to make a specific response when a specific pair of pictures is presented (a stimulus–response association). Reward-predictive neural activity could then reflect an AO association, indicating knowledge of the reward that is associated with that response. A second problem is that choice behavior requires at least two components: a discriminative cue that indicates the choice options (which can be a stimulus or an action) and a means for the animal to demonstrate its choice by selecting the preferred outcome. These components are not always separated. For example, in a T-maze the action serves as both the discriminative cue (indicating which outcome will be received) and the mechanism by which to select the preferred outcome.

To dissociate these processes, we designed a choice task that enabled us to cue possible outcomes via either AO or SO associations and to separate these processes from the response related to the choice. We predicted that OFC and ACC would encode SO and AO associations respectively, consistent with previous neuropsychological findings.

Materials and Methods

Animal preparation. Two male rhesus monkeys (*Macaca mulatta*), age 6 and 7, weighing 7.5 kg and 12.5 kg, performed the AO and SO tasks in a computerized system. They sat in a primate chair facing a computer monitor. Experimental control and behavioral data were displayed and collected using the NIMH Cortex program. Eye position was recorded via an infrared camera (ISCAN). In each recording session, neurons from

Received Oct. 19, 2012; revised Nov. 20, 2012; accepted Nov. 27, 2012.

Author contributions: C.-H.L. and J.D.W. designed research; C.-H.L. performed research; C.-H.L. and J.D.W. analyzed data; C.-H.L. and J.D.W. wrote the paper.

The project was funded by NIDA Grant R01DA19028 and NINDS Grant P01NS040813 to J.D.W., and NIDA Training Grant F31DA026352 to C.-H.L.

The authors declare no competing financial interests.

Correspondence should be addressed to Jonathan D. Wallis, Helen Wills Neuroscience Institute and Department of Psychology, University of California at Berkeley, 132 Barker Hall, Berkeley, CA 94720-3190. E-mail: wallis@berkeley.edu.

DOI:10.1523/JNEUROSCI.4920-12.2013

Copyright © 2013 the authors 0270-6474/13/331864-08\$15.00/0

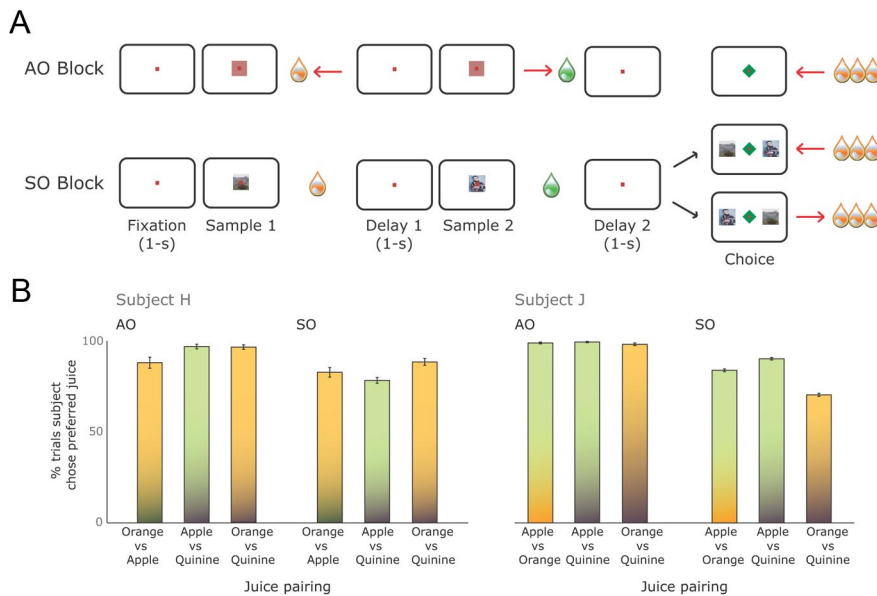


Figure 1. The behavioral task and performance. **A**, The experiment consisted of two tasks presented in alternating blocks. In the AO task, the subject moved the lever (indicated by red arrows) in two directions to receive two different liquid rewards. He then repeated one of the movements to receive more of that corresponding reward. For instance, the subject could move left to receive orange juice, then right to receive apple juice. If the subject prefers orange juice, he will repeat the leftward movement in the choice period to gain more orange juice. In the SO block, the subject viewed two pictures that were paired with different rewards. Then he selected the picture associated with his preferred reward via a lever movement. Pictures were equally likely to appear on either side of the screen at the choice phase. **B**, Mean (\pm SEM) proportion of trials in which the subject selected his more preferred juice in each juice pairing.

Berkeley Animal Care and Use Committee and met the National Institutes of Health guidelines.

Behavioral task. The experiment comprised two tasks: an AO and an SO task (Fig. 1A). In both tasks, each trial consisted of two distinct phases: a sampling phase and a choice phase. During the sampling phase, the subject experienced two events that were each predictive of a specific outcome. In the SO task, the subject was presented with two pictures sequentially and each picture was followed by one of three juices. In the AO task, the subject performed two actions sequentially and each action was followed by one of three juices. The relationship between the predictive events and the juices was randomly varied from trial to trial. The sampling phase was followed by a choice phase in which the subject made a choice guided by the predictive events of the sampling phase. To choose optimally the subject had to remember which juice was paired with which predictive event during the sampling phase of the trial.

Each trial began with a 1 s fixation period, in which the subject fixated within $\pm 2.5^\circ$ of a fixation spot. In an SO trial, during the sampling phase, the subject viewed two different pictures that yielded two different rewards. The pictures were $5^\circ \times 5^\circ$ isoluminant, natural scene images and they were presented for 0.6 s. Reward was presented such that its offset coincided with the offset of the picture. Then in the choice phase, the subject saw both pictures side by side and moved a lever in the direction of the picture associated with the more preferred reward. The left/right position of the pictures at the choice period was counterbalanced and randomized across trials. The SO contingencies between the picture and reward changed trial to trial. All SO pairings were equally likely to occur. Thus, in order for the subject to obtain more of his preferred reward at the choice phase, he needed to remember which picture was linked with which outcome during the sampling phase of that particular trial. The AO trial was similar, but instead of viewing pictures, the subject had to execute two different lever movements and then repeated one of them in the choice phase. During the sampling phase, reward was delivered immediately following the detection of a lever movement. For both tasks, fixation was required throughout the trial, except during delivery of reward during the sampling phase and following onset of the “choice” cue. The intertrial interval was 3 s.

Rewards were apple juice, orange juice, and quinine. Reward size and concentration were tailored to each subject so as to ensure that they showed consistent preferences and so that they received their daily fluid aliquot within a single recording session. Sample juices lasted ~ 0.4 s, which resulted in the delivery of 0.25 ml of juice. Choice rewards lasted ~ 1250 s, which resulted in the delivery of 0.77 ml of juice.

The two tasks were organized into blocks: 35 trial-long blocks of AO trials and 45 trial-long blocks of SO trials. Because behavioral performance was lower on the SO task, this ensured that we had a similar number of correct trials for each task. We varied the order-

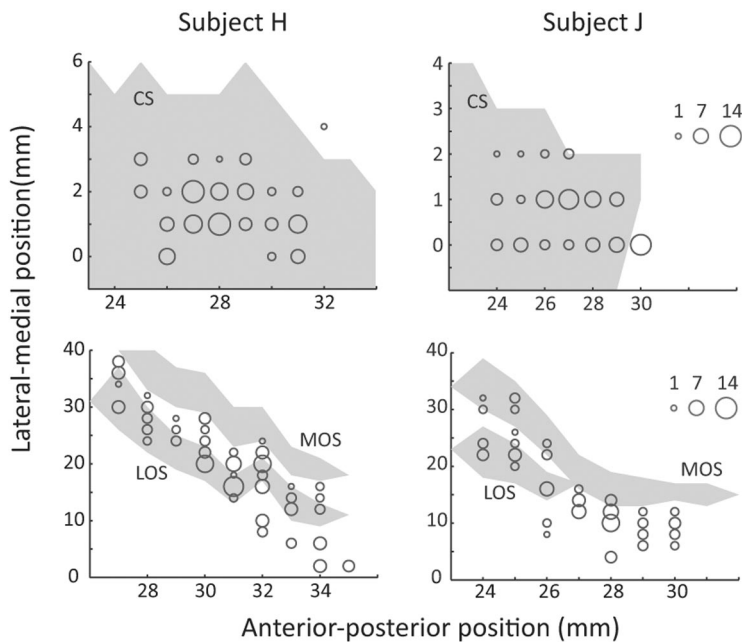


Figure 2. Number of neurons recorded at each location in ACC (top) and OFC (bottom). The anterior–posterior position is measured from the interaural line. In subject H, the genu of the corpus callosum was at AP 24 mm and in subject J it was at AP 23 mm. In the ACC plot, the lateral–medial position extended from the fundus of the cingulate sulcus (0 mm) to more medial positions within the dorsal bank of the cingulate sulcus. In the OFC plot, the lateral–medial position extended from the ventral bank of the principal sulcus (0 mm), around the inferior convexity, and onto the orbitofrontal surface. The extent of sulci is shown by the gray shading. CS, Cingulate sulcus; MOS, medial orbital sulcus; LOS, lateral orbital sulcus.

both OFC and ACC were randomly sampled and their activity was recorded (Plexon Instruments). Methods for recording neuronal activity have been previously reported in detail (Lara et al., 2009). All procedures used were approved by the University of California at

ing of the predictive events (pictures in the SO task and movements in the AO task) across blocks, but within a block, the order of the predictors did not change. Consequently, there were two different blocks of the AO task and two different blocks of the SO task. In the AO task, the subject had to initially determine the AO block type through trial and error on the first trial, but would then know the order for the remaining trials. For example, if the first trial of an AO block required a leftward movement followed by a rightward movement, then that ordering would remain consistent for the rest of the block. During a session, blocks alternated from being one of the AO blocks to being one of the SO blocks, and across sessions we alternated between beginning with either an AO or an SO block. Other than those constraints, the ordering of the blocks was random.

Analysis of neuronal data during the sampling phase. We excluded error trials from our statistical analyses. Subjects made two types of errors: premature breaks of eye fixation and movements in the wrong direction during the AO task. Errors resulted in 5 s timeout, after which the trial resumed from the point where the subject was before the error. If the error occurred before the delivery of the first sample reward, then the trial was included in our statistical analyses, but all other error trials were excluded. In practice, error trials only counted for a small proportion of the trials within a recording session (subject H, $3 \pm 0.4\%$; subject J, $15 \pm 0.8\%$).

For each task, we visualized spike density histograms by averaging activity across the appropriate conditions using a sliding window of 150 ms. To analyze encoding of the first reward, we focused on neuronal selectivity that began at reward onset and ended at the end of the first delay period. This corresponded to the period of time in which the subject needed to encode information about which reward was presented and with which predictive event it was associated (which picture in the SO task and which action in the AO task). For each neuron, we separated its data by task. In each task, we performed a “sliding” analysis. We took a 200 ms window of time, beginning at reward onset, and performed a two-way ANOVA on the neuron’s mean firing rate during that window with factors of Predictor (which action was made in the AO task or which picture was shown in the SO task) and Outcome (which reward was given). We then advanced the window by 10 ms and analyzed the next 200 ms window of time, and continued in this fashion until the end of the first delay period. We defined each neuron’s selectivity according to which factors reached significance at any point during this period. We controlled for multiple comparisons in this, and all subsequent sliding analyses, by calculating a false discovery rate using baseline neuronal data taken from the first second of the intertrial interval. We did this by generating 1000 artificial datasets by randomly shuffling the relationship between trial number and experimental condition for each neuron. For each neuron and each artificial dataset, we then performed the sliding analysis and determined a threshold which fewer than 5% of neurons reached.

We used this same sliding analysis to determine the neuron’s maximal selectivity and to examine the time course of neuronal encoding. For each neuron and each time point we calculated the percentage of variance in the neuron’s firing rate that each experimental factor explained (percentage explained variance or PEV). We defined the neuron’s maximal selectivity as the largest PEV value. We defined the earliest latency of selectivity for an experimental factor as the first time bin that reached our

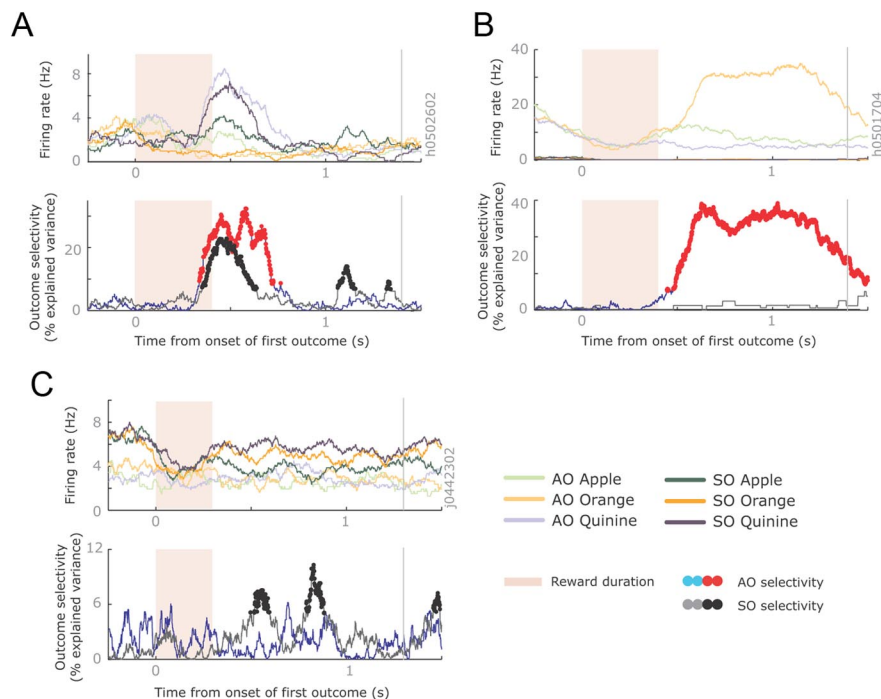


Figure 3. Example neurons with different outcome selectivity. **A**, Spike density histogram illustrating an OFC neuron encoding the first reward across both tasks. It had its highest firing rate when the reward was quinine (blue lines) and its lowest firing rate when the reward was orange juice (orange lines). This neuron was recorded from subject H, and so the ordering was the inverse of this subject’s preferences. The magnitude of outcome selectivity, defined as the percentage of variance in the neuron’s firing rate attributable to the outcome, is shown below the histogram. Significant encoding is denoted with red and black dots in the AO and SO tasks, respectively. The gray vertical line illustrates the end of the delay period. **B**, An ACC neuron that encoded outcomes only in the AO task. (The activity on the SO task was so low that it is barely visible on this plot.) This neuron was recorded from subject H and the ordering of the juices reflected the subjects’ preferences. **C**, An OFC neuron that encoded outcomes only in the SO task. This neuron was recorded from subject J and was the inverse of the subject’s preferences.

Table 1. Percentage of recorded neurons with significant encoding of the predictive event or the outcome it predicted at any point from the delivery of the first outcome until the end of the first delay

Selectivity	Area	Predictor	Outcome	Predictor × Outcome
AO only	OFC	12	17	6
	ACC	12	11	11
SO only	OFC	6	18	8
	ACC	6	19	5
Both	OFC	1	14	1
	ACC	2	17	0

Neurons are classified according to whether they showed a significant main effect (with no significant interactions) or a significant interaction. Neurons are further divided according to whether they showed the selectivity in only one task or both tasks. None of the proportions listed in the table differed between the areas (χ^2 tests, $p > 0.05$ in all cases). Bold indicates that the percentage of selective neurons was significantly higher than that expected by chance (binomial test, $p < 0.05$).

threshold for selectivity. We contrasted the strength and latency of neuronal encoding across brain areas and tasks using ANOVA.

Analysis of neuronal data during the choice phase. To characterize neuronal selectivity at the choice phase of the task, we performed two sliding ANOVAs for each neuron in each task. For AO trials, we used a two-way ANOVA with the dependent variable of mean firing rate and factors of chosen action and chosen reward. For SO trials, we used a three-way ANOVA with the additional factor of stimulus, the spatial ordering of the pictures (i.e., which picture was on the left and which was on the right). In both analyses, we removed trials in which quinine was picked since this occurred too infrequently to meaningfully analyze the data. Each analysis was performed from the onset of the second reward until the time of the chosen action. We determined this for each recording session by calculating the subject’s median reaction time for the AO and SO task independently. Bins were 200 ms long and advanced by 10 ms. We de-

fined each neuron's selectivity according to which factors reached significance at any point during this period.

Finally, we determined which neurons encoded the chosen outcome following the delivery of the final reward with another sliding analysis. For each neuron, we performed a one-way ANOVA on mean firing rate with the factor of chosen outcome. The analysis window began 200 ms after the onset of the reward and ended at its offset. Again, bins were 200 ms and advanced in 10 ms increments.

Results

Behavioral preferences across tasks

Two rhesus monkeys (H and J) performed sequential choice tasks using AO or SO associations (Fig. 1A). In the AO task, during the sample phase, the subject made two arm movements, each of which was followed by a different drop of juice. During the choice phase, the subject then repeated one of the movements to receive a larger amount of the juice that was associated with that movement during the sample phase. Which juice was associated with which movement was randomly determined on each trial. Therefore, to receive his preferred juice at the choice phase, the subject needed to remember which juice was paired with which movement during the sample phase. The SO task was analogous, but instead of making arm movements, the subject saw two pictures appear sequentially, each of which was followed by a small drop of juice. At the choice phase, the two pictures reappeared and the subject selected one, to receive the larger amount of juice.

Both subjects showed clear preferences on both tasks, picking their preferred juice on 88% of trials (Fig. 1B). Subject H preferred orange juice to apple juice and quinine, and apple juice to quinine. Subject J preferred apple juice to both orange juice and quinine, and orange juice to quinine. For both subjects, their preferences were consistent across both tasks. We quantified their ability to perform the tasks by determining the number of choices that were consistent with their preferences. Subject H's performance on the AO task (mean consistent choices = $94 \pm 1.6\%$) was significantly better than his performance on the SO task (mean = $83 \pm 1.7\%$; one-way ANOVA, ($F_{(1,160)} = 21, p < 1 \times 10^{-5}$). Subject J's performance was also better on the AO task (mean = $99 \pm 0.5\%$) than the SO task (mean = $82 \pm 1.3\%$; one-way ANOVA, ($F_{(1,118)} = 158, p < 1 \times 10^{-15}$).

Subjects' reaction times showed clear differences between the two tasks. Subject H took 199 ms (median) to respond in the AO task, which was much faster than the 744 ms in the SO task (Wilcoxon rank-sum test, $p < 1 \times 10^{-15}$). Subject J had a similar pattern: 266 ms in the AO task and 709 ms in the SO task (Wilcoxon rank-sum test, $p < 1 \times 10^{-15}$). These differences likely related to differences in the cognitive processes occurring at the choice phase in each task. In the AO task, no decision necessarily occurred, since the subject could have planned his choice action already, indeed from the moment he received the second juice. But in the SO task, the subject could not plan ahead. He needed to observe both pictures presented side by side, select the one that was associated with the more preferred option, and only then could he plan his action.

Encoding of AO and SO associations during the sampling phase

We recorded the activity of 215 ACC neurons (H, 125; J, 90) and 249 OFC neurons (H, 145; J, 104). Our ACC recordings were from the dorsal bank of the cingulate cortex. There is considerable disagreement regarding the correct cytoarchitectonic designation of this area, with it labeled as area 9 (Vogt et al., 2005), 32 (Petrides and Pandya, 1994), 9/32 (Paxinos et al., 2000), or 24b

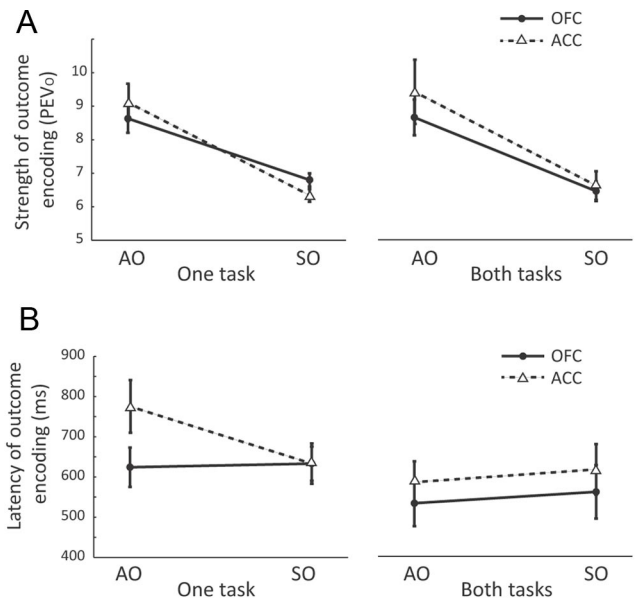


Figure 4. Strength and latency of outcome encoding in the AO and SO tasks. **A**, To quantify the strength of outcome encoding, for each neuron with a significant main effect of Outcome, we measured the maximum percentage of variance in its firing rate attributable to this factor (PEV₀). To compare the strength of this encoding across our neuronal populations, we then ran a three-way ANOVA on the PEV₀ values with factors of brain area, task, and whether the cell was selective in one or both tasks. We found that outcome selectivity was stronger in the AO task than the SO task ($F_{(1,247)} = 53, p < 1 \times 10^{-11}$), but there were no other significant main effects or interactions. In particular, there was no evidence that outcome selectivity was stronger in ACC during the AO task and OFC during the SO task (Task \times Area interaction, $F_{(1,247)} = 1.4, p > 0.1$). **B**, For each outcome-selective neuron, we used the sliding analysis to determine the latency with which it first encoded outcome information. We then performed a three-way ANOVA on the neuronal latencies with factors of brain area, task, and whether neurons were selective in one or both tasks. Outcome selectivity was encoded significantly faster by neurons that were selective in both tasks (mean = 576 ± 30 ms) compared with those selective in only one task (mean = 651 ± 25 ms, $F_{(1,247)} = 5.3, p < 0.05$). There were no other significant main effects or interactions. In particular, there was no evidence that there was any difference between the brain areas in terms of the latency to encode outcome information on the AO or SO task (Task \times Area interaction, $F_{(1,247)} < 1, p > 0.1$).

(Carmichael and Price, 1994). It most likely represents a transition zone between the cingulate and prefrontal cortices. It has also been the predominate focus of previous neurophysiological studies of ACC (Seo and Lee, 2007; Luk and Wallis, 2009; Hayden and Platt, 2010; Kennerley et al., 2011; Cai and Padoa-Schioppa, 2012). Our OFC recordings were from areas 11, 12, and 13. Figure 2 illustrates the precise recording locations.

Both tasks required the subject to remember the first juice outcome across the first delay. Many neurons encoded this information across the first delay, and showed similar outcome encoding across the two tasks. Figure 3A illustrates an OFC neuron that encoded quinine across the delay of both the SO and AO task. However, other outcome-selective neurons showed very different patterns of selectivity between the two tasks. Figure 3B illustrates an ACC neuron that encoded the first outcome during the AO task but did not fire at all during the SO task. In contrast, Figure 3C illustrates an OFC neuron that encoded the outcome during the SO task but not during the AO task.

We quantified the prevalence of neurons that encoded the first juice outcome, the first predictive event, or a specific combination of juice and predictor, using a sliding two-way ANOVA (see Materials and Methods). There were two potential patterns of result that would be consistent with our original hypothesis. First, we might expect more neurons to encode specific SO or AO

associations in OFC and ACC, respectively. Such neurons would show a significant Predictor \times Outcome interaction in the appropriate task. There was some evidence to suggest that this was the case (Table 1). The proportion of neurons encoding Predictor \times Outcome interactions in ACC was significantly above chance during the AO task, but not the SO task, while the opposite was true in OFC. However, there was only a small proportion of neurons showing significant Predictor \times Outcome interactions in either ACC or OFC, and so statistical tests aimed at determining whether the proportions were significantly different between the two brain areas did not reach significance (χ^2 tests, $p > 0.1$ in both cases).

In contrast to the small number of neurons encoding Predictor \times Outcome interactions, a larger proportion of neurons encoded just the Outcome (Table 1). A second pattern of encoding that would also have been consistent with our original hypothesis was if such outcome-selective neurons were more prevalent in OFC during the SO task and more prevalent in ACC during the AO task. However, there was no evidence that this was the case (Table 1). We also examined whether there was any evidence that Outcome encoding was stronger and/or earlier in OFC compared with ACC for the SO task and vice versa for the AO task, but these analyses did not show any differences between the areas (Fig. 4).

In summary, there was evidence to support our original hypothesis, but it was rather weak. Although some ACC neurons encoded specific AO associations and some OFC neurons encoded specific SO associations, they were a relatively small proportion of the overall neuronal population. Instead, neurons in both areas tended to encode the outcome independent of its association with the predictive event, and this encoding was similar in both areas for both tasks.

Encoding of actions during the choice phase

During the choice phase, many neurons encoded the chosen action. Figure 5 illustrates two examples of neurons that modulated their activity depending on the chosen action. The ACC neuron in Figure 5A showed increased activity during the AO task when the subject intended to choose the rightward response compared with the leftward response, but did not encode the expected reward associated with the chosen action. There was no evidence that the neuron encoded the chosen action in the SO task. In contrast, the OFC neuron in Figure 5B had a higher firing rate when the subject chose to make a rightward movement, but only in the SO task. It did not encode the spatial position of the pictures or which reward the subject expected.

To quantify this selectivity, we performed two sliding ANOVAs. For the AO task, we used a two-way ANOVA with the dependent variable of mean firing rate and factors of chosen action and chosen reward. For the SO task, we used a three-way ANOVA adding the additional factor of the position of the pictures on the screen at the choice phase. We focused on those neurons that only

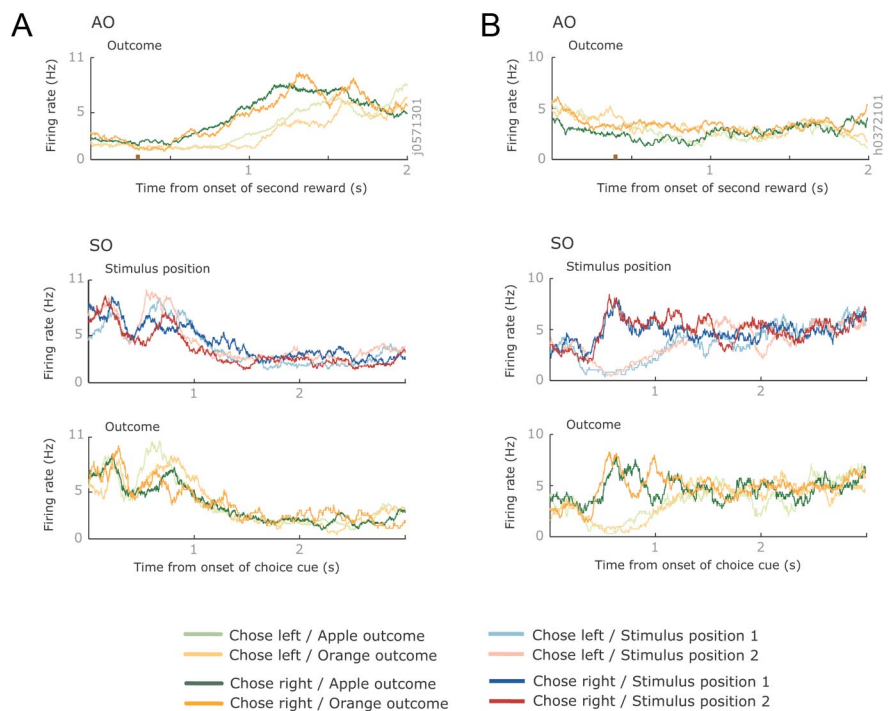


Figure 5. Example neurons encoding the chosen action in only one task. Spike density histograms are sorted according to the action chosen by the subject as well as the reward they expected for both the AO and SO tasks. In addition, for the SO task, spike density histograms are sorted according to the action chosen by the subject as well as the left/right position of the pictures on the screen. Neural activity is shown from the time of onset of the second reward for the AO task (the brown spot on the x-axis indicates time of reward offset) and from the time of the onset of the choice cue for the SO task. **A**, An ACC neuron that encoded the chosen action in the AO task, but not the SO task. **B**, An OFC neuron that encoded the chosen action only in the SO task.

encoded the chosen action, i.e., showed a significant main effect of Action with no significant interactions. Encoding of the chosen action began shortly after the delivery of the second juice in the AO task, but not until the presentation of the choice cues in the SO task (Fig. 6). This is consistent with the reaction time data suggesting that the action is selected earlier in the trial in the AO task relative to the SO task. In addition, we observed a clear double dissociation. Encoding of the chosen action was significantly more prevalent in ACC (34 of 215 or 16%) in the AO task compared with OFC (24 of 249 or 10%, χ^2 test = 4.0, $p < 0.05$). In contrast, in the SO task, encoding of the chosen action was significantly more prevalent in OFC (50 of 249 or 20%) compared with ACC (22 of 215 or 10%, χ^2 test = 8.5, $p < 0.005$). These action-encoding neurons were not simply motor neurons, since only a minority (OFC, 9 of 249 or 4%; ACC, 5 of 215 or 2%) encoded the action in both tasks.

We also quantified the degree to which neurons encoded other task parameters during the choice phase (Table 2). In the SO task, many neurons in both OFC and ACC encoded the action in relation to the position of the stimuli on the screen. In other words, they would encode a specific action but only when it was directed toward a specific picture i.e., a significant Stimulus \times Action interaction. There was little encoding of other task parameters.

Neurons encode the final juice reward in a task-dependent manner

Once the choice was made, the subject received the reward that he had selected. Many neurons responded to the final delivery of this reward in a way that depended on the task. Figure 7A shows an OFC neuron that showed a higher firing rate when orange was

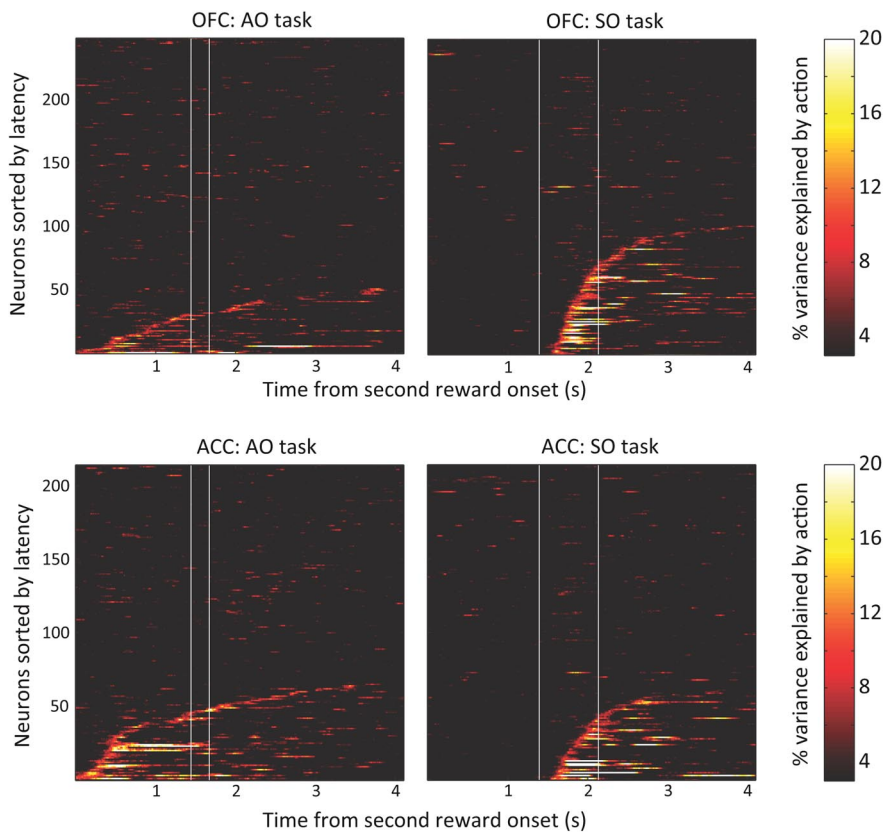


Figure 6. Time course of encoding chosen action across the neuronal population. Each plot shows the percentage of variance in each neuron’s firing rate that can be explained by the chosen action for the AO task and the SO task. Neural selectivity is shown from the onset of the second reward. The first vertical line indicates the median time of the choice cue onset and the second vertical line indicates the median time of the chosen action. Each horizontal line is the selectivity from a single neuron, and they have been sorted on the basis of the latency of that selectivity. In the AO task, the chosen action begins to be encoded from the time of the onset of the second reward, and it is encoded by a significantly larger population of neurons in ACC than OFC. In the SO task, the chosen action cannot be encoded until the spatial position of the pictures is revealed. However, once this occurs, the information is encoded by a significantly larger population of neurons in OFC than ACC.

Table 2. Percentage of recorded neurons with significant encoding of Stimulus (left/right position of the pictures), Action (the animal’s choice: left or right) or Outcome (the type of juice associated with the choice) during the delivery of the second reward up until the time of the chosen action (based on the subject’s median reaction time for each task)

Selectivity	Area	AO task	SO task
Stimulus	OFC	—	1
	ACC	—	0
Action	OFC	10*	20*
	ACC	16*	10*
Outcome	OFC	9	7
	ACC	10	6
S × A	OFC	—	12
	ACC	—	18
S × O	OFC	—	2
	ACC	—	4
A × O	OFC	3	5
	ACC	5	7
S × A × O	OFC	—	8
	ACC	—	10

Neurons are classified according to whether they showed a significant main effect (with no significant interactions) or a significant interaction. Note that for the AO task, there were no pictures on the screen, and consequently the Stimulus parameter was not manipulated. Bold indicates that the percentage of selective neurons was significantly higher than that expected by chance (binomial test, $p < 0.05$). Asterisk indicates that the proportion of selective neurons was significantly different between the two areas (χ^2 tests, $p < 0.05$).

selected relative to apple, but only in the AO task. Figure 7B shows an ACC neuron that distinguished between the delivery of orange juice and apple juice, but only in the SO task. To quantify these effects, we performed a sliding one-way ANOVA with outcome (orange or apple juice) as the independent variable. Although over half of the neurons encoded outcomes (OFC, 145 of 249 or 58%; ACC, 107 of 215 or 50%), the majority did so in just one task (Table 3).

Discussion

Our original hypothesis predicted that OFC neurons would be responsible for encoding SO associations and ACC would be responsible for encoding AO associations. This was based on neuropsychological studies in rats (Balleine and Dickinson, 1998; Pickens et al., 2003; Ostlund and Balleine, 2007), monkeys (Rudebeck et al., 2008) and humans (Camille et al., 2011), that have demonstrated a dissociation in the role of ACC and OFC in using AO and SO associations to guide behavior. Our results are consistent with this division of labor, but they also paint a more nuanced picture for the specific roles of the two areas. We only saw a small population of neurons encoding the specific AO or SO associations, although the relative proportion of these neurons in OFC and ACC was consistent with our original hypothesis. In contrast, we saw robust encoding of the action necessary to make the choice. Furthermore, encoding of this action was more prevalent in OFC in the SO task and more prevalent in ACC in the AO task. Our results

suggest that the function of the two areas is not so much the encoding of specific AO or SO associations per se, but rather using those associations to guide choice.

A prominent theory of OFC function argues that decision-making occurs by comparing SO associations: choices are held to be made in a “goods” space (Padoa-Schioppa, 2011). Part of the motivation for this claim is the relative paucity of action encoding in OFC. An alternative theory states that choices are made downstream of OFC in an action space by integrating the value of the potential outcome (e.g., a candy bar) with the value of the action necessary to acquire the outcome (e.g., the effort necessary to go to the store and buy the candy bar) (Rangel and Hare, 2010). Our results suggest a more complex picture. Both OFC and ACC neurons are capable of encoding the chosen action, but OFC is more likely to do so when the choice is guided by environmental stimuli, whereas ACC is more likely to do so when the choice is guided by the subject’s own actions. These results are consistent with recent neuroimaging studies that have also found that there is a good deal of flexibility in the neuronal mechanisms underlying choice, with decisions being made in either goods or action space depending on the demands of the task (Wunderlich et al., 2010). Furthermore, decisions in action space activate ACC (Wunderlich et al., 2009), whereas decisions in goods space activate ventromedial prefrontal cortex (Wunderlich et

al., 2010), an area in which a BOLD response is frequently observed during tasks that involve OFC (Wallis, 2012). These conclusions are remarkably similar to our own.

There are a couple of reasons why we may not have seen a large number of neurons encoding specific AO or SO associations. One possibility is that the actual storage of these associations takes place in other brain areas independent of OFC and ACC. With regard to SO associations, an obvious candidate is the amygdala, which has long been held to be important for this process (Baxter and Murray, 2002; Everitt et al., 2003; Balleine and O'Doherty, 2010). AO associations might be stored in downstream motor areas, such as the cingulate motor area (Shima and Tanji, 1998), or in the striatum (Lau and Glimcher, 2007; Kim et al., 2009; Stalnaker et al., 2010). It is also possible that a different task design might have been better able to detect encoding of AO and SO associations. Our design focused on how such associations are held in working memory, as several theories have speculated that frontal cortex may be particularly important for “online” decision-making (Wallis, 2007; Zald, 2007; Padoa-Schioppa, 2011; Hunt et al., 2012). It is possible that we might have observed a dissociation between OFC and ACC in the encoding of SO and AO associations had we focused on long-term associative encoding, a possibility that is worth exploring in future experiments.

The relative paucity of encoding relating to the AO or SO associations was in marked contrast to the robust neuronal encoding of the action responsible for the subject's ultimate choice. The most surprising aspect of this result was that OFC neurons encoded the action necessary to choose the final reward on the SO task. Many previous studies that examined OFC neuronal responses during decision-making in monkeys have concluded that OFC neurons do not encode actions (Tremblay and Schultz, 1999; Padoa-Schioppa and Assad, 2006; Ichiara-Takeda and Funahashi, 2008; Kennerley and Wallis, 2009; Abe and Lee, 2011), although there have been exceptions to this consensus (Tsujimoto et al., 2009). This raises the question as to why these other choice tasks did not see action encoding in OFC, particularly given that they used stimuli to indicate available outcomes. One possibility is that tasks that could use Pavlovian approach responses rather than goal-directed actions may have less encoding of the choice action. There are a couple of reasons that may have biased previous choice tasks toward being solved through Pavlovian mechanisms. Most tasks have used a consistent mapping between the stimulus and its outcome, focusing on long-term storage processes. Thus, a specific stimulus can acquire a specific incentive value and attract Pavlovian approach responses. In contrast, our SO task could not be performed in this way. Both pictures could predict any of the three rewards depending on the trial. Consequently, neither picture could be consistently associated with a particular outcome, and therefore could not acquire a specific incentive value. A recent study in humans has also found that frontal decision-making mechanisms become less involved with repeated presentations of a choice, consistent with more low-level selection mechanisms taking over as choices become practiced (Hunt et al., 2012).

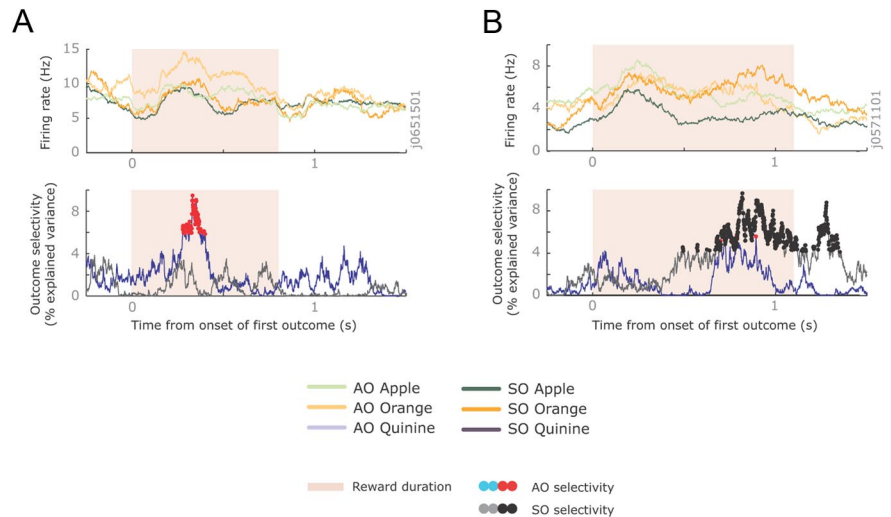


Figure 7. Example neurons selective for outcomes after the decision is executed. **A**, An OFC neuron that encoded chosen outcomes only in the AO task. It had elevated firing when the subject selected orange juice (light orange line). The magnitude of outcome selectivity is shown below the histogram. Both plots follow the same format as Figure 2. **B**, An ACC neuron that encoded chosen outcomes only in the SO task.

Table 3. Percentage of recorded neurons with significant encoding of outcome during juice delivery

AO only	OFC	26
	ACC	20
SO only	OFC	19
	ACC	22
Both	OFC	13
	ACC	7

Neurons are divided according to whether they showed the selectivity in only one task or both tasks. None of the proportions listed in the table differed between the areas (χ^2 tests, $p > 0.05$ in all cases). Bold indicates that the percentage of selective neurons was significantly higher than that expected by chance (binomial test, $p < 0.05$).

A second factor that could influence whether choices are made through Pavlovian responses is the way in which the subject informs the experimenter of his choice. Some tasks require the subject to simply look at the chosen stimulus. Eye movements could be a reflexive orientation response, much like a Pavlovian approach response. In contrast, other studies (including our SO task) require the subject to make a lever movement in the direction of the stimulus, a more arbitrary learned mapping between what is on the screen and the physical movement of the lever. In a recent study, we had one subject choose between pictures associated with different outcomes using eye movements and another subject using a lever movement (Kennerley et al., 2009). The subject who used the lever had many more action-encoding neurons in OFC than the subject who made his choice using eye movements. Thus, the more arbitrary mapping between the pictures and the choice action may have prevented Pavlovian responses being used to solve the task, thereby necessitating action encoding in OFC. Notably, the previous study to observe action encoding in monkey OFC also required choices to be made based on abstract stimulus–response mappings (Tsujimoto et al., 2009).

Finally, our findings caution against trying to understand the functions of ACC and OFC solely in terms of calculating scalar value signals, such as those observed in neuromodulatory systems. For example, prominent theories of OFC function argue that it is responsible for calculating the value of goods in our environment along an abstract value scale (Rangel and Hare, 2010; Padoa-Schioppa, 2011). Yet our results show that neurons

in both OFC and ACC can respond to the same good (juice, in this case) in very different ways depending on the task in which the subject is engaged. These findings are similar to recent studies in rats (Takahashi et al., 2011), which have suggested that OFC is responsible for encoding a “state representation,” whereby information about the context in which the organism finds itself can be used to derive more accurate value signals. Thus, similar to other frontal regions (Miller and Cohen, 2001), the response of neurons in ACC and OFC is likely to be a complex interplay between external stimuli, internal motivations, and the context in which they occur.

References

- Abe H, Lee D (2011) Distributed coding of actual and hypothetical outcomes in the orbital and dorsolateral prefrontal cortex. *Neuron* 70:731–741. [CrossRef Medline](#)
- Balleine BW, Dickinson A (1998) Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology* 37:407–419. [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](#)
- Baxter MG, Murray EA (2002) The amygdala and reward. *Nat Rev Neurosci* 3:563–573. [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 (1994) Architectonic subdivision of the orbital and medial prefrontal cortex in the macaque monkey. *J Comp Neurol* 346:366–402. [CrossRef Medline](#)
- Everitt BJ, Cardinal RN, Parkinson JA, Robbins TW (2003) Appetitive behavior: impact of amygdala-dependent mechanisms of emotional learning. *Ann N Y Acad Sci* 985:233–250.
- Feierstein CE, Quirk MC, Uchida N, Sosulski DL, Mainen ZF (2006) Representation of spatial goals in rat orbitofrontal cortex. *Neuron* 51:495–507. [CrossRef Medline](#)
- Furuyashiki T, Holland PC, Gallagher M (2008) Rat orbitofrontal cortex separately encodes response and outcome information during performance of goal-directed behavior. *J Neurosci* 28:5127–5138. [CrossRef Medline](#)
- Hayden BY, Platt ML (2010) Neurons in anterior cingulate cortex multiplex information about reward and action. *J Neurosci* 30:3339–3346. [CrossRef Medline](#)
- Hunt LT, Kolling N, Soltani A, Woolrich MW, Rushworth MF, Behrens TE (2012) Mechanisms underlying cortical activity during value-guided choice. *Nat Neurosci* 15:470–476. [CrossRef Medline](#)
- Ichihara-Takeda S, Funahashi S (2008) Activity of primate orbitofrontal and dorsolateral prefrontal neurons: effect of reward schedule on task-related activity. *J Cogn Neurosci* 20:563–579. [CrossRef Medline](#)
- Ito S, Stuphorn V, Brown JW, Schall JD (2003) Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science* 302:120–122. [CrossRef Medline](#)
- Kennerley SW, Wallis JD (2009) Encoding of reward and space during a working memory task in the orbitofrontal cortex and anterior cingulate sulcus. *J Neurophysiol* 102:3352–3364. [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 H, Sul JH, Huh N, Lee D, Jung MW (2009) Role of striatum in updating values of chosen actions. *J Neurosci* 29:14701–14712. [CrossRef Medline](#)
- Lara AH, Kennerley SW, Wallis JD (2009) Encoding of gustatory working memory by orbitofrontal neurons. *J Neurosci* 29:765–774. [CrossRef Medline](#)
- Lau B, Glimcher PW (2007) Action and outcome encoding in the primate caudate nucleus. *J Neurosci* 27:14502–14514. [CrossRef Medline](#)
- Luk CH, Wallis JD (2009) Dynamic encoding of responses and outcomes by neurons in medial prefrontal cortex. *J Neurosci* 29:7526–7539. [CrossRef Medline](#)
- Matsumoto K, Suzuki W, Tanaka K (2003) Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science* 301:229–232. [CrossRef Medline](#)
- Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24:167–202. [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 (2011) Neurobiology of economic choice: a good-based model. *Ann 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](#)
- Paxinos G, Huang X, Toga AW (2000) The rhesus monkey brain in stereotaxic coordinates. San Diego: Academic.
- Petrides M, Pandya DN (1994) Comparative architectonic analysis of the human and macaque frontal cortex. In: *Handbook of neuropsychology* (Boller F, Grafman J, eds), pp 17–57. New York: Elsevier.
- Pickens CL, Saddoris MP, Setlow B, Gallagher M, Holland PC, Schoenbaum G (2003) Different roles for orbitofrontal cortex and basolateral amygdala in a reinforcer devaluation task. *J Neurosci* 23:11078–11084. [CrossRef Medline](#)
- Rangel A, Hare T (2010) Neural computations associated with goal-directed choice. *Curr Opin Neurobiol* 20:262–270. [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](#)
- Seo H, Lee D (2007) Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. *J Neurosci* 27:8366–8377. [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](#)
- Stalnaker TA, Calhoun GG, Ogawa M, Roesch MR, Schoenbaum G (2010) Neural correlates of stimulus-response and response-outcome associations in dorsolateral versus dorsomedial striatum. *Front Integr Neurosci* 4:12. [CrossRef Medline](#)
- Sul JH, Kim H, Huh N, Lee D, Jung MW (2010) Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making. *Neuron* 66:449–460. [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](#)
- Tremblay L, Schultz W (1999) Relative reward preference in primate orbitofrontal cortex. *Nature* 398:704–708. [CrossRef Medline](#)
- Tsujimoto S, Genovesio A, Wise SP (2009) Monkey orbitofrontal cortex encodes response choices near feedback time. *J Neurosci* 29:2569–2574. [CrossRef Medline](#)
- van Wingerden M, Vinck M, Lankelma JV, Pennartz CM (2010) Learning-associated gamma-band phase-locking of action-outcome selective neurons in orbitofrontal cortex. *J Neurosci* 30:10025–10038. [CrossRef Medline](#)
- Vogt BA, Vogt L, Farber NB, Bush G (2005) Architecture and neurocytology of monkey cingulate gyrus. *J Comp Neurol* 485:218–239. [CrossRef Medline](#)
- Wallis JD (2007) Orbitofrontal cortex and its contribution to decision-making. *Annu Rev Neurosci* 30:31–56. [CrossRef Medline](#)
- Wallis JD (2012) Cross-species studies of orbitofrontal cortex and value-based decision-making. *Nat Neurosci* 15:13–19. [CrossRef Medline](#)
- Wallis JD, Miller EK (2003) Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *Eur J Neurosci* 18:2069–2081. [CrossRef Medline](#)
- 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* 7:1370–1375. [CrossRef Medline](#)
- Wunderlich K, Rangel A, O’Doherty JP (2009) Neural computations underlying action-based decision making in the human brain. *Proc Natl Acad Sci U S A* 106:17199–17204. [CrossRef Medline](#)

Wunderlich K, Rangel A, O'Doherty JP (2010) Economic choices can be made using only stimulus values. *Proc Natl Acad Sci U S A* 107:15005–15010. [CrossRef](#) [Medline](#)

Zald DH (2007) Orbital versus dorsolateral prefrontal cortex: anatomical insights into content versus process differentiation models of the prefrontal cortex. *Ann N Y Acad Sci* 1121:395–406. [CrossRef](#)