

Increased Firing to Cues That Predict Low-Value Reward in the Medial Orbitofrontal Cortex

Amanda C. Burton^{1,2}, Vadim Kashtelyan¹, Daniel W. Bryden^{1,2} and Matthew R. Roesch^{1,2}

¹Department of Psychology, ²Program in Neuroscience and Cognitive Science, University of Maryland, College Park, MD 20742, USA

Address correspondence to Matthew R. Roesch. Email: mroesch@umd.edu

Anatomical, imaging, and lesion work have suggested that medial and lateral aspects of orbitofrontal cortex (OFC) play different roles in reward-guided decision-making, yet few single-neuron recording studies have examined activity in more medial parts of the OFC (mOFC) making it difficult to fully assess its involvement in motivated behavior. Previously, we have shown that neurons in lateral parts of the OFC (lOFC) selectively fire for rewards of different values. In that study, we trained rats to respond to different fluid wells for rewards of different sizes or delivered at different delays. Rats preferred large over small reward, and rewards delivered after short compared with long delays. Here, we recorded from single neurons in rat rostral mOFC as they performed the same task. Similar to the lOFC, activity was attenuated for rewards that were delivered after long delays and was enhanced for delivery of larger rewards. However, unlike lOFC, odor-responsive neurons in the mOFC were more active when cues predicted low-value outcomes. These data suggest that odor-responsive mOFC neurons signal the association between environmental cues and unfavorable outcomes during decision making.

Keywords: discounting, inhibition, orbitofrontal cortex, prediction, reward, single unit, value

Introduction

Orbitofrontal cortex (OFC) is involved in learning and reward-based decision-making (Kringelbach 2005; Schoenbaum and Roesch 2005; Murray et al. 2007; Wallis 2007; Kable and Glimcher 2009; Schoenbaum et al. 2009; Padoa-Schioppa 2011). Although OFC is often treated as a unitary structure, anatomical and imaging studies have suggested dissociable functions within subregions of the OFC (Carmichael and Price 1996; Elliott et al. 2000; O'Doherty et al. 2001; Kringelbach and Rolls 2004; McClure et al. 2004, 2007; Hoover and Vertes 2011; Kahnt et al. 2012; Wallis 2012). This dissociation has become clearer as researchers start to apply focal lesions to different aspects of the OFC in rats and primates (Iversen and Mishkin 1970; Noonan et al. 2010; Rygula et al. 2010; Mar et al. 2011; Rudebeck and Murray 2011a, 2011b). For example, work in nonhuman primates has shown that lateral OFC (lOFC) is critical for updating the value of objects during selective satiation, whereas medial OFC (mOFC) appears to be more critical for stopping responding when previously rewarded objects are no longer rewarded during extinction (Rudebeck and Murray 2011a, 2011b). Other primate labs report that lateral, not medial, OFC is critical for reward-credit assignment, whereas mOFC is necessary for normal reward-guided decision-making (Noonan et al. 2010).

In rats, a similar story is starting to develop (St Onge and Floresco 2010; Mar et al. 2011; Stopper et al. 2014). For example, a recent report showed that lesions to mOFC and lOFC make rats less and more impulsive, respectively, during performance of a standard delay-discounting task (Mar et al.

2011). In this task, rats chose between a large delayed reward and a small immediate reward. Over several trial blocks, the delay that preceded the large reward increased from 0 to 60 s. In tasks like these, rats initially choose the large reward, but gradually stop selecting it when the delay becomes longer. The delay at which the rat stops selecting the large reward reflects the impulsivity level of the rat. Mar and colleagues found that rats with mOFC lesions were less impulsive after extended postlesion training (i.e., continued to choose the large reward at longer delays compared with controls), whereas rats with lOFC lesions were more impulsive (i.e., selecting the large delayed reward less often than controls).

These datasets suggest that models of decision making that include the OFC must be revised to account for the functional dissociation between mOFC and lOFC. Unfortunately, the precise nature of the mOFC's involvement in decision making is still unclear, in part because few studies have examined activity in mOFC in behaving animals. The differential effects observed after lesions to more medial and lateral subregions suggest that neural correlates related to decision making and reward evaluation in the mOFC must be different than those that have been characterized in more lateral portions (Tremblay and Schultz 1999; Wallis and Miller 2003; Roesch and Olson 2004, 2005; Schoenbaum and Roesch 2005; Padoa-Schioppa and Assad 2006; Roesch and Olson 2007; Simmons et al. 2007; van Duuren et al. 2007; Wallis 2007; Kennerley and Wallis 2009; van Duuren et al. 2009; Bouret and Richmond 2010; Kennerley et al. 2011; Morrison and Salzman 2011; Morrison et al. 2011; Padoa-Schioppa 2011). Alternatively, neural processing related to these functions might be similar between these 2 structures and the differential loss of function after lesions might simply reflect the output structures that they project to (Morecraft et al. 1992; Carmichael and Price 1995a, 1995b, 1996; Price et al. 1996; Price 2007; Saleem et al. 2008; Schilman et al. 2008).

To address this issue, we recorded from single neurons in the rostral mOFC while rats performed an odor-guided task in which they chose between differently valued rewards. Value was manipulated by independently varying the expected delay to and size of the reward. At the time of reward delivery, reward-responsive neurons showed elevated firing for immediate and larger rewards. Unlike lOFC—and most reward-related brain areas for that matter—odor-responsive neurons in the mOFC fired significantly more strongly for odor cues that predicted a low value.

Materials and Methods

Subjects

Male Long-Evans rats ($n = 7$) were obtained at 175–200 g from Charles River Labs, Wilmington, MA, USA. Rats were tested at

the University of Maryland, College Park, in accordance with the University of Maryland and National Institute of Health guidelines.

Surgical Procedures and Histology

Rats had a drivable bundle of 10 25 μm diameter FeNiCr wires chronically implanted in the left or right hemisphere dorsal to mOFC ($n = 7$; 4.7 mm anterior to bregma, 0.5 mm laterally, and 2 mm ventral to the brain surface; Bryden, Johnson, Diao, et al. 2011). Electrode wires were housed in 27-G cannula. Immediately prior to implantation, the wires were freshly cut with surgical scissors to extend approximately 1 mm beyond the cannula and electroplated with platinum (H_2PtCl_6 , Aldrich, Milwaukee, WI, USA) to an impedance of approximately 300 k Ω . Brains were removed and processed for histology using standard techniques.

We define mOFC as rostral portions of the frontal cortex that include both ventral and medial aspects of the OFC according to Paxinos and Watson (1997). Solid gray bars in Figure 1e represent the estimated location of the recording electrodes based on histology. Electrode penetrations that crossed the coronal plane at which the forceps minor corpus callosum became visible and/or extended more laterally than 1.5 mm were excluded. Two rats were excluded due to the misplacement of electrodes (Fig. 1e, open boxes).

Behavioral Task

On each trial, a nose poke into the odor port after house light illumination resulted in delivery of an odor cue to a hemicylinder located behind this opening (Bryden, Johnson, Diao, et al. 2011; Roesch and Bryden 2011). One of 3 different odors (2-octanol, pentyl acetate, or carvone) was delivered to the port on each trial. One odor instructed the rat to go to the left to get reward, a second odor instructed the rat to go to the right to get reward, and a third odor indicated that the rat could obtain the reward at either well. Odors were counterbalanced across rats. The meaning of each odor did not change across sessions. Odors were presented in a pseudorandom sequence such that the free-choice odor was presented on 7 of 20 trials and the left/right odors were presented in equal in the remaining trials.

During recording, one well was randomly designated as short (500 ms) and the other long (1–7 s) at the start of the session (Fig. 1a: Block 1). In the second block of trials, these contingencies were switched (Fig. 1a: Block 2). The length of the delay under long conditions abided by the following algorithm: the side designated as long started off as 1 s and increased by 1 s every time that side was chosen on a free-choice odor (up to a maximum of 7 s). If the rat chose the side designated as long fewer than 8 of the previous 10 free-choice trials, the delay was reduced by 1 s for each trial to a minimum of 3 s. The reward delay for long free-choice trials was yoked to the delay in forced-choice trials during these blocks. In later blocks, we held the delay preceding reward delivery constant (500 ms) while manipulating the size of the expected reward (Fig. 1a: Blocks 3 and 4). The reward was a 0.05-mL bolus of 10% sucrose solution. For a big reward, an additional bolus or two was delivered 500 ms after the first bolus. At least 60 trials per block were collected for each neuron. Size blocks were always performed in Blocks 3 and 4 to offset changes in motivation that might occur due to satiety. Essentially there were 4

basic trial types (short, long, big, and small) by 2 directions (left and right). Conditions were pseudorandomly interleaved, so that no more than 3 trial types occur consecutively. Rats were water deprived (~20–30 min of free water per day) with free access on weekends.

Single-Unit Recording

Procedures were the same as described previously (Bryden, Johnson, Diao, et al. 2011; Bryden, Johnson, Tobia, et al. 2011). Wires were screened for activity daily; if no activity was detected, the rat was removed, and the electrode assembly was advanced 40 or 80 μm . Otherwise, active wires were selected to be recorded, a session was conducted, and the electrode was advanced at the end of the session. Neural activity was recorded using 4 identical Plexon Multichannel Acquisition Processor systems (Dallas, TX, USA) interfaced with odor discrimination training chambers. Signals from the electrode wires were amplified $\times 20$ by an op-amp headstage (Plexon, Inc., HST/8o50-G20-GR), located on the electrode array. Immediately outside the training chamber, the signals were passed through a differential preamplifier (Plexon, Inc., PBX2/16sp-r-G50/16fp-G50), where the single-unit signals were amplified $\times 50$ and filtered at 150–9000 Hz. The single-unit signals were then sent to the Multichannel Acquisition Processor box, where they were further filtered at 250–8000 Hz, digitized at 40 kHz, and amplified at $\times 1$ –32. Waveforms ($>2.5:1$ signal-to-noise) were extracted from active channels and recorded to disk by an associated workstation with event timestamps from the behavior computer and sorted in Offline Sorter using template matching (Plexon).

Data Analysis

Analysis epochs were computed by taking the total number of spikes and dividing by the time over which spikes were counted (firing rate). Neurons were first characterized by comparing firing rate during baseline with that in response to odors and rewards, averaged over all trial types (t -test, $P < 0.05$). Odor-related neural firing was examined over an analysis epoch that started 100 ms after onset of the odor and ended when the rat exited the odor port (“odor epoch”). To analyze reward-related activity (“reward epoch”) and licking rate, we examined activity on short- and long-delay trials starting at reward delivery and ending 1 s later. For large and small trials, the reward epoch started 500 ms after the delivery of the first bolus (i.e., delivery of the second bolus on large trials), and lasted 2 s to capture activity related to consumption of the additional reward. On average, rats spent 3.7 and 5.1 s in the well after reward delivery on small and big reward trials, respectively. Thus, this comparison captures activity when the rats were experiencing the extra boli, but were still in the fluid well. Finally, activity during the 500 ms prior to reward delivery on long-delay trials was examined to determine whether mOFC neurons fired in anticipation of the delayed reward. Baseline activity was taken during 1 s starting 2 s before odor onset (“baseline”).

These epochs (odor and reward) were used for each neuron to compute difference scores (value indices) between differently valued outcomes (i.e., short minus long; large minus small). Wilcoxon tests ($P < 0.05$) were used to measure significant differences between trial types at the population level, and to measure significant shifts from zero in distributions of value

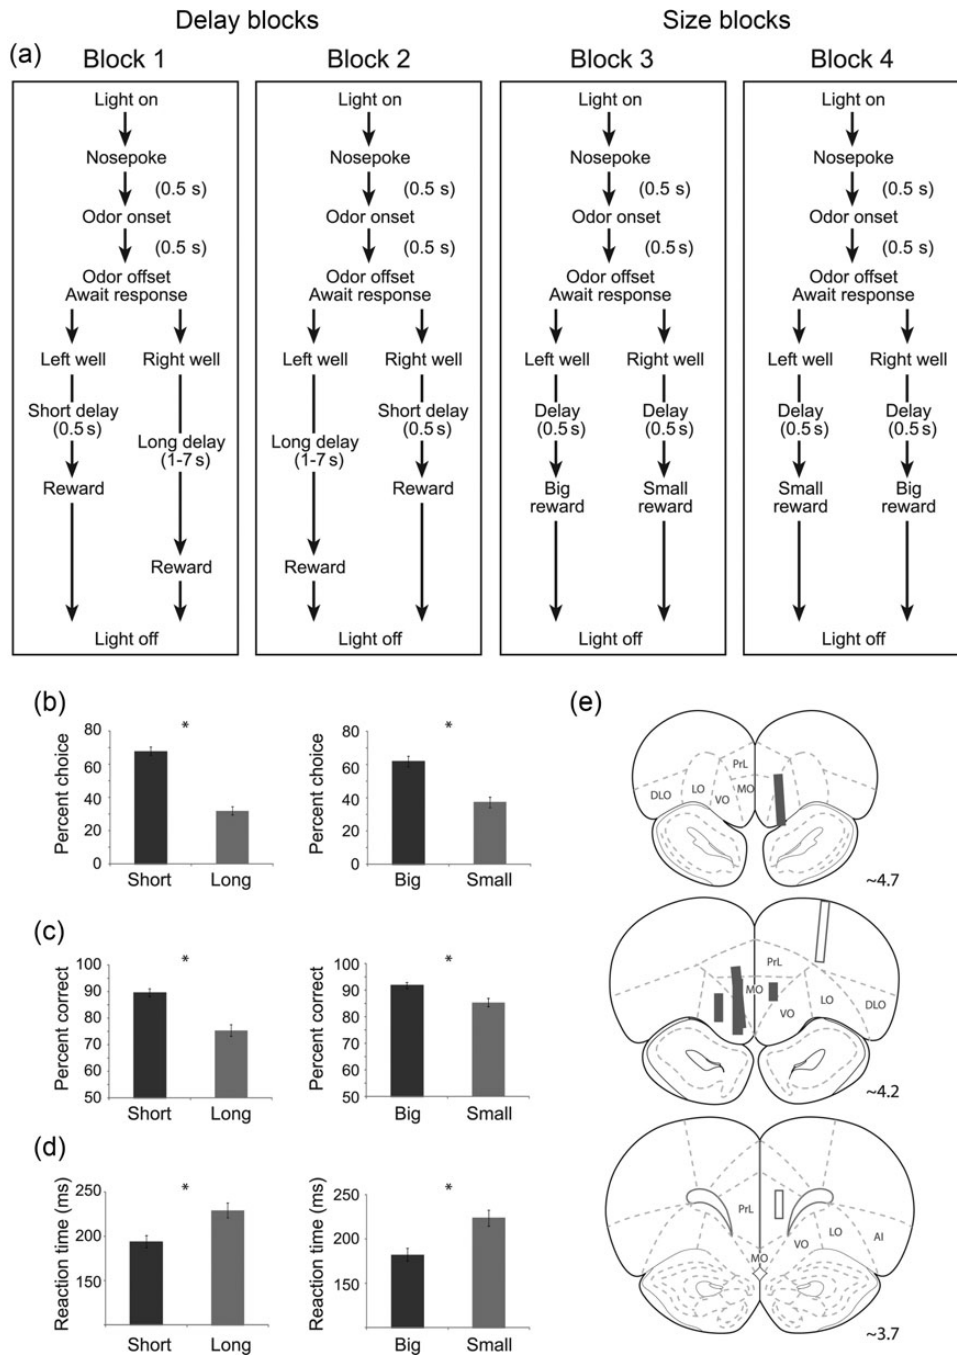


Figure 1. Task, behavior, and recording sites. (a) The sequence of events in each trial block. For each recording session, one fluid well was arbitrarily designated as short (a short 500-ms delay before reward) and the other designated as long (a relatively long 1- to 7-s delay before reward) (Block 1). After the first block of trials (~60 trials), contingencies unexpectedly reversed (Block 2). With the transition to Block 3, the delays to reward were held constant across wells (500 ms), but the size of the reward was manipulated. The well designated as long during the previous block now offered 2–3 fluid boli, whereas the opposite well offered 1 bolus. The reward stipulations again reversed in Block 4. (b) The impact of delay length (right) and reward size (left) manipulations on choice behavior during free-choice trials. (c) Impact of value on forced-choice trials for short versus long delay (left) and big versus small rewards (right). (d) Reaction times (odor offset to nose unpoke from odor port) on forced-choice trials (expressed in ms) comparing short- versus long-delay trials and big- versus small-reward trials. Only rats that contributed to the neural dataset were included in the behavioral analysis (b–d; $n = 5$). (e) Location of recording sites. Filled gray boxes mark the locations of electrodes based on histology. Electrode wires were housed in a 27-G cannula. Shown are representative slices of 4.7-, 4.2-, and 3.7-mm coronal sections anterior to bregma from Paxinos and Watson (1997). The center of the majority of recording electrodes fell in between 4.2 and 4.7 mm anterior to bregma. One electrode was more anterior, centered roughly at 4.5–4.7 mm anterior to bregma. Rats were excluded from analysis if the electrode track crossed the plane at which the forceps minor corpus callosum became visible to avoid the contribution of more posterior medial prefrontal cortical regions (~3.7). Open gray boxes represent recording sites excluded due to being too lateral or too posterior. Asterisks indicate planned comparisons revealing statistically significant differences (t -test, $P < 0.05$). Error bars indicate standard errors of the mean (SEMs). PrL: prelimbic; MO: medial orbital; VO: ventral orbital; LO: lateral orbital; DLO: dorsolateral orbital; AI: agranular insular.

indices, which were not normally distributed (Jarque-Bera; $P < 0.05$). Analyses of variance (ANOVAs) and t -tests were used to measure differences between baseline and analysis epochs,

and between trial types at the single-cell level ($P < 0.05$). The activity of neurons for which we examined differences between trials types at the single-cell level only violated normality

(Jarque-Bera; $P < 0.05$) in 7% of neurons, which is fewer than expected from chance alone (χ^2 ; $P = 0.32$). Multiple regression analyses with delay and size as factors were performed on individual units during the odor and reward epochs ($P < 0.05$).

Results

Rats were trained on the behavioral choice task in which we manipulated reward size and delay. The sequence of events is described in the Methods and depicted in Figure 1. Rats started each trial by nose poking into a central odor port. After 500 ms, 1 of 3 odors was delivered. Two of the odors signaled to the rat to move to the left or right to receive the reward (forced-choice odors). A third odor indicated that they could choose either well to receive the reward (free-choice). The delay to and size of reward were independently varied in different trial blocks (Fig. 1*a*). On average, rats performed 244 correct trials per session.

Rats perceived differently delayed and sized rewards as having different values across all 4 trial blocks. On free-choice trials, rats chose the well associated with large reward and short delay significantly more often than the well associated with small reward and long delay, respectively. This was significant across all recording sessions (Fig. 1*b*; t -test; $t_{68}' > 10$; $P < 0.05$) and individually for each rat (t -test; $t_{68}' > 4$; $P < 0.05$). On forced-choice trials, rats were more accurate and faster on large-reward and short-delay trials, when compared with their respective counterparts (Fig. 1*c,d*; t -test; $t_{68}' > 5$; $P < 0.05$). The impact of expected value on forced-choice trials was also consistent across rats. Each individual rat exhibited significantly faster responding on short-delay and large-reward trials (t -test; $t_{68}' > 3$; $P < 0.05$). Finally, rats were motivated across all 4 trial blocks; reaction times were not significantly different between delay and size blocks (t -test; $t_{68} = 1.2$; $P = 0.24$), and the differences between differently valued outcomes were significant in all 4 trial blocks (t -test; $t_{68}' > 3$; $P < 0.05$). Thus, performance on free- and forced-choice trials was modulated by the predicted outcomes in both size and delay trial blocks.

Reward-Related Activity in the mOFC Was Stronger for an Immediate and Large Reward

We recorded from 251 rostral mOFC neurons in 5 rats ($n' = 9, 31, 36, 83,$ and 92) from 69 sessions. We first characterized neurons as being reward-responsive by asking how many neurons showed significantly higher firing during reward delivery (reward epoch) compared with baseline (baseline epoch; t -test; $P < 0.05$). The average baseline firing was 4.15 spikes/s ($n = 251$; 1 s epoch before nose poke). Of the 251 neurons, 56 neurons [22%; $n = 3$ (33%), 5 (16%), 12 (33%), 15 (18%), and 21 (23%)] showed significant increases in firing over baseline, which is more than expected by chance alone (χ^2 ; $P < 0.05$).

As illustrated by firing of the single-cell examples in Figure 2*a,b*, and for the population ($n = 56$) in Figure 2*c,d*, activity was often higher when the reward was large (Fig. 2*a,d*; dark gray) or delivered after a short delay (Fig. 2*a-c*; dark gray), compared with when the reward was small or delayed by several seconds (dark gray vs. gray; reward epoch; single unit: t -test, $P < 0.05$; population: Wilcoxon; $P < 0.05$). To quantify these effects, for each reward-responsive neuron, we plotted difference scores between firing during short and long,

and large and small rewards, and asked in how many neurons was there significant differential firing within each value manipulation on forced-choice trials (ANOVA; regression; $P < 0.05$; reward epoch). The distributions of value indices are plotted in Figure 2*e*. For both delay (Fig. 2*ei*) and size (Fig. 2*eiii*), the distributions were significantly shifted (Wilcoxon; $P < 0.05$) in the positive direction, indicating that the majority of mOFC neurons fired more strongly for high- compared with low-value reward (i.e., short and large over long and small, respectively). The 2 effects were not correlated (Fig. 2*eii*; $P = 0.37$; $r^2 = 0.01$), suggesting that neurons that increased firing for one value manipulation did not show the same change for the other value manipulation as one would expect if activity in the mOFC reflected some sort of common-currency encoding (Roesch et al. 2006).

Finally, the counts of reward-responsive neurons that fired significantly more strongly for high- compared to low-value outcomes (ANOVA; $P < 0.025$; Bonferroni) were in the significant majority [Fig. 2*e*; black bars; 20 (49%) vs. 9 (22%); χ^2 ; $P < 0.05$]. To further illustrate the significance of this result at the single-unit level, we performed a multiple regression analysis with delay and size as factors (Roesch et al. 2006). During the reward epoch, 29 (52%) and 14 (25%) of reward-responsive neurons showed a positive and negative correlation with a value in the multiple regression, respectively ($P < 0.05$).

Firing to Delayed Rewards

Reward-related activity in the mOFC appeared to reflect the anticipation of reward. During trials in which the delay was only 500 ms (i.e., short, big, and small), activity started to rise prior to reward onset. For long-delayed rewards, significant increases in firing did not occur until after reward delivery. This is evident in the single-cell example shown in Figure 2*a* and across the population (Fig. 2*c*). In Fig. 2*a,c*, activity for rewards delivered after short delay increased firing during the 500 ms preceding reward delivery (dark gray), whereas activity did not show a change in firing during long-delay trials until reward was actually delivered at time zero (light gray). This likely reflects the difficulty that animals have in timing rewards that are delayed by several seconds as measured by anticipatory licking (Fiorillo et al. 2008; Kobayashi and Schultz 2008; Takahashi et al. 2009).

Indeed, rats in our study licked more in anticipation of reward delivery on short- compared with long-delay trials. The average lick rate during the 250 ms before reward delivery was significantly higher for short-delay trials (vs. long-delay; t -test; $t_{55} = 8.0$; $P < 0.05$), suggesting that they could better anticipate the more immediate reward. The strength of this difference was correlated with the difference in firing in the mOFC during short- and long-delay trials during the reward epoch (Fig. 2*f*; $P = 0.05$; $r^2 = 0.07$), suggesting that when rats could better anticipate reward delivery, activity was stronger in the mOFC.

This correlation might suggest that activity in the mOFC simply reflected motor commands that are coupled to value, such as licking, orofacial movements, and swallowing (Gutierrez et al. 2006). Although it is difficult to rule this out, we do not think this is the case because the average licking rate during the reward epoch was not correlated with the average firing rate during the same period (reward epoch; $P = 0.63$; $r^2 = 0.004$) during performance of size blocks. In addition, as we will describe below, activity in the mOFC represented the

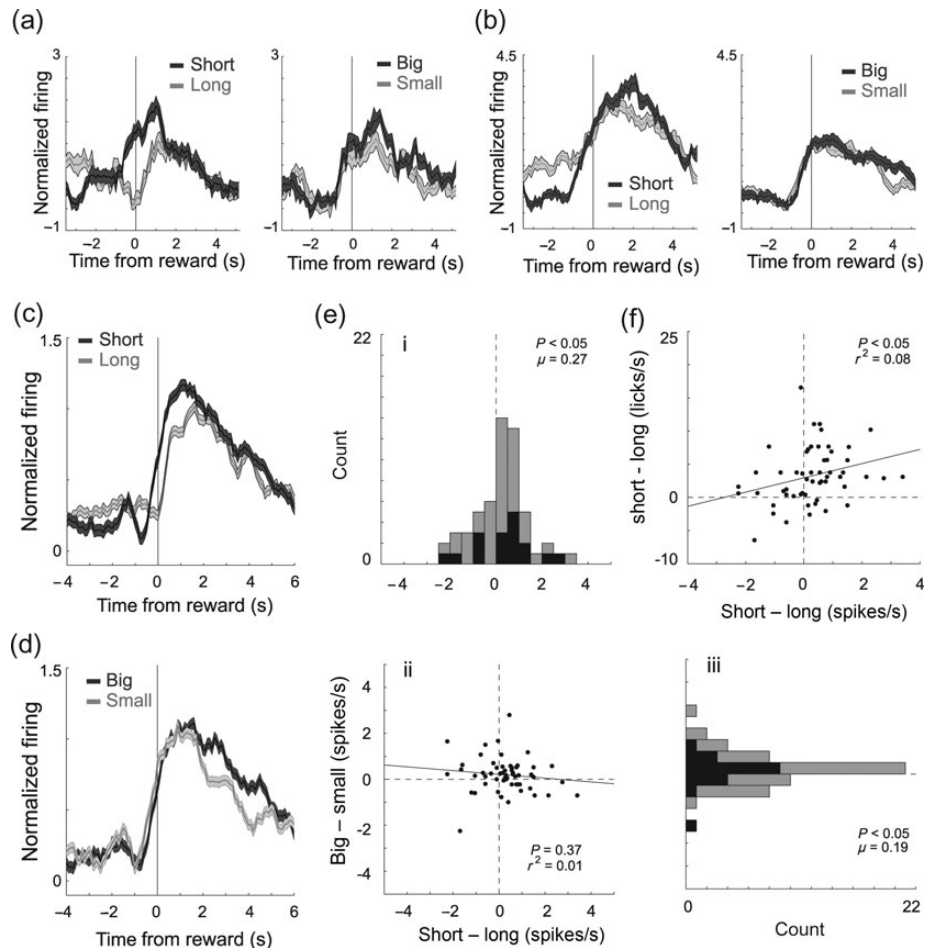


Figure 2. Reward-related activity in the OFC was stronger for an immediate and large reward. (a–d) Histograms representing the activity of single cells (a–b) and across the population (c–d) of reward-responsive neurons ($n = 56$; 22%) in the mOFC during task performance of delay (dark gray = short; light gray = long) and size (dark gray = big; light gray = small) blocks. Activity is aligned to reward delivery (time zero). For short, big, and small trials, well entry occurred 500 ms before reward delivery. On long-delay trials, well entry was 1–7 s before reward delivery. Neurons were selected by comparing activity during the reward epoch when compared with baseline (see text; t -test; $P < 0.05$). Activity is normalized by subtracting the mean and dividing by the standard deviation (z -score). Bins are 100 ms. Thickness of line reflects standard error of the mean (SEM). Note that activity that precedes the 500 ms before reward delivery for rewards that were delivered after a short delay (short, big, and small) cannot be directly compared with activity that precedes 500 ms before reward delivery on long-delay trials, because task events (well entry, port exit, etc.) occur at different time points across these trial types. (e) Correlation (ii) between difference scores for size and delay blocks (i.e., short minus long (i) and large minus small (iii)). Neural activity was taken during the reward epoch. Black bars in distribution histogram represent neurons that showed a significant difference between differently valued outcomes ($P < 0.05$; main or interaction effect of value in a 2-factor ANOVA; reward epoch). (f) Correlation between value indices for the licking rate in anticipation of reward (250 ms before reward delivery) and for the firing rate during the reward epoch (1 s after reward) on short- and long-delay trials (short-long). All data are taken from forced-choice trials.

spatial direction of the movement, whereas the licking rate was not significantly modulated by the spatial location (reward epoch; t -test; $t_{55} = 0.06$ $P = 0.95$). To further address this issue, we examined activity and licking 2–5 s after reward delivery, during which time activity in the mOFC might have reflected prolonged licking associated with consumption of the large reward. Even during this extended period, the correlation between licking and firing rates was not significant ($P = 0.18$; $r^2 = 0.04$).

Although reward-related activity was attenuated on long- relative to short-delay trials as in Figure 2a, other neurons did maintain firing during anticipation of rewards delayed by several seconds (long-delay trials). This is best illustrated by the single-cell example in Figure 2b (long). Activity immediately preceding reward delivery (500 ms) was significantly higher compared with baseline when rats were waiting for the delayed reward (light gray). Of the 56 reward-responsive neurons, 27 (48%; t -test; $P < 0.05$) exhibited significantly higher firing during the

500 ms before reward delivery compared with baseline, whereas only 7 fired significantly less (χ^2 ; $P < 0.05$), demonstrating that many single neurons in the mOFC fired in anticipation of the delayed reward.

Odor-Evoked Activity in the OFC Was Stronger for Cues That Predict Long Delay and Small Reward

Then, we examined activity during odor sampling. Of the 251 total mOFC neurons, 41 [16%; $n = 1$ (11%), 1 (3%), 4 (11%), 17 (21%), and 18 (20%)] fired significantly more strongly during odor sampling compared with baseline (odor epoch; t -test; $P < 0.05$; χ^2 ; $P < 0.05$). Surprisingly, neurons in the mOFC fired significantly more strongly for odor cues that predicted low-value outcomes. This is illustrated in the single-cell examples and across the entire population of odor-responsive mOFC neurons in Figure 3a–d. Immediately after odor onset, before initiation of the behavioral response (port exit), population activity was significantly higher when small versus large

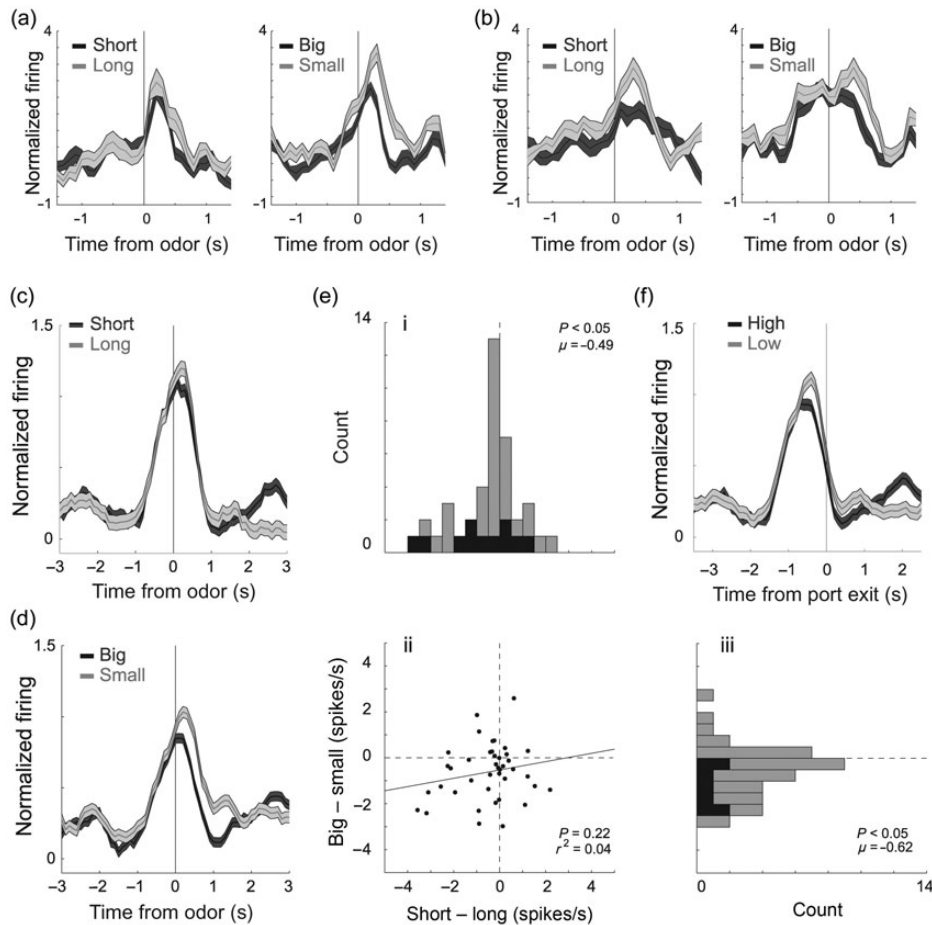


Figure 3. Odor-evoked activity in the OFC was stronger for cues that predict long delay and small reward. (a–d) Histograms representing the activity of single cell (a–b) and across the population (c–d) of odor-responsive neurons ($n = 41$; 16%) in the mOFC during task performance of delay (dark gray = short; light gray = long) and size (dark gray = big; light gray = small) blocks aligned to odor onset. For short, big, and small trials reward occurred several seconds later. This included the 500 ms of odor delivery and prefluid delay plus the intervening time taken to respond to the odor and move to the odor port. On long-delay trials, reward occurred an additional 0.5–6.5 s later, thus cannot be examined in this figure. (e) Correlation (ii) between difference scores for size and delay blocks [i.e., short minus long (i) and large minus small (iii)]. Activity was taken during the odor epoch (100 ms after odor onset to odor port exit). Black bars in distribution histogram represent neurons that showed a significant difference between differently valued outcomes ($P < 0.05$; main or interaction effect of value in a 2-factor ANOVA; odor epoch). (f) Activity is aligned to odor port exit to show that differences in activity between high- and low-value outcomes was not a product of different reaction times. Neurons were selected by comparing activity during the odor epoch with baseline (1 s before nosepoke; t -test; $P < 0.05$). Activity is normalized by subtracting the mean and dividing by the standard deviation (z-score). Bins are 100 ms. Thickness of line reflects standard error of the mean (SEM). All data are taken from forced-choice trials.

reward was predicted (Fig. 3d; dark vs. light gray; odor epoch; Wilcoxon; $P < 0.05$) and when long versus short delay was predicted (Fig. 3c; dark vs. light gray; odor epoch; Wilcoxon; $P < 0.05$).

To further quantify this effect, we computed the difference scores between high- and low-value rewards and asked in how many odor-responsive units did forced-choice odors that predicted low-value reward elicit significantly stronger firing (Fig. 3e; ANOVA; odor epoch; black bars). The number of neurons that fired more strongly for odors that predicted a low value (ANOVA; $P < 0.025$; Bonferroni) were in the significant majority [11 (26%) vs. 2 (5%); $\chi^2 = 3.8$; $P < 0.05$]. To further illustrate the significance of this result at the single-unit level, we performed a multiple regression analysis with delay and size as factors during the odor epoch (Roesch et al. 2006). As expected from the ANOVA, 1 (2.4%) and 10 (24%) showed a positive and negative correlation with the value, respectively ($P < 0.05$).

At the population level, both delay (Fig. 3ei) and size (Fig. 3eiii) distributions were significantly shifted in the negative direction (Wilcoxon; $P < 0.05$). Although delay effects appeared weaker than size effects, the 2 distributions were not significantly different (Wilcoxon; $P = 0.75$). As we will describe below, stronger differences emerge when trials are broken down by the direction.

Although the 2 effects appeared to be correlated—most of the cells fell in the bottom left quadrant—the correlation between size and delay indices was not significant (Fig. 3eii; $P = 0.20$; $r^2 = 0.04$). This indicates that those neurons that fired more strongly for cues that predicted longer delays were not necessarily those that fired more strongly when the same cue predicted small reward, and vice versa, even though the overall effect was one of higher firing for longer delay and smaller reward predicting cues. As in the IOFC, this suggests that encoding in the mOFC does not reflect some sort of common-currency encoding for expected rewards (Roesch et al. 2006).

Increased activity during odors that predict a low-value reward does not simply reflect the fact that rats spent more time in the odor port. Activity in the mOFC does not continue to fire until port exit. This is illustrated in Figure 3*f*, which aligns activity on odor port exit for high- and low-value trials averaged over size and delay manipulations. Note that activity on low- and high-value trials peak and come back together before port exit. Further, if we repeat the analysis that examines difference scores described in the previous paragraph with an analysis epoch that is cut off at 100 ms after odor offset, instead of being variable to port exit, the results remain the same; both delay and size distributions are significantly shifted in the negative direction (Wilcoxon; $P < 0.05$).

Thus, activity during reward delivery and odor sampling in the mOFC carry different signs in relation to rewarding outcomes, with odor- and reward-related activity being stronger and weaker for low-value outcomes, respectively. To determine whether the 2 effects were correlated, we performed a regression analysis on value indices during the odor and reward epoch for all odor- and reward-responsive neurons ($n = 97$). The correlation between the 2 was not significant ($r^2 = 0.03$; $P = 0.10$), suggesting that neurons that tended to fire more strongly for cues that predicted a low-value reward did not tend to fire more strongly during delivery of high-value outcomes.

Encoding of Response Direction in the mOFC

Previous results have shown that activity in rat IOFC responds differently depending on the direction of the behavioral

response (Feierstein et al. 2006; Roesch et al. 2006, 2007). Here, we asked if activity in the mOFC was also modulated by response direction. Of the 41 odor-responsive neurons, 9 (22%) showed a significant main or interaction effect with response direction in a 2-factor ANOVA ($P < 0.025$; χ^2 ; $P < 0.05$) as illustrated by the single-cell example in Figure 4*a*; activity was highest when odor cues predicted a small reward in the left well. To further qualify this, we broke down the population activity into each cell's preferred and nonpreferred direction (Fig. 4*b*), as defined by the direction that elicited the strongest response (e.g., left in the example). Here, "preferred" refers to the direction that elicited the highest activity, not the outcome favored by the animal. Across the population of odor-responsive cells, the difference between high- and low-value outcomes appeared to be stronger for responses made in the preferred direction (filled dark vs. light gray) compared with the nonpreferred direction (open dark vs. light gray). In the preferred direction, the distribution of value indices was shifted in the negative direction, indicating stronger firing for a lower-value reward (Fig. 4*c*; black; Wilcoxon; $P < 0.05$). In the nonpreferred direction, the distribution of value indices was not significantly shifted (Fig. 4*c*; gray; Wilcoxon; $P = 0.88$) and was significantly different than that of indices obtained from preferred direction trials (Fig. 4*c*; Wilcoxon; $P < 0.05$).

Of the 56 reward-responsive neurons, 21 (38%) showed a significant main or interaction effect with response direction in the ANOVA ($P < 0.025$; χ^2 ; $P < 0.05$). In addition, the difference between high- and low-value outcomes appeared to be stronger for responses made in the preferred direction (Fig. 5*b*;

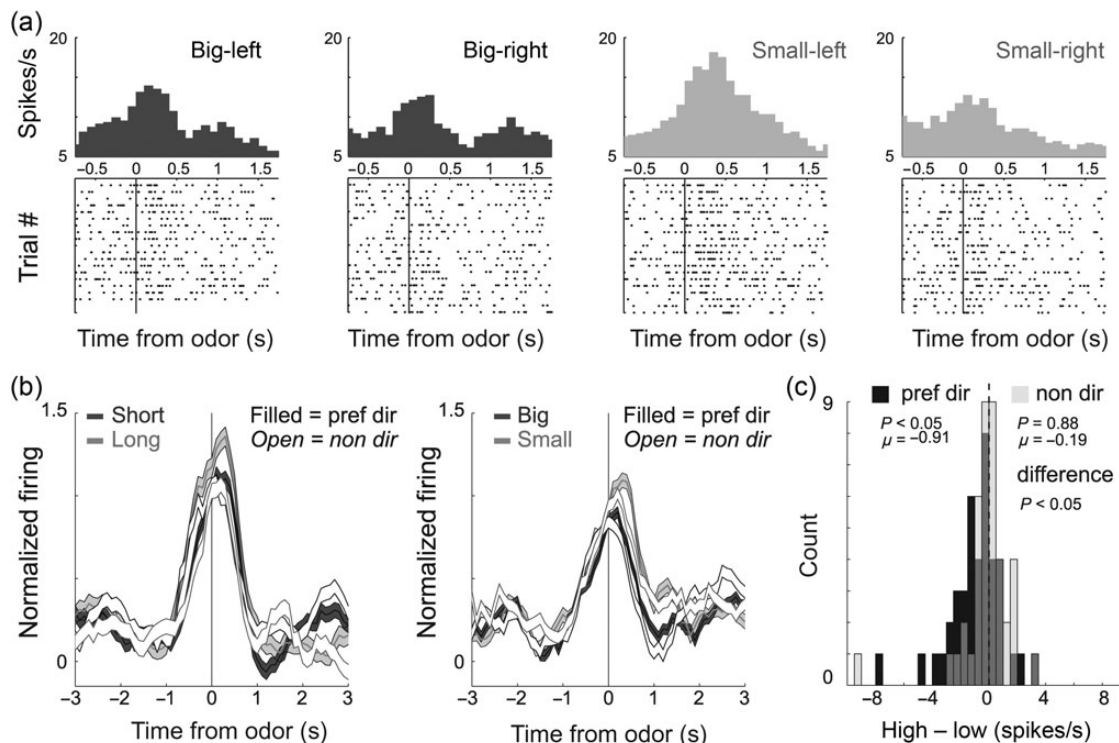


Figure 4. Odor-responsive neurons in the mOFC were directionally selective. (a) Activity of a single cell during size blocks demonstrating higher firing when odor cues predicted a small reward on the left. (b) Average firing rate over all 41 odor-responsive neurons broken down by preferred and nonpreferred response direction. Preferred direction was defined for each cell by determining which trial type elicited the strongest firing. Filled = preferred direction; open = nonpreferred direction; dark gray = high value (short and big); light gray = low value (long and small). Activity is normalized by subtracting the mean and dividing by the standard deviation (z-score). Bins are 100 ms. Thickness of line reflects standard error of the mean (SEM). (c) Distribution of value indices taken during the odor epoch (see Methods) independently for preferred (black) and nonpreferred response directions (light gray). Light gray distributions are transparent, and dark gray thus indicates where black (preferred) and light gray (nonpreferred) overlap. The Wilcoxon test were used to determine whether the 2 distributions were significantly different from zero and from each other ($P < 0.05$).

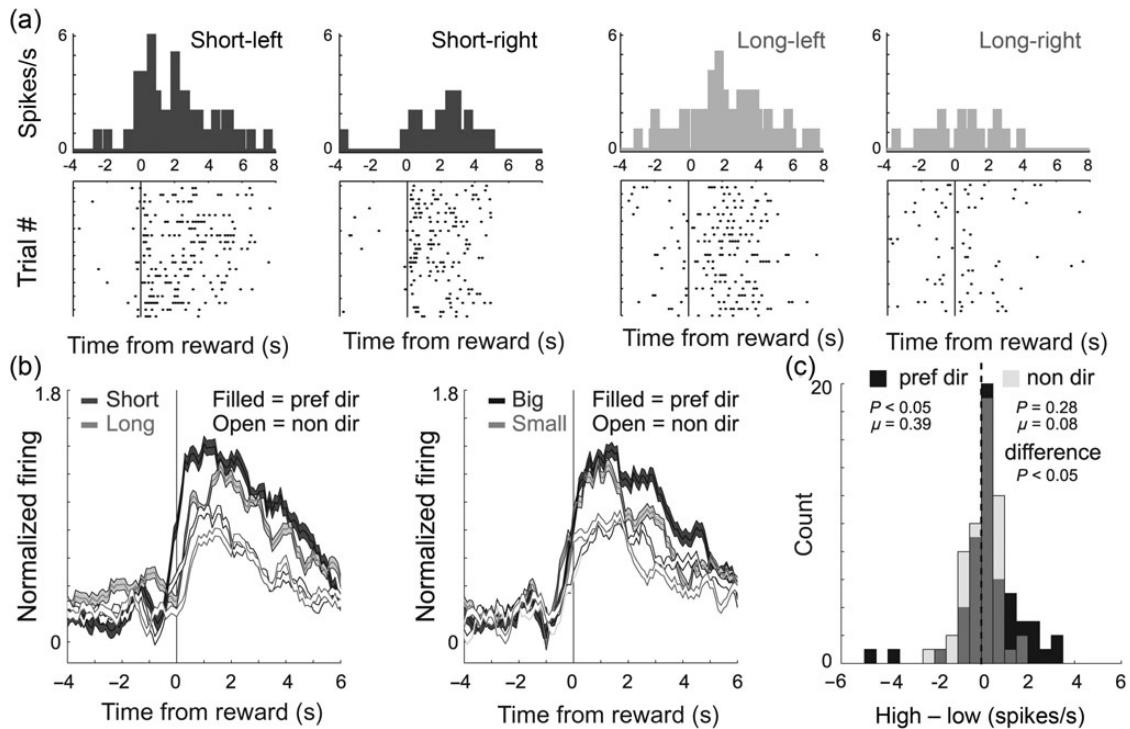


Figure 5. Reward-responsive neurons in the mOFC were directionally selective. (a) Activity of a single cell during delay blocks demonstrating higher firing during short- versus long-delay trials at the time of reward delivery. (b and c) Same as b–c in Figure 4, except for the 56 reward-responsive neurons (reward epoch).

filled dark gray vs. light gray). In the preferred direction, values were shifted in the positive direction indicating stronger firing for higher-value reward (Fig. 5c; black; Wilcoxon; $P < 0.05$). The distribution of value indices in the nonpreferred direction was not significantly shifted (Fig. 5c; gray; Wilcoxon; $P = 0.28$) and was significantly different from that of value indices obtained on preferred direction trials (Wilcoxon; $P < 0.05$). Thus, we conclude that odor- and reward-responsive mOFC neurons showed enhanced value encoding in the cell's preferred response direction.

Emergence of Outcome Selectivity During Odor Sampling and its Relation to Behavior

In a final analysis, we examined how selectivity for cues that predicted low-value outcomes emerged during learning and whether activity in the mOFC was correlated with reaction time. Figure 6a,b plots the average firing rate over delay and size blocks for movements made in the preferred direction during the first and last 5 trials for each trial type. Consistent with the previous sections, activity after learning was stronger for low-value outcomes (Fig. 6a,b; solid gray vs. black). Interestingly, this selectivity developed as a result of decreased firing on high-value trials that occurred with learning (Fig. 6a, b; black dashed vs. black solid). That is, activity was significantly lower for short-delay and large-reward trials at the end of the trial block relative to the beginning (odor epoch; Wilcoxon; P 's < 0.05). This relationship did not exist for cues that predicted low-value rewards (Fig. 6a,b; gray dashed vs. gray solid). Thus, selectivity emerged through a reduction in firing for cues that predicted a high-value reward.

This is further quantified in Figure 6c, which plots the normalized difference between high- and low-value outcomes during the first 5 (black) and last 5 (gray) trials. The

distribution is significantly shifted below zero only after learning (Wilcoxon; $P < 0.05$) and significantly different than during early trials (Wilcoxon; $P < 0.05$). Differences in firing reflected the rats' behavior in that value-induced differences in reaction time (faster for high-value reward) were present during late (Fig. 6d; gray; $P < 0.05$), but not early, trials (Fig. 6d; black; Wilcoxon; $P = 0.81$). Furthermore, the change in firing that occurred over the course of the trial block was significantly correlated with the strength of learning that occurred during the session as measured by changes in reaction time (Fig. 6e; $r^2 = 0.09$; $P < 0.05$).

These results suggest that mOFC may serve to alter behavior when low-value rewards are predicted by forced-choice cues. If true, then neural selectivity observed after learning, during odor sampling, might be correlated with reaction time differences observed between cues that predict high- and low-value outcomes. Consistent with this hypothesis, strong enhancement of the firing rate during odor sampling was correlated with slower behavioral responses. This is illustrated in Figure 7, which plots the value index (high – low/high + low) computed on average firing rates during the odor epoch against the value index computed for reaction times during those trials. As expected from the analysis above, both indices were negative, indicating slower reaction times and higher firing on low-value trials. Furthermore, both were correlated, demonstrating that when rats showed stronger reaction time differences, neural selectivity in the mOFC was enhanced ($P < 0.05$; $r^2 = 0.21$).

Discussion

Consistent with imaging and anatomical studies, recent lesion work in rats and primates has shown that subregions of the

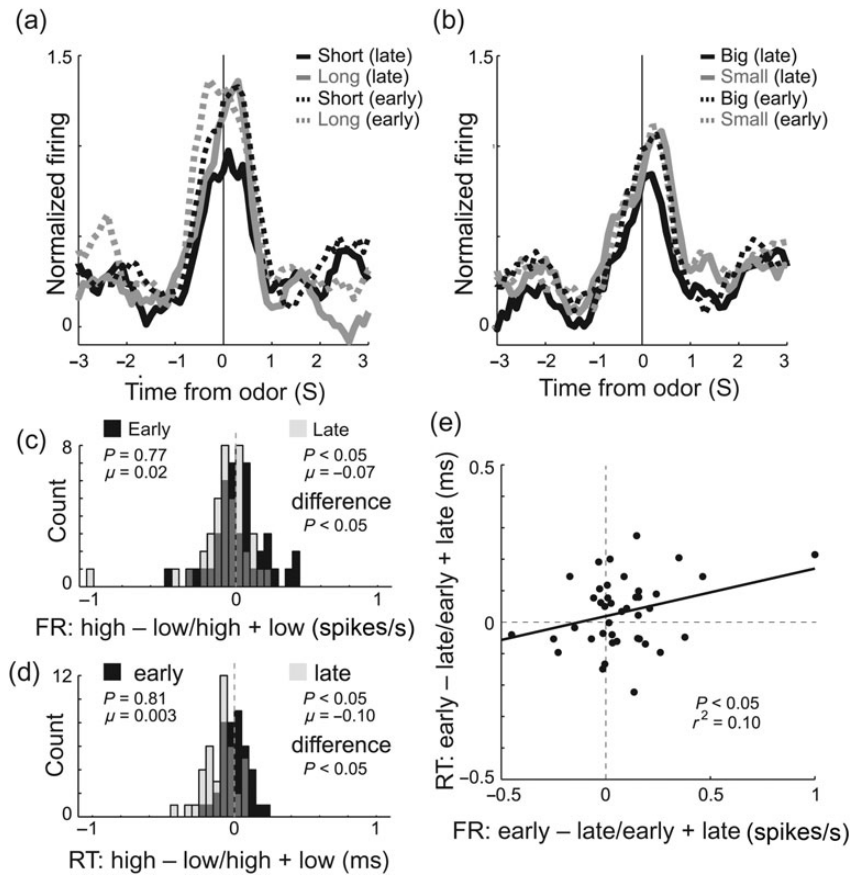


Figure 6. Emergence of cue selectivity during learning. (a and b) Population activity for the 41 odor-responsive neurons for responses made in the preferred direction, averaged over free- and forced-choice trials, for delay (a) and size (b) blocks. For each trial type, the average of the first (dashed) and last (solid) 5 trials in a block are shown. Black = short or large; gray = long or small. (c) Distribution of value indices (high – low/high + low) reflecting the firing rate (odor epoch) difference between high- and low-value outcomes, early (first 5 trials; black) and late (last 5 trials; late) during learning. (d) Distribution of value indices (high – low/high + low) reflecting the reaction time difference between high- and low-value outcomes, early (first 5 trials; black) and late (last 5 trials; late) during learning. (e) Scatter plot represents the correlation between changes in firing and in reaction time that occur during learning (early – late/early + late) on high-value reward trials. FR: firing rate; RT: reaction time. The Wilcoxon test were used to determine whether the 2 distributions were significantly different from zero and from each other ($P < 0.05$).

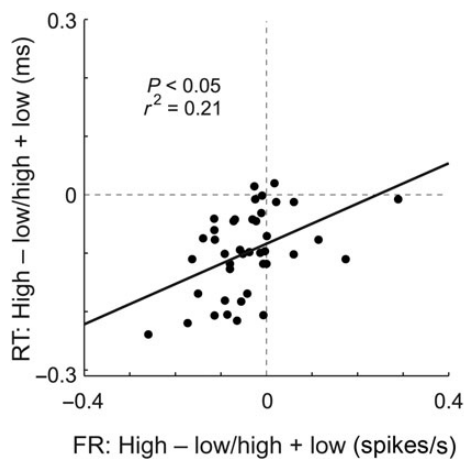


Figure 7. Correlation between reaction time (RT) and firing rate (FR) collapsed across both value manipulations. Scatter plot represents the correlation between high- and low-value trial-type differences for reaction time (odor offset to odor port exit) and neural firing (odor epoch) averaged across value manipulation and direction. Value index = high – low/high + low.

OFC perform different functions (Noonan et al. 2010; Mar et al. 2011; Rudebeck and Murray 2011a, 2011b). However, few

studies have examined activity in the mOFC making it difficult to understand the exact nature that mOFC plays in reward-guided decision-making. Here, we demonstrate that reward and cue-evoked responses in the mOFC are modulated by the size of and delay to reward, 2 value manipulations that clearly impact decision making. At the time of reward delivery, activity was higher when outcomes were of higher value. During odor sampling, the opposite effect was observed, that is, firing was higher for odor cues that predicted low-value outcomes in odor-responsive neurons. Below we discuss these results, comparing mOFC to activity in other areas, including our own work in IOFC, with the caveat that these comparisons are made across studies, in different rats, and that neurons might have been sampled from different layers considering the structure of these 2 areas.

In previous reports, we characterized firing in the IOFC as rats performed the same task (Roesch et al. 2006; Takahashi et al. 2009; Roesch et al. 2012). Activity related to reward expectancy and delivery was similar across mOFC and IOFC in that overall activity was reduced when rewards were delayed. The major difference between these 2 subregions emerged during the sampling of odors that predicted different outcomes. Although neurons that show increased firing for cues

that predict a low-value reward have been described previously in the IOFC, the proportion of neurons in the mOFC showing this effect were in the majority, and the population response was stronger over all neurons when cues signaled a low value. This makes mOFC unique among brain areas thought to be critical in reward-guided decision-making; most reward-related regions in the brain fire more strongly for cues that predict a more valuable reward, including IOFC (Tremblay and Schultz 1999; Roesch and Olson 2004, 2005; Schoenbaum and Roesch 2005; Padoa-Schioppa and Assad 2006; Roesch and Olson 2007; Simmons et al. 2007; van Duuren et al. 2007; Wallis 2007; Kennerley and Wallis 2009; van Duuren et al. 2009; Bouret and Richmond 2010; Kennerley et al. 2011; Padoa-Schioppa 2011).

Such a signal might be important for inhibitory control and/or complement more common response bias signals that are elevated when animals expect better rewards. Consistent with this idea, reports in humans and rats have shown that mOFC dysfunction causes subjects to make more risky decisions, possibly due to a disruption in inhibitory control over biases to select riskier rewards (Clark et al. 2008; St Onge and Floresco 2010; Zeeb et al. 2010; Stopper et al. 2014). Furthermore, our data indicate that increased firing in the mOFC during the sampling of odors that predicted low-value outcomes was positively correlated with differences in reaction time, indicating that when activity was high, reaction times were slow. However, we must exercise some caution here because increased firing on low-value trials might also be interpreted as signals that allow for a behavioral response to occur, albeit away from the more valued outcome.

In none of our studies, have we found a single brain area that increased population firing to cues that predicted a low-value reward (Roesch and Bryden 2011). This includes brain areas that are in relatively close proximity to our recording sites, such as medial prefrontal cortex (mPFC) and IOFC (Roesch et al. 2006, 2012; Gruber et al. 2010). Unfortunately, in this study, we cannot dissociate between medial and ventral OFC because our sample size was too low. With that said, we observed no obvious differences between the 2 regions, but future work is necessary to determine if they carry different signals. Although connections of the medial and ventral OFC do overlap, recent work based on connectivity has suggested that ventral and medial OFC might play different roles comparable with IOFC and mPFC, respectively, and that both ventral and medial aspects of OFC might serve as a link between IOFC and mPFC (Hoover and Vertes 2011; Kahnt et al. 2012; Wallis 2012). Findings such as these make it difficult to draw a hard line between mOFC and mPFC. Regardless of whether you consider this region part of the OFC or PFC, we show that predicted outcome encoding is considerably different relative to the IOFC (Roesch et al. 2006, 2012) and PFC (Gruber et al. 2010), consistent with recent lesion work targeting this specific region (Mar et al. 2011; Stopper et al. 2014).

To the best of our knowledge, there is only one other single-unit study that has shown elevated firing in the majority of cue-responsive neurons when animals anticipate a small reward. In that paper, monkeys performed a go/no-go task for a predicted large or small reward (Minamimoto et al. 2005). They found increased activity in the centromedian nucleus of the thalamus (CM) when monkeys made actions (go or no-go)

for a small compared with a large reward. Further, they showed that stimulation of CM caused typically speeded reactions to be slow, demonstrating a role for CM in a mechanism complementary to more common signals that are thought to bias animals toward better reward. Although connectivity between mOFC and CM is relatively light, it is possible that interactions between them are critical for reward-guided behaviors (Hoover and Vertes 2011; Vertes et al. 2012).

For decades, it was thought that OFC was critical for response inhibition because damage to OFC made animals and humans lose aspects of inhibitory control and become more impulsive in their actions (Damasio et al. 1994; Bechara et al. 2000; Berlin et al. 2004; Torregrossa et al. 2008; Schoenbaum et al. 2009). Here, we show that activity was high during situations in which the animal had to inhibit responding at the beginning of trial blocks and when forced-choice trials instructed the rat to respond away from the desired reward toward the low value well. Loss of signaling of unfavorable outcomes during decision making and learning could account for many of the deficits thought to reflect deficits in inhibition. Interestingly, if the role of this signal is to inhibit behavioral output, it appears to do so in an outcome-specific manner, because the correlation between selective firing during size and delay blocks was not significant (Fig. 3e), suggesting that OFC does not output a simple general/global inhibition signal.

One common way to assess response inhibition and impulsivity is to conduct delay-discounting procedures in which animals choose between small immediate rewards and large rewards delivered after long delays (Cardinal et al. 2004; Kalenscher and Pennartz 2008; Zeeb et al. 2010). Although the involvement of OFC in impulsive choice is indisputable, the exact role it plays is still unclear due to conflicting findings from several different labs (Kheramin et al. 2002; Mobini et al. 2002; Winstanley et al. 2004; Rudebeck et al. 2006; Winstanley 2007; Clark et al. 2008; Churchwell et al. 2009; Sellitto et al. 2010; Zeeb et al. 2010; Mar et al. 2011).

To add to the complexity of this story, more recent work has shown that different regions of the OFC serve opposing functions during delay discounting (Mar et al. 2011). In this study, Mar and colleagues showed that lesions of the mOFC made rats discount less, encouraging responding to the delayed reward after extended postlesion training (i.e., less impulsive), whereas IOFC lesions make rats discount more, decreasing preference for the delayed reward (i.e., more impulsive). Still, others have reported no impact of mOFC inactivation on delay discounting (Stopper et al. 2014).

Here, we show that, like IOFC, activity in mOFC reward-responsive neurons was attenuated for delayed reward. These results suggest that distinctive deficits observed after focal lesions to mOFC and IOFC cannot be explained by differences in firing during reward delivery. However, unlike IOFC, odor-responsive neurons in the mOFC signal low-value outcomes at the time of the decision. We suggest that mOFC's role in classic delay-discounting tasks is to signal low-value outcomes during decision making. Although these effects were significant, they were not dramatic, suggesting that other tasks are necessary to fully uncover the roles that mOFC plays in behavior, such as tasks that require more inhibitory-related functions (e.g., stop-signal) and those that require decisions made under risk (e.g., probability and uncertainty).

Funding

This work was supported by grants from the National Institute on Drug Abuse (R01DA031695, M.R.).

Notes

Conflict of Interest: None declared.

References

- Bechara A, Damasio H, Damasio AR. 2000. Emotion, decision making and the orbitofrontal cortex. *Cereb Cortex*. 10:295–307.
- Berlin HA, Rolls ET, Kischka U. 2004. Impulsivity, time perception, emotion and reinforcement sensitivity in patients with orbitofrontal cortex lesions. *Brain*. 127:1108–1126.
- Bouret S, Richmond BJ. 2010. Ventromedial and orbital prefrontal neurons differentially encode internally and externally driven motivational values in monkeys. *J Neurosci*. 30:8591–8601.
- Bryden DW, Johnson EE, Diao X, Roesch MR. 2011. Impact of expected value on neural activity in rat substantia nigra pars reticulata. *Eur J Neurosci*. 33:2308–2317.
- Bryden DW, Johnson EE, Tobia SC, Kashtelyan V, Roesch MR. 2011. Attention for learning signals in anterior cingulate cortex. *J Neurosci*. 31:18266–18274.
- Cardinal RN, Winstanley CA, Robbins TW, Everitt BJ. 2004. Limbic corticostriatal systems and delayed reinforcement. *Ann N Y Acad Sci*. 1021:33–50.
- Carmichael ST, Price JL. 1996. Connectional networks within the orbital and medial prefrontal cortex of macaque monkeys. *J Comp Neurol*. 371:179–207.
- Carmichael ST, Price JL. 1995a. Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *J Comp Neurol*. 363:615–641.
- Carmichael ST, Price JL. 1995b. Sensory and premotor connections of the orbital and medial prefrontal cortex of macaque monkeys. *J Comp Neurol*. 363:642–664.
- Churchwell JC, Morris AM, Heurtelou NM, Kesner RP. 2009. Interactions between the prefrontal cortex and amygdala during delay discounting and reversal. *Behav Neurosci*. 123:1185–1196.
- Clark L, Bechara A, Damasio H, Aitken MR, Sahakian BJ, Robbins TW. 2008. Differential effects of insular and ventromedial prefrontal cortex lesions on risky decision-making. *Brain*. 131:1311–1322.
- Damasio H, Grabowski T, Frank R, Galaburda AM, Damasio AR. 1994. The return of Phineas Gage: clues about the brain from the skull of a famous patient. *Science*. 264:1102–1105.
- Elliott R, Dolan RJ, Frith CD. 2000. Dissociable functions in the medial and lateral orbitofrontal cortex: evidence from human neuroimaging studies. *Cereb Cortex*. 10:308–317.
- Feierstein CE, Quirk MC, Uchida N, Sosulski DL, Mainen ZF. 2006. Representation of spatial goals in rat orbitofrontal cortex. *Neuron*. 51:495–507.
- Fiorillo CD, Newsome WT, Schultz W. 2008. The temporal precision of reward prediction in dopamine neurons. *Nat Neurosci*. 11:966–973.
- Gruber AJ, Calhoun GG, Shusterman I, Schoenbaum G, Roesch MR, O'Donnell P. 2010. More is less: a disinhibited prefrontal cortex impairs cognitive flexibility. *J Neurosci*. 30:17102–17110.
- Gutierrez R, Carmona JM, Nicolelis MA, Simon SA. 2006. Orbitofrontal ensemble activity monitors licking and distinguishes among natural rewards. *J Neurophysiol*. 95:119–133.
- Hoover WB, Vertes RP. 2011. Projections of the medial orbital and ventral orbital cortex in the rat. *J Comp Neurol*. 519:3766–3801.
- Iversen SD, Mishkin M. 1970. Perseverative interference in monkeys following selective lesions of the inferior prefrontal convexity. *Exp Brain Res*. 11:376–386.
- Kable JW, Glimcher PW. 2009. The neurobiology of decision: consensus and controversy. *Neuron*. 63:733–745.
- Kahnt T, Chang LJ, Park SQ, Heinzle J, Haynes JD. 2012. Connectivity-based parcellation of the human orbitofrontal cortex. *J Neurosci*. 32:6240–6250.
- Kalenscher T, Pennartz CM. 2008. Is a bird in the hand worth two in the future? The neuroeconomics of intertemporal decision-making. *Prog Neurobiol*. 84:284–315.
- Kennerley SW, Behrens TE, Wallis JD. 2011. Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nat Neurosci*. 14:1581–1589.
- 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.
- Kheramin S, Body S, Mobini S, Ho MY, Velazquez-Martinez DN, Bradshaw CM, Szabadi E, Deakin JF, Anderson IM. 2002. Effects of quinolinic acid-induced lesions of the orbital prefrontal cortex on inter-temporal choice: a quantitative analysis. *Psychopharmacology (Berl)*. 165:9–17.
- Kobayashi S, Schultz W. 2008. Influence of reward delays on responses of dopamine neurons. *J Neurosci*. 28:7837–7846.
- Kringelbach ML. 2005. The human orbitofrontal cortex: linking reward to hedonic experience. *Nat Rev Neurosci*. 6:691–702.
- Kringelbach ML, Rolls ET. 2004. The functional neuroanatomy of the human orbitofrontal cortex: evidence from neuroimaging and neuropsychology. *Prog Neurobiol*. 72:341–372.
- Mar AC, Walker AL, Theobald DE, Eagle DM, Robbins TW. 2011. Dissociable effects of lesions to orbitofrontal cortex subregions on impulsive choice in the rat. *J Neurosci*. 31:6398–6404.
- McClure SM, Ericson KM, Laibson DI, Loewenstein G, Cohen JD. 2007. Time discounting for primary rewards. *J Neurosci*. 27:5796–5804.
- McClure SM, Laibson DI, Loewenstein G, Cohen JD. 2004. Separate neural systems value immediate and delayed monetary rewards. *Science*. 306:503–507.
- Minamimoto T, Hori Y, Kimura M. 2005. Complementary process to response bias in the centromedian nucleus of the thalamus. *Science*. 308:1798–1801.
- Mobini S, Body S, Ho MY, Bradshaw CM, Szabadi E, Deakin JF, Anderson IM. 2002. Effects of lesions of the orbitofrontal cortex on sensitivity to delayed and probabilistic reinforcement. *Psychopharmacology (Berl)*. 160:290–298.
- Morecraft RJ, Geula C, Mesulam MM. 1992. Cytoarchitecture and neural afferents of orbitofrontal cortex in the brain of the monkey. *J Comp Neurol*. 323:341–358.
- Morrison SE, Saez A, Lau B, Salzman CD. 2011. Different time courses for learning-related changes in amygdala and orbitofrontal cortex. *Neuron*. 71:1127–1140.
- Morrison SE, Salzman CD. 2011. Representations of appetitive and aversive information in the primate orbitofrontal cortex. *Ann N Y Acad Sci*. 1239:59–70.
- Murray EA, O'Doherty JP, Schoenbaum G. 2007. What we know and do not know about the functions of the orbitofrontal cortex after 20 years of cross-species studies. *J Neurosci*. 27:8166–8169.
- 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 USA*. 107:20547–20552.
- O'Doherty J, Kringelbach ML, Rolls ET, Hornak J, Andrews C. 2001. Abstract reward and punishment representations in the human orbitofrontal cortex. *Nat Neurosci*. 4:95–102.
- Padoa-Schioppa C. 2011. Neurobiology of economic choice: a good-based model. *Annu Rev Neurosci*. 34:333–359.
- Padoa-Schioppa C, Assad JA. 2006. Neurons in the orbitofrontal cortex encode economic value. *Nature*. 441:223–226.
- Paxinos G, Watson C. 1997. *The Rat Brain in Stereotaxic Coordinates*, Compact 3rd ed. London: Academic Press. p 11–15.
- Price JL. 2007. Definition of the orbital cortex in relation to specific connections with limbic and visceral structures and other cortical regions. *Ann N Y Acad Sci*. 1121:54–71.
- Price JL, Carmichael ST, Drevets WC. 1996. Networks related to the orbital and medial prefrontal cortex: a substrate for emotional behavior? *Prog Brain Res*. 107:523–536.
- Roesch MR, Bryden DW. 2011. Impact of size and delay on neural activity in the rat limbic corticostriatal system. *Front Neurosci*. 5:130.

- Roesch MR, Bryden DW, Cerri DH, Haney ZR, Schoenbaum G. 2012. Willingness to wait and altered encoding of time-discounted reward in the orbitofrontal cortex with normal aging. *J Neurosci.* 32:5525–5533.
- Roesch MR, Calu DJ, Burke KA, Schoenbaum G. 2007. Should I stay or should I go? Transformation of time-discounted rewards in orbitofrontal cortex and associated brain circuits. *Ann N Y Acad Sci.* 1104:21–34.
- Roesch MR, Olson CR. 2005. Neuronal activity in primate orbitofrontal cortex reflects the value of time. *J Neurophysiol.* 94:2457–2471.
- Roesch MR, Olson CR. 2007. Neuronal activity related to anticipated reward in frontal cortex: does it represent value or reflect motivation? *Ann N Y Acad Sci.* 1121:431–446.
- Roesch MR, Olson CR. 2004. Neuronal activity related to reward value and motivation in primate frontal cortex. *Science.* 304:307–310.
- Roesch MR, Taylor AR, Schoenbaum G. 2006. Encoding of time-discounted rewards in orbitofrontal cortex is independent of value representation. *Neuron.* 51:509–520.
- Rudebeck PH, Murray EA. 2011a. Balkanizing the primate orbitofrontal cortex: distinct subregions for comparing and contrasting values. *Ann N Y Acad Sci.* 1239:1–13.
- Rudebeck PH, Murray EA. 2011b. Dissociable effects of subtotal lesions within the macaque orbital prefrontal cortex on reward-guided behavior. *J Neurosci.* 31:10569–10578.
- Rudebeck PH, Walton ME, Smyth AN, Bannerman DM, Rushworth MF. 2006. Separate neural pathways process different decision costs. *Nat Neurosci.* 9:1161–1168.
- Rygh R, Walker SC, Clarke HF, Robbins TW, Roberts AC. 2010. Differential contributions of the primate ventrolateral prefrontal and orbitofrontal cortex to serial reversal learning. *J Neurosci.* 30:14552–14559.
- Saleem KS, Kondo H, Price JL. 2008. Complementary circuits connecting the orbital and medial prefrontal networks with the temporal, insular, and opercular cortex in the macaque monkey. *J Comp Neurol.* 506:659–693.
- Schilman EA, Uylings HB, Galis-de Graaf Y, Joel D, Groenewegen HJ. 2008. The orbital cortex in rats topographically projects to central parts of the caudate-putamen complex. *Neurosci Lett.* 432:40–45.
- Schoenbaum G, Roesch M. 2005. Orbitofrontal cortex, associative learning, and expectancies. *Neuron.* 47:633–636.
- Schoenbaum G, Roesch MR, Stalnaker TA, Takahashi YK. 2009. A new perspective on the role of the orbitofrontal cortex in adaptive behaviour. *Nat Rev Neurosci.* 10:885–892.
- Sellitto M, Ciaramelli E, di Pellegrino G. 2010. Myopic discounting of future rewards after medial orbitofrontal damage in humans. *J Neurosci.* 30:16429–16436.
- Simmons JM, Ravel S, Shidara M, Richmond BJ. 2007. A comparison of reward-contingent neuronal activity in monkey orbitofrontal cortex and ventral striatum: guiding actions toward rewards. *Ann N Y Acad Sci.* 1121:376–394.
- St Onge JR, Floresco SB. 2010. Prefrontal cortical contribution to risk-based decision making. *Cereb Cortex.* 20:1816–1828.
- Stopper CM, Green EB, Floresco SB. 2014. Selective involvement by the medial orbitofrontal cortex in biasing risky, but not impulsive, choice. *Cereb Cortex.* 24:154–162.
- Takahashi YK, Roesch MR, Stalnaker TA, Haney RZ, Calu DJ, Taylor AR, Burke KA, Schoenbaum G. 2009. The orbitofrontal cortex and ventral tegmental area are necessary for learning from unexpected outcomes. *Neuron.* 62:269–280.
- Torregrossa MM, Quinn JJ, Taylor JR. 2008. Impulsivity, compulsivity, and habit: the role of orbitofrontal cortex revisited. *Biol Psychiatry.* 63:253–255.
- Tremblay L, Schultz W. 1999. Relative reward preference in primate orbitofrontal cortex. *Nature.* 398:704–708.
- van Duuren E, Escamez FA, Joosten RN, Visser R, Mulder AB, Pennartz CM. 2007. Neural coding of reward magnitude in the orbitofrontal cortex of the rat during a five-odor olfactory discrimination task. *Learn Mem.* 14:446–456.
- van Duuren E, van der Plasse G, Lankelma J, Joosten RN, Feenstra MG, Pennartz CM. 2009. Single-cell and population coding of expected reward probability in the orbitofrontal cortex of the rat. *J Neurosci.* 29:8965–8976.
- Vertes RP, Hoover WB, Rodriguez JJ. 2012. Projections of the central medial nucleus of the thalamus in the rat: node in cortical, striatal and limbic forebrain circuitry. *Neuroscience.* 219:120–136.
- Wallis JD. 2012. Cross-species studies of orbitofrontal cortex and value-based decision-making. *Nat Neurosci.* 15:13–19.
- Wallis JD. 2007. Orbitofrontal cortex and its contribution to decision-making. *Annu Rev Neurosci.* 30:31–56.
- 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.
- Winstanley CA. 2007. The orbitofrontal cortex, impulsivity, and addiction: probing orbitofrontal dysfunction at the neural, neurochemical, and molecular level. *Ann N Y Acad Sci.* 1121:639–655.
- Winstanley CA, Theobald DE, Cardinal RN, Robbins TW. 2004. Contrasting roles of basolateral amygdala and orbitofrontal cortex in impulsive choice. *J Neurosci.* 24:4718–4722.
- Zeeb FD, Floresco SB, Winstanley CA. 2010. Contributions of the orbitofrontal cortex to impulsive choice: interactions with basal levels of impulsivity, dopamine signalling, and reward-related cues. *Psychopharmacology (Berl).* 211:87–98.