Behavioral/Cognitive

Medial Orbitofrontal Neurons Preferentially Signal Cues Predicting Changes in Reward during Unblocking

Nina Lopatina,[1,3] Michael A. McDannald,[2] Clay V. Styer,[1] Jacob F. Peterson,[1] Brian F. Sadacca,[1] Joseph F. Cheer,[4] and Geoffrey Schoenbaum[1,4,5]

[1]National Institute on Drug Abuse, Intramural Research Program, Baltimore, Maryland 21224, [2]Department of Psychology, Boston College, Chestnut Hill, Massachusetts 02467, [3]University of Maryland School of Medicine, Program in Neuroscience, Baltimore, Maryland 21201, [4]University of Maryland School of Medicine, Department of Anatomy and Neurobiology, Baltimore, Maryland 21201, and [5]Johns Hopkins University, Department of Neuroscience, Baltimore, Maryland 21205

The orbitofrontal cortex (OFC) has been broadly implicated in the ability to use the current value of expected outcomes to guide behavior. Although value correlates have been prominently reported in lateral OFC, they are more often associated with more medial areas. Further, recent studies in primates have suggested a dissociation in which the lateral OFC is involved in credit assignment and representation of reward identity and more medial areas are critical to representing value. Previously, we used unblocking to test more specifically what information about outcomes is represented by OFC neurons in rats; consistent with the proposed dichotomy between the lateral and medial OFC, we found relatively little linear value coding in the lateral OFC (Lopatina et al., 2015). Here we have repeated this experiment, recording in the medial OFC, to test whether such value signals might be found there. Neurons were recorded in an unblocking task as rats learned about cues that signaled either more, less, or the same amount of reward. We found that medial OFC neurons acquired responses to these cues; however, these responses did not signal different reward values across cues. Surprisingly, we found that cells developed responses to cues predicting a change, particularly a decrease, in reward value. This is consistent with a special role for medial OFC in representing current value to support devaluation/revaluation sensitive changes in behavior.

Key words: orbitofrontal; Pavlovian; rat; single unit

Significance Statement

This study uniquely examines encoding in rodent mOFC at the single-unit level in response to cues that predict more, less, or no change in reward in rats during training in a Pavlovian unblocking task, finding more cells responding to change-predictive cues and stronger activity in response to cues predictive of less reward.

Introduction

The orbitofrontal cortex (OFC) has been broadly implicated in the ability to use the current value of expected outcomes to guide behavior (Schoenbaum et al., 1998; Delamater, 2007; Ostlund and Balleine, 2007; Steiner and Redish, 2012; Luk and Wallis, 2013). This requires the ability to predict what specific outcomes are available in the environment and the ability to use that infor-

mation to retrieve those outcomes' current values to guide behavior (Pickens et al., 2003; West et al., 2011; Jones et al., 2012). Although general or abstract value correlates have been prominently reported in lateral OFC (lOFC) (Padoa-Schioppa and Assad, 2006, 2008; Padoa-Schioppa, 2009, 2013), they are arguably more often identified in medial OFC (mOFC) (Plassmann et al., 2007, 2010; Hare et al., 2008, 2010; Levy and Glimcher, 2011; McNamee et al., 2013; Strait et al., 2014). Further, recent studies in primates have suggested a dissociation in which the lOFC is involved in credit assignment and representation of reward identity and more medial areas are critical to representing value (Noonan et al., 2010, 2012). This suggests a model in which subregions of the OFC interact in the above two stage process, with lOFC tracking the associative model of the task, representing specific associations between cues and specific outcomes, and mOFC accessing those representations to retrieve a current value.

We have recently used an unblocking to test more specifically what information about outcomes is represented by OFC neurons. Across two experiments, we recorded neurons in the lOFC

in rats learning about cues that signaled changes in an expected outcome. In the first study (McDannald et al., 2014), we compared firing to cues that signaled either a shift in reward flavor or shift in reward value. To our surprise, we found similar representation of these two cues, even though one reflected a valueless change in reward features. This result was more consistent with a feature representation than with a representation tied to value. To explore this further, we conducted a second study (Lopatina et al., 2015), in which we compared firing in lOFC neurons to cues that signaled more reward, less reward, or no change in reward. Despite the linear change in value across cues, a change reflected in the rats' behavior, we failed to find neural correlates that reflected reward value across cues. Instead, we found dissociable populations of lOFC neurons that developed firing to each of the three cues, including the cue that predicted no change in reward. Again, this result seemed more consistent with a representation of outcome features than with a representation tied to any abstract or outcome-independent value.

Here, we have repeated this second experiment, recording in the mOFC, to test whether abstract value representations might be found there. mOFC neurons were recorded in an unblocking task as rats learned about cues that signaled more, less, or the same amount of reward. mOFC neurons acquired responses to these cues, albeit in significantly lower proportions than in lOFC. However, again, these responses did not correlate with abstract value across cues. Instead, most associative neurons developed responses to a single cue, most strongly to the cue that predicted a decrease in reward. These results are not consistent with simple models in which mOFC represents either general value or outcome features. Instead, they suggest that mOFC is signaling information (emotional reactions, error signals, or strategy shifting) elicited by a change, particularly a reduction, in the value of an expected reward. This would be consistent with a special role for mOFC in representing current value to support devaluation/revaluation-sensitive changes in behavior.

Materials and Methods

Subjects. Thirteen Male Long–Evans rats were obtained at 200–250 g from Charles River Laboratories. Rats were tested at the Intramural Research Program at the National Institute on Drug Abuse in accordance with National Institutes of Health guidelines.

Surgery and histology. Using aseptic, stereotaxic surgical methods, a drivable bundle of 16 25-μm-diameter FeNiCr wires (Stablohm 675, California Fine Wire) was chronically implanted in the left hemisphere at medial OFC at a 13° angle, 4.4 mm anterior to bregma, 1.58 mm lateral, and 2.78 mm ventral to each rat's brain surface. We implanted these microelectrodes in medial orbital cortex because this area has recently been reported to be homologous to medial orbital areas in primates based on a comparison of connectivity with striatum, hippocampus, and amygdala (Heilbronner et al., 2016). These wires were cut at an angle with surgical scissors immediately before implantation, to extend ∼1.8– 2.5 mm beyond the cannula, with a range of ∼0.3 mm between wires. Current was passed through each electrode immediately before implantation to lower the impedance to ∼300–400 kOhms. Rats were anesthetized with isoflurane. Subcutaneous injections of 0.1 ml lidocaine and 0.1 ml carprofen diluted in saline were used for analgesia. At the study's conclusion, a 15 μA current was passed through each electrode to mark the final position. Following perfusion of the rats, their brains were extracted and processed for histology using standard techniques.

Blocking task. Recording was conducted in grounded aluminum chambers ∼18 in. on each side with sloping walls narrowing to an area of 12 in. × 12 in. at the bottom. An odor port was located centrally above a fluid well on a panel in the right wall of each chamber. Above the panel were two lights. To allow rapid delivery of olfactory cues to the odor port, it was connected to an airflow dilution olfactometer. Odors were chosen

from compounds obtained from International Flavors and Fragrances. The fluid well was connected to lines controlling the independent delivery of liquid rewards. A computer running a behavioral program written in C2+ implemented control of the task. Following implantation with microelectrodes, rats were water deprived by restricting access to 10 min daily. Following 2 d of water deprivation, rats were shaped, in stages, to hold in the odor port for 1 s to receive a water reward at the well. Each trial started with house light illumination, following which rats had 3 s to enter the odor port. A failure to enter the odor port caused restart of the trial. Rats were required to hold for 1 s in the odor port, and upon exit had 3 s to enter the reward well. Again, failure to hold for 1 s or to make reward well entry within 3 s resulted in restart of the trial. Following shaping, rats were trained until they proficiently responded for the initial odor to receive a medium-sized bolus of diluted chocolate milk solution; this comprised up to 15 sessions, with a maximum of 170 trials per session. Completion of ∼150 trials per session was characterized as proficient responding.

Once rats were deemed proficient at initial training and single units were isolated, the unblocking procedure began. On each of the two learning days, rats received four trial types. The first trial type was a reminder of initial training. The remaining trial types comprised a 200 ms presentation of the initial odor followed by one of three 800 ms, novel, differentiable odors: one signaling the same medium-sized bolus of chocolate milk used in prior training, a second signaling a larger bolus, and a third signaling a smaller bolus. The behavioral requirements for each of trial type were exactly as in initial training. Rats completed 20–40 trials with each novel odor per session during unblocking. Then, on the probe test day, rats received 10 reminder trials of each type, followed by up to 10 trials of each novel odor alone without reward, interleaved with rewarded presentations of the initial odor to maintain responding. During the unrewarded, novel-odor extinction trials, both requirements to sample the odor for 1 s and respond to the reward well were lifted. This unblocking procedure was repeated two or three times per rat, using a new set of blocked, upshift, and downshift odors each time.

Single-unit recording. Neural activity was recorded using four identical Plexon Multichannel Acquisition Processor systems interfaced with odor discrimination training chambers described above. Following recovery from surgery and proficiency in shaping, electrodes were advanced daily until activity was obtained. Rats received reminder training using the pretrained initial odor, as described above, during this process. Once rats showed proficient responding and single units were isolated, the rat began unblocking. During this 3 d procedure, the electrode was moved ∼167 μm between the first and second learning days and again following each 3 d unblocking procedure in all rats. This was done between unblocking days and before repetition of this process in new odor cues to acquire neurons in a new location in OFC.

Statistical data analysis. Units were sorted using Offline Sorter software from Plexon using a k-means algorithm. Sorted files were next processed in Neuroexplorer to extract relevant event markers and unit timestamps. These data were then analyzed in MATLAB (The MathWorks). To analyze activity in response to the novel odors, we examined activity between 300 and 1300 ms subsequent to initial odor onset, which corresponded approximately with the novel odor delivery to the odor port. The intertrial interval (ITI) was defined as the 2 s preceding illumination of the house light. Normalized firing was calculated by subtracting firing rate during the ITI from the period of interest: Normalized firing = (Period spikes/s) − (ITI spikes/s). Odor-responsive neurons were identified as units that showed an increase in firing from baseline during odor sampling (t tests, $p < 0.01$) on at least one of the four trial types. Odor period analyses and figures show data beginning with the eighth trial, excluding heat plots over trials and classification accuracy over trials. Analyses began on the eighth trial to average a 10 and 5 trial cutoff for putative sensory firing on unblocking days 1 and 2, respectively.

Neurons were classified as putative sensory neurons and excluded from further consideration if they significantly increased firing to all odors and were not value coding. Neurons that were not eliminated by this screen were classified as upshift-responsive if they increased firing significantly to the upshift cue. Alternatively, units were classified as downshift-responsive or blocked-responsive if they increased firing to
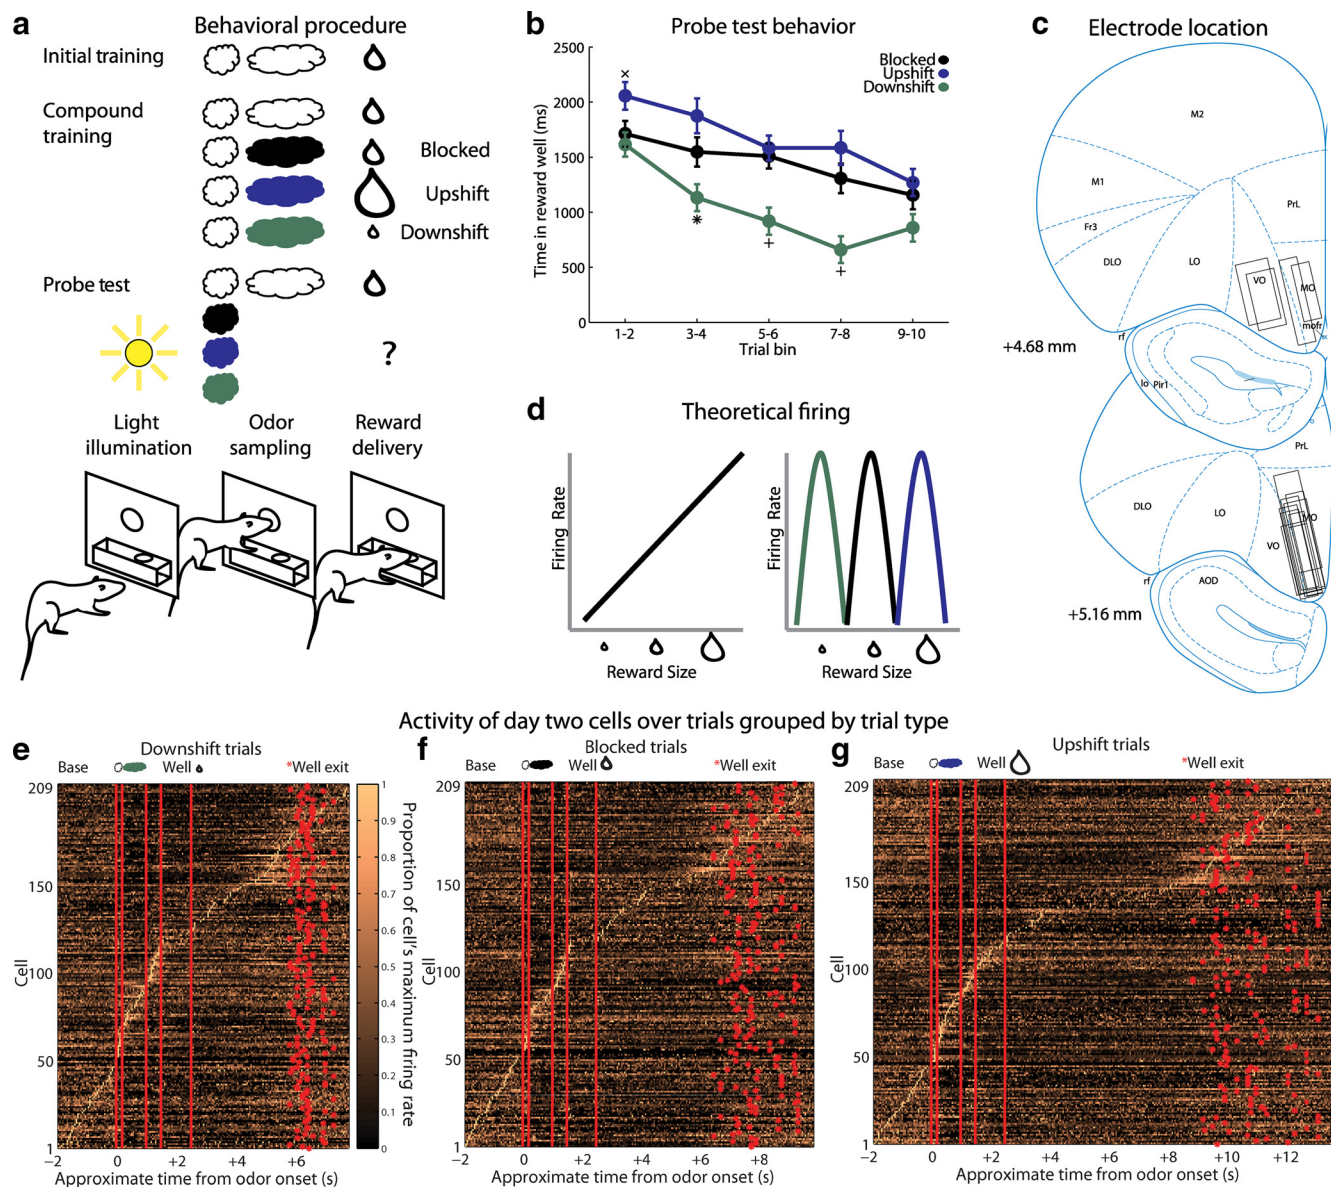
Figure 1. Experimental outline, behavior summary, recording sites, and recorded cells. ***a***, Thirsty rats were initially trained to enter an odor port after a house light lit up, then to go to the reward well below to receive a drop of chocolate milk. There were four trial types in the unblocking session. The first was a reminder of initial training. On the other three trial types, the originally trained odor was briefly presented, followed by one of three novel odors. The reward following the novel odors was unchanged (black; blocked trials), larger in size (blue; upshift trials), or smaller in size (green; downshift trials). In the probe test stage, we assessed learning by presenting the novel odors without a subsequent reward. ***b***, Time in the reward well on the probe test trials. ANOVA for time spent in the reward well with odor (blocked, upshift, downshift), and trial (1–10) as factors found a significant effect of odor (ANOVA, $F_{(2,66)} = 18.88, p = 3.27 \times 10^{-7}$) and trial (ANOVA, $F_{(9,297)} = 9.94, p = 2.43 \times 10^{-13}$). Planned comparisons confirmed that, in the first two trial block, rats spent significantly more time in the reward well following the upshift odor ($p = 0.0078$) relative to the blocked odor. Rats also spent less time in the reward well following the downshift odor on Trials 3– 8 relative to the blocked odor ($p = 0.0306, p = 0.0009$, and $p = 0.0003$, respectively). $*p < 0.05$. $\times p < 0.01$. $+p < 0.001$. Error bars indicate SEM. ***c***, Single-unit activity was recorded from the medial orbital cortex with some recordings in or overlapping ventral orbital cortex. Locations are shown at 4.68 and 5.16 mm anterior to bregma. MO, Medial orbital cortex; VO, ventral orbital cortex. ***d***, Theoretical firing pattern for cells that signal value monotonically (left) or are tuned to individual reward sizes (right). ***e–g***, Activity of day 2 cells with a baseline firing rate $<$10 Hz sorted by when each cell reaches its maximal firing rate over the course of a trial on (***e***) downshift, (***f***) blocked, and (***g***) upshift trials. Cells are sorted independently in the order of earliest maximal firing rate within each trial type; thus, the cell numbers do not correspond between trial types.

the downshift or blocked cues, respectively. Single-unit and population activity was plotted in 50 ms bins. Population activity was additionally analyzed with repeated-measures ANOVA with bin (50 ms) and odor trial type (initial, blocked, upshift, and downshift) as factors. Single-unit and population firing were smoothed by taking a four bin average moving in 50 ms increments moving away from the novel odor period.

Heat plots were constructed in 150 ms sliding windows moving away from the novel odor onset in 50 ms increments. Warmer colors (dark red) indicated positive difference scores, whereas cooler colors (dark blue) indicated negative difference scores. This was done for the first 10

trials of the identified population. Significance of differential firing to identified odors was determined by performing a one-tailed t test comparing differential firing to zero in the exact same 150 ms sliding windows for each of the 10 trials.

Classification accuracy was calculated using a linear classifier on instantaneous firing rate measured in 100 ms bins. The goal of the classification was to predict the trial type. Classification accuracy by trial was performed using a sliding window. Accuracy was calculated during the 300 ms window following novel odor presentation. Bins were counted as additional trials to offset low trial number. Statistical significance was calculated based on a binomial distribution (Combrisson and Jerbi,
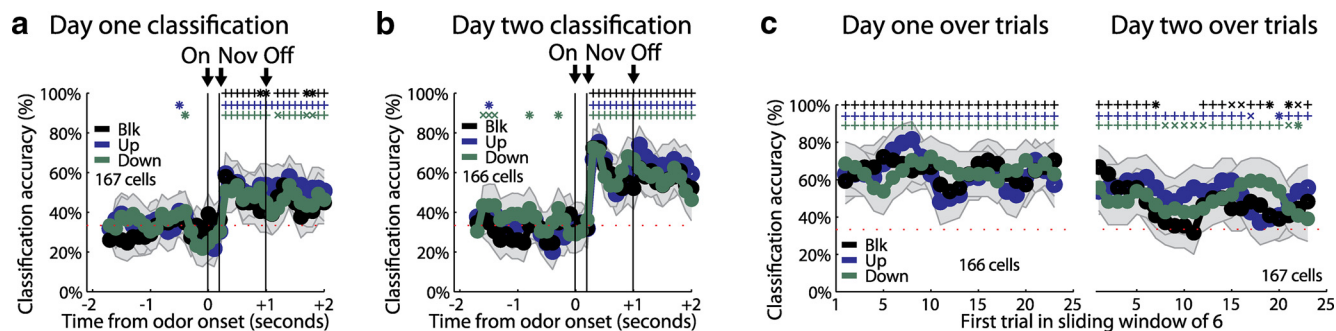
Figure 2. Recorded units not restricted to odor-responsive units. *a*, *b*, Classification accuracy for all cells with firing <10 Hz for day 1 cells (*a*) and day 2 cells (*b*). Dashed red line indicates chance. *c*, Classification accuracy over a sliding window of 6 trials for day 1 (left) and day 2 (right) cells. Classification accuracy significance above chance is indicated above time bins in a color matching the trial type. $*p < 0.05$. $^{\times}p < 0.01$. $^{+}p < 0.001$. Error bars indicate SEM.

2015). Classification accuracy during the odor period was smoothed by taking a three bin average moving in 100 ms increments away from the novel odor period. Classification accuracy over trials plots were smoothed by taking a three trial average.

Results

We recorded single-unit activity in the mOFC in 13 rats during an odor-based unblocking task (Fig. 1*a*). After implantation of microelectrodes in the mOFC, rats were trained to sample an odor in a central port following house light illumination and then respond to a reward well below for a single medium-sized drop of chocolate milk. This training was extensive, lasting for at least 4 d, and was meant to establish the initial odor as a reliable predictor of this specific outcome. Each rat then underwent 2 or 3 rounds of unblocking.

Each round of unblocking began with 2 d of training and consisted of four trial types (Fig. 1*a*, compound training). One type was a reminder: the initially trained odor was followed by the expected outcome. On the other three trial types (upshift, downshift, blocked), rats were presented with the initially trained odor, followed immediately by one of three novel odors. On blocked trials, the novel odor was followed by the expected medium-sized drop of milk, whereas on upshift and downshift trials, the novel odor was followed by a noticeably larger or smaller drop of milk, respectively. As a result, the novel odor added on blocked trials predicted no change in reward value or features, whereas the novel odors added on the other two trial types predicted explicit changes in reward value. Further, because the reward identity remained the same, as did the number, location, and timing of delivery, any associative information that accrues to these novel odors should be primarily driven by the changes in value and possibly with the feature of "larger" or "smaller." Within the constraints of this design, we reasoned that, if neurons responding to the odor cues were indeed signaling general or abstract value, then they should exhibit rank ordered differences in firing to the upshifted, blocked, and downshifted odor cues (Fig. 1*d*, left). On the other hand, if cue-responsive neurons were signaling unique attributes of each outcome (e.g., larger or smaller), then we might expect them to be tuned to a particular cue (Fig. 1*d*, right).

Rats learned to respond differently to the upshifted and downshifted cues

In the unblocking sessions, rats were sensitive to presentation of the novel odors, exhibiting longer latencies to respond at the reward well following odor sampling on these three trial types. Longer latencies to the novel odors were most apparent on the very first trial of each session, particularly on day 1 (data not shown; ANOVA revealed a main effect of trial ($F_{(19,608)} = 7.6, p < 1 \times 10^{-4}$), cue ($F_{(2,64)} = 9.27, p = 2.91 \times 10^{-4}$), and a trial × day interaction ($F_{(19,608)} = 2.17, p = 0.0028$)). In addition to this effect, the rats also learned that two of these odors predicted meaningful changes in the outcome. This was evident in the extinction probe test in which they initially spent more time in the fluid well following sampling of the upshift odor and less time following sampling of the downshift odor, versus the blocked odor, as if expecting more and less reward respectively on upshift and downshift trials (Fig. 1*b*).

Upshift-responsive neurons in the mOFC do not signal general value and are not selective for the upshift cue

We recorded 188 single units during the first day of unblocking and 212 units on the second unblocking day in 34 rounds of training across all 13 rats (Fig. 1*c*). Characteristically, cells in these populations fired to all of the events that characterized the trials in our task, including prominently during odor sampling (Fig. 1*e*–*g*). The proportion of odor-responsive cells was approximately half that in the lateral OFC: 54 of 400 (13.5%) compared with 172 of 680 (25.3%). This difference is significant by χ^2 test ($\chi^2 = 21.17, p < 0.001$). To address our hypothesis, units with a baseline firing rate <10 Hz were screened for phasic responses to one of the four odors using a t test comparing firing rates during the ITI and novel odor period (significance level = $p < 0.01$). We used a 10 Hz high pass cutoff to exclude fast spiking cells as we had in our previous study in lOFC. This screen excluded 4 cells, none of which showed value coding or any other meaningful coding. Our screen for phasic excitatory responses over ITI targeted odor responses we had previously found were associative rather than sensory. Similar analyses for inhibitory responses did not find significant decreases from baseline consistent with associative signaling (data not shown). Excluding nonassociative neurons, which fired from the start of training to all of the odors, our screen for phasic responses identified 44 units (day 1 = 23, day 2 = 21) that showed a significant increase in firing to at least one of the odor cues but not all four odors similarly (for activity of all units with a baseline firing rate <10 Hz regardless of odor response, see Fig. 2).

We began by asking whether upshift-responsive neurons were responsive to general value or whether they were signaling features, focusing our initial analysis on the neurons (21 of 44) that showed a significant phasic response to the upshifted odor. As in lOFC, this population response appeared to grow stronger with learning (Fig. 3*f*,*g*), demonstrating that it was not driven strictly

by physical properties of the odors. There is no significant firing above baseline at the beginning of the first trial, and the number of bins with significant firing above baseline is greater on subsequent trials than the first trial. Although there were insufficient cells in this subpopulation to permit a statistical verification of this visual impression, we did find a statistically reliable effect when we combined the upshift-responsive subpopulation with the downshift- and blocked-responsive subpopulations in Figures 4 and 5. ANOVA with factors trial, bin, and subpopulation found a main effect of trial ($F_{(9,441)} = 1.92$, $p = 0.0477$) and interaction between trial and bin ($F_{(135,6615)} = 1.479$, $p = 3.028 \times 10^{-4}$). There was no main effect of subpopulation ($F_{(2,49)} = 0.847$, $p = 0.435$) and no interaction between trial and subpopulation ($F_{(18,441)} = 0.738$, $p = 0.772$).

However, unlike upshift-responsive cells in the lOFC, which fired most strongly to the upshifted cue as a group (Fig. 6a, left), these cells fired as strongly to the other two novel cues as to the upshift cue overall: ANOVA restricted to only the three novel odors found no effect of odor (Fig. 3d). Further examination of individual units showed that these neurons exhibited a variety of patterns of activity to the other odor cues (Fig. 3a–c). Thus, firing across the upshift-responsive population did not appear to provide the linear value correlate predicted by a value-coding hypothesis (Fig. 1d, left). Of

←

colors as in raster plots. Shaded areas represent SEM. ANOVA with bin (50 ms) and odor as factors found significant effects of bin ($F_{(59,1180)} = 6.18$, $p < 1 \times 10^{-11}$), odor ($F_{(3,60)} = 3.31$, $p = 0.026$), and a bin × odor interaction ($F_{(177,3540)} = 1.88$, $p = 8.12 \times 10^{-11}$). ANOVA restricted to the novel odor period with odor and time (first 500 ms vs second 500 ms; top right, inset) as factors found a main effect of odor ($F_{(3,60)} = 3.20$, $p = 0.0296$). ANOVA restricted to only the three novel odors found no effect of odor ($F_{(2,40)} = 0.79$, $p = 0.46$). *e*, Scatter plot of blocked-cue and baseline-normalized upshift and downshift firing rate for all cells in the upshift-responsive population found a nonsignificant positive correlation between upshift and downshift firing. *f*, Baseline and initial cue-normalized firing to the upshift cue for the upshift-responsive population on the first 10 trials over both unblocking day ($n = 21$). Firing was calculated in a 150 ms sliding window for each 50 ms bin moving away from novel odor onset. The firing rate over ITI was then plotted, with dark red bins indicating maximal firing to the upshift cue and blue indicating minimal firing. Heat plot values shown on right of heat plot. *g*, The significance of the increased firing to the upshift cue was determined by performing a one-tailed *t* test comparing firing rate to 0, using a significance of $p < 0.01$ and a sliding window as in *f*. Black bins represent significant elevations in firing to the upshift cue.
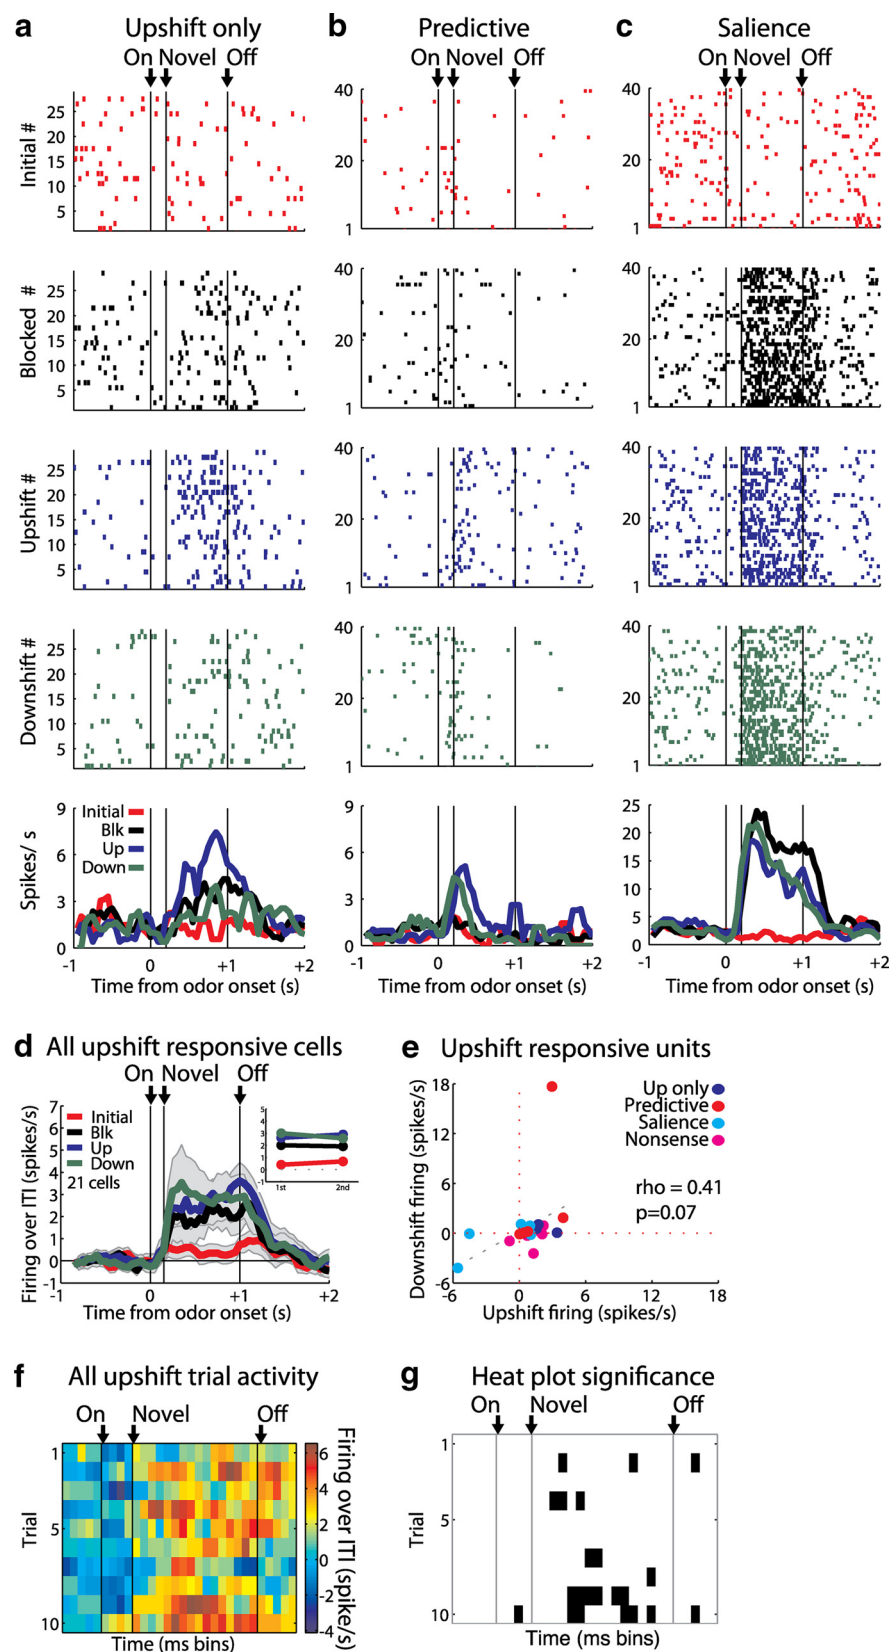
Figure 3. Single-unit and population firing of upshift-responsive neurons. *a–c*, Raster plots for firing of single units on initial (red), blocked (black), upshift (blue), and downshift (green) trials. First vertical line indicates odor onset (On). Second vertical line indicates novel odor onset (Novel). Third vertical line indicates odor offset (Off). Each tick represents a spike. Average activity across all trials is plotted by odor (bottom) showing the following: (*a*) selective firing to the upshift odor on day 2, (*b*) putative predictive firing to the upshift and downshift odors on day 2, (*c*) putative salience firing to all three novel odors but not the initially trained odor on day 1, and (*d*) mean neural activity (novel odor epoch, ITI) for the upshift-responsive population ($n = 21$) is plotted. Line
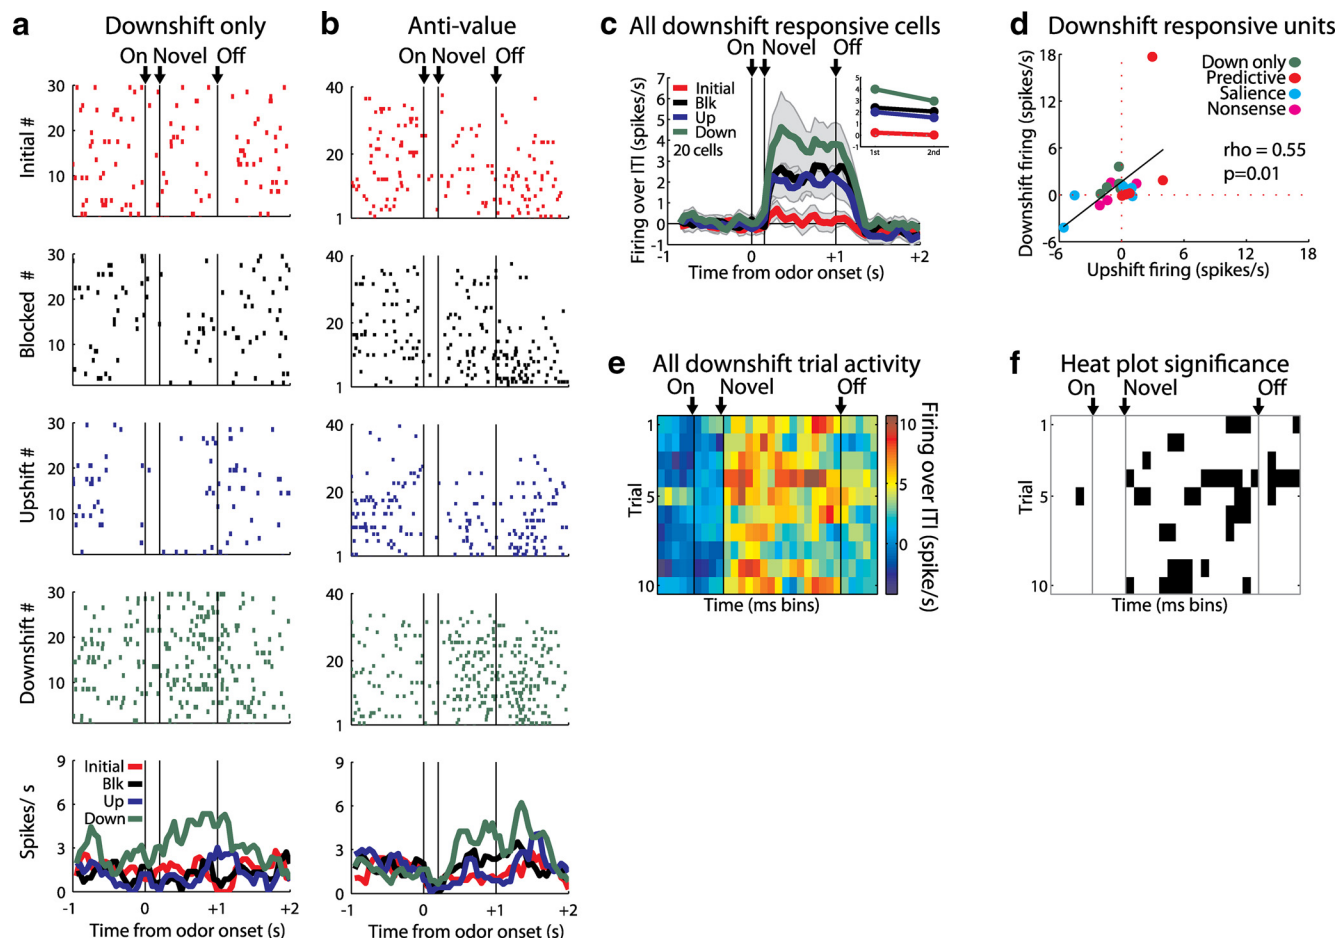
Figure 4. Single-unit and population firing of downshift-responsive neurons. *a*, *b*, Raster plots as in Figure 3*a*–*c* for day 1 cells showing (*a*) selective firing to the downshift cue and (*b*) putative anti-value firing, exhibiting a monotonic firing pattern negatively correlated with reward value. *c*, Mean neural activity (novel odor epoch, ITI) as in Figure 3*d* for the downshift-responsive population ($n = 20$). ANOVA with bin and odor as factors found significant effects of bin ($F_{(59,1121)} = 9.44$, $p < 1 \times 10^{-11}$), odor ($F_{(3,57)} = 6.75$, $p = 5.637 \times 10^{-4}$), and the bin $\times$ odor interaction ($F_{(177,3363)} = 2.49$, $p < 1 \times 10^{-11}$). ANOVA restricted to the novel odor period with odor and time (first 500 ms vs second 500 ms; top right, inset) as factors found a main effect of odor ($F_{(3,57)} = 6.30$, $p = 9.2 \times 10^{-4}$). ANOVA restricted to only the three novel odors found a main effect of odor ($F_{(2,38)} = 3.25$, $p = 0.0498$). *d*, Scatter plot of blocked-cue and baseline-normalized upshift and downshift firing rate for all cells in the downshift-responsive population as in Figure 3*e*. *e*, *f*, Heat plots and *p* value plots as in Figure 3*f*, *g* for all downshift-responsive cells on unblocking days 1 and 2 ($n = 20$).

course, the heterogeneity of response types could obscure a simple value correlate in a subpopulation, so next we asked whether any of the 21 upshift-responsive neurons exhibited such a firing pattern. Value coding was operationally defined as an excitatory response to at least one odor and significantly more activity to the upshift cue than the blocked cue as well as less activity to the downshift cue than the blocked cue. We found no neurons exhibiting firing that clearly rank ordered the cues according to their value (0 of 21; 0%). We found several neurons that showed isolated firing to the upshifted odor alone (as in Fig. 3*a*; 5 of 21; 24%). We also found an equivalent number of upshift-responsive neurons that were equally responsive to the other two novel odor cues (as in Fig. 3*c*; 5 of 21; 24%), replicating the finding of salience or novelty encoding in our prior study. Additionally, we saw an equivalently sized population of neurons that fired similarly to the upshifted and downshifted cues and did not respond to the blocked cue (as in Fig. 3*b*; 5 of 21; 24%). This population was reminiscent of the neurons in our prior study that fired similarly to the two cues predicting changes in the outcome. These three groups were the largest subtypes. Notably, none of these populations showed a relationship between firing to the upshift and downshift cues consistent with value coding across

them (Fig. 3*e*). Our results here largely replicated what we observed in the lateral OFC for this specific population, with one notable exception: these upshift-responsive cells were not particularly selective for the upshift cue, even though that was how we isolated them.

Cells in the mOFC respond strongly to the downshift odor
Given the unexpected result of equal firing to all three novel cues in the upshift-responsive population, we next identified all of the cells that responded to the downshifted cue. There were 20 cells that fit this profile (Fig. 4*c*), including 10 that also responded to the upshift cue. As in the upshift-responsive population, the response of these cells appeared to grow stronger over time, with the population exhibiting longer and earlier periods of significantly elevated firing across the first 10 trials of unblocking (Fig. 4*e*,*f*). Interestingly, the cells in this group responded most strongly to the downshift cue versus the other two novel cues (upshift and blocked), although only one unit exhibited an "anti-value" firing pattern (Fig. 4*b*; 1 of 20, 5%). Whereas ANOVA restricted to only the three novel odors found no effect of odor in the upshift-responsive group, this effect was found in the downshift responsive population

$(F_{(2,38)} = 3.25, p = 0.0498)$. There were 6 cells that responded only to the downshift cue (as in Fig. 4a; 6 of 20, 30%), the most prevalent correlate in this population. The response pattern of this population is also not consistent with a simple relationship between reward value and firing rate.

mOFC neurons respond least to the blocked odor cue

We next identified cells that responded to the blocked cue. We identified 11 cells that fired to this cue (Fig. 5b), including 7 that responded to the upshift cue and 7 that responded to the downshift cue. As in the upshift- and downshift-responsive populations, the response of these cells appeared to grow stronger over time, with the population exhibiting longer and earlier periods of significantly elevated firing across the first 10 trials of the unblocking (Fig. 5d,e). However, the cells in this group responded similarly to all three novel cues (Fig. 5b) as the upshift-responsive population had; and only two responded solely to the blocked cue (Fig. 5a). The neurons that responded uniquely to the blocked cue in mOFC comprised the smallest subgroup within this population, and this number ($n = 2$; Fig. 6b, center) was a smaller proportion than the number that responded uniquely to either the upshift ($\chi^2 = 5.345, p < 0.05$) or the downshift cues ($\chi^2 = 4.243, p < 0.05$).

mOFC response to size shifting cues differs from lOFC response

Overall, the above suggests that there are significant differences in how the lOFC and mOFC represent associative information. To summarize these differences more clearly, we plotted the different populations and their respective overlap in Figure 6. As this figure illustrates, the mOFC response to the novel odor cues differs in several ways from the lOFC response in our previous study. For starters, there were simply fewer neurons responsive to the novel odor cues. This is noted at the outset of Results and is also evident in the smaller numbers on the diagram, despite approximately similar numbers of recorded neurons. In addition, whereas the three novel cues were represented by similar proportions of cells in lOFC (Fig. 6a, similar size circles), neurons in the mOFC were primarily associated with cues predicting a change in reward, underrepresenting the blocked cue (Fig. 6b, middle). Indeed, the number of neurons that responded uniquely to the blocked cue in mOFC ($n = 2$; Fig. 6b) was a smaller proportion than the number that responded uniquely to either the upshift or the downshift
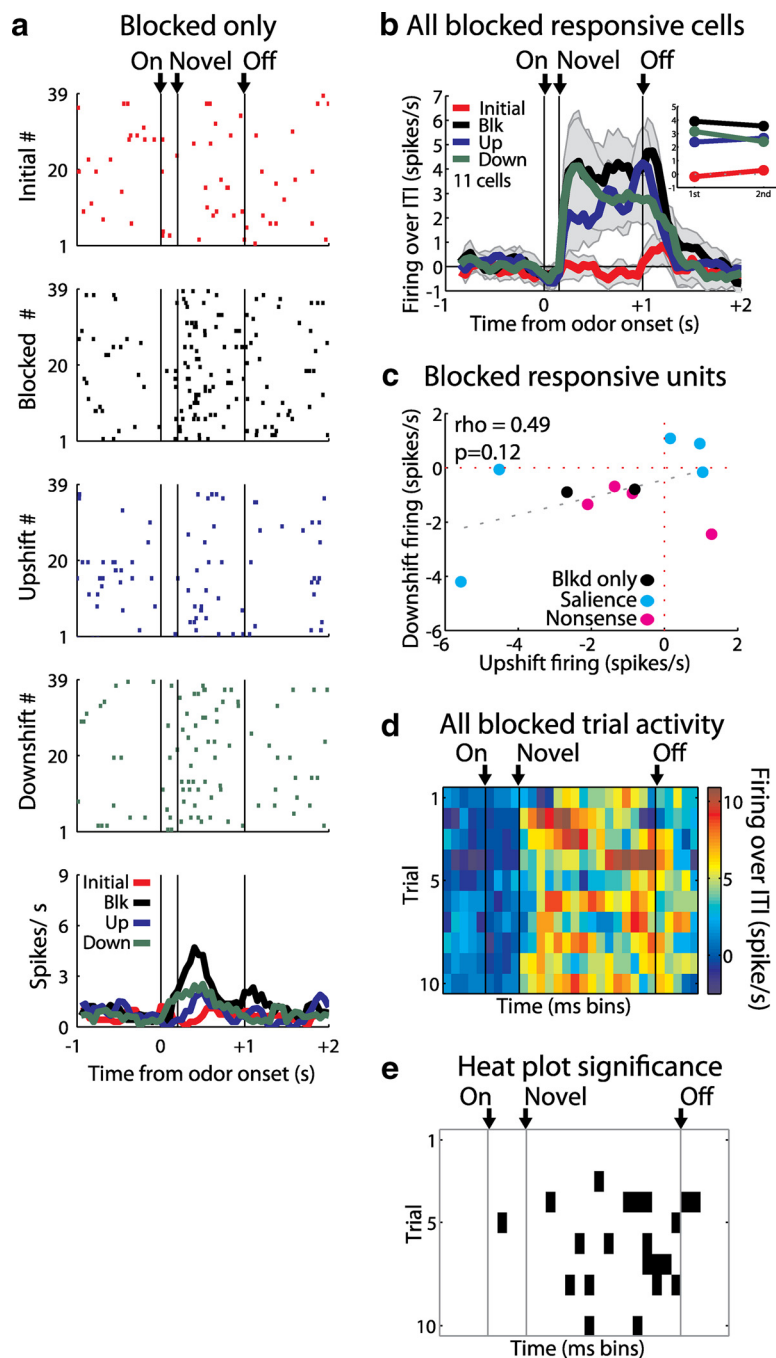

Figure 5. Single-unit and population firing of blocked-responsive neurons. **a**, Raster plots as in Figure 3a–c for firing of single units that are selective for the blocked cue. **b**, Mean neural activity (novel odor epoch, ITI) as in Figure 3d for the blocked-responsive population ($n = 11$). ANOVA with bin and odor as factors found significant effects of bin ($F_{(59,590)} = 4.49, p < 1 \times 10^{-11}$), odor ($F_{(3,30)} = 8.07, p = 4.36 \times 10^{-4}$), and the bin × odor interaction ($F_{(177,1770)} = 2.53, p < 1 \times 10^{-11}$). ANOVA restricted to the novel odor period with odor and time (first 500 ms vs second 500 ms, shown in top right, inset) as factors found a main effect of odor ($F_{(3,30)} = 6.57, p < 0.01$). ANOVA restricted to only the three novel odors found no effect of odor ($F_{(2,20)} = 2.24, p = 0.1325$). **c**, Scatter plot of blocked-cue and baseline-normalized upshift and downshift firing rate for all cells in the blocked-responsive population as in Figure 3e. **d, e**, Heat plots and p value plots as in Figure 3f, g for all blocked-responsive cells on unblocking days 1 and 2 ($n = 11$).

cues. This was not the case in lOFC, where the blocked cue received similar representation to the other populations (Fig. 6a). Last, we compared the two populations' value encoding with a regression across all three reward amounts at $p < 0.01$ without any prior or subsequent screening. We found 10 mOFC cells that fit the regression: 3 positive and 9 negative (2.5% of all recorded
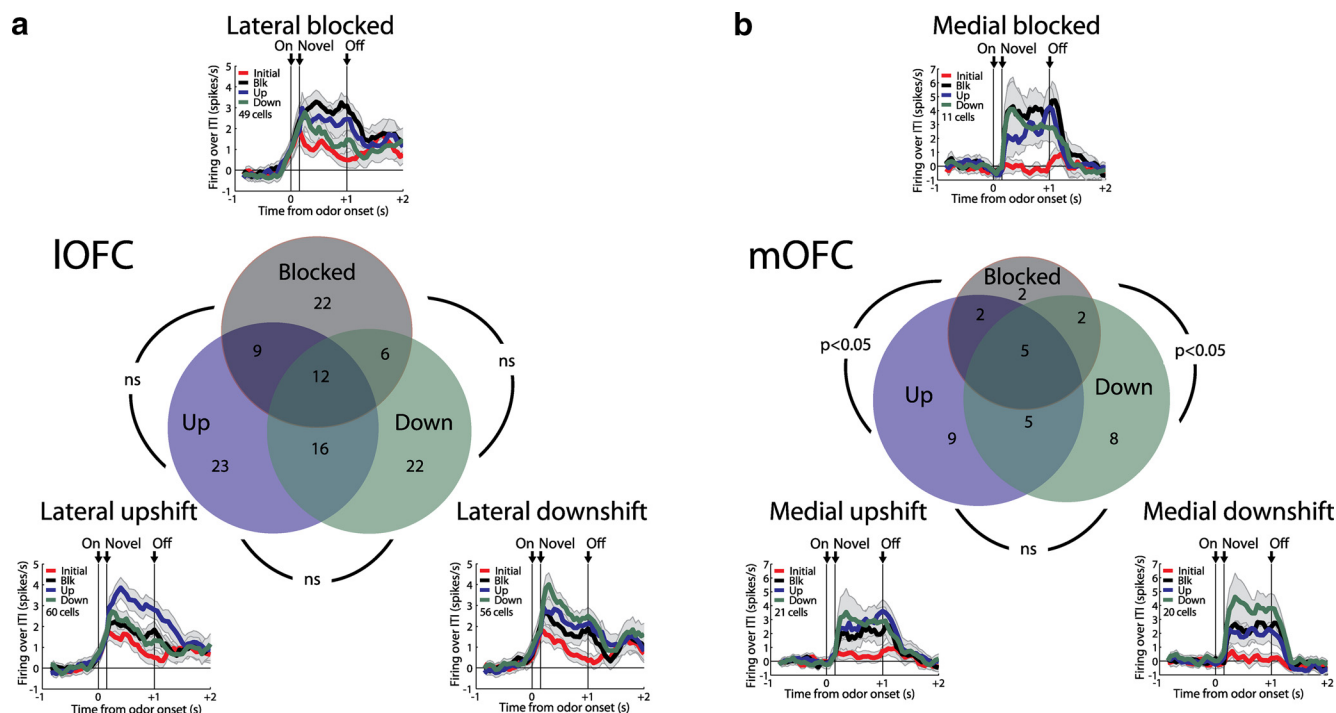
Figure 6. lOFC and mOFC comparison. **a**, Venn diagram of lOFC response types. Each circle is proportionate to the number of cells that respond to each respective cue. Overlap between circles indicates cells that respond to multiple cues. Numbers represent the number of cells that respond to each cue or combination of cues. Curved lines indicate results of comparison of unique numbers in each group by χ^2. Inset, Population plots show mean neural activity (novel odor epoch, ITI) as in Figure 3d for populations in the lOFC that were as follows: Left, Upshift-responsive (n = 60). Right, Downshift-responsive (n = 56). Top, Blocked-responsive (n = 49). **b**, Venn diagram of mOFC response types as in **a**. Mean neural activity (novel odor epoch, ITI) as in Figure 3d for populations in the mOFC, replicated from previous figures, which were as follows: Left, Upshift-responsive (n = 21). Right, Downshift-responsive (n = 20). Top, Blocked-responsive (n = 11).

cells). We had found 1 negative value coding cell using our previous screen, screening first for a phasic odor response and subsequently for a significant difference in firing between blocked and both upshift and downshift firing. In lOFC, this same regression across all three reward amounts yielded 41 cells (6% of all recorded cells): 17 positive and 24 negative. We had found 8 cells with our previous screen, 5 positive and 3 negative. These relative proportions in mOFC and lOFC, 2.5% and 6%, are similar to the proportion of odor-responsive cells in the two regions: 13.5% and 25.3%, respectively, suggesting no difference in rank value coding over all cells.

Discussion

In a previous study (Lopatina et al., 2015), we tested the hypothesis that lOFC single-unit responses that accompany cues predictive of reward size underlie feature-based rather than value-based characteristics. We recorded in lOFC in an unblocking task in which novel cues predicted more, less, or no change in the amount of reward delivered. We reasoned that, if lOFC neurons signal the abstract or general value predicted by the cues after learning, then neurons that fired more to a cue predicting more reward should fire less to a cue predicting less reward, with firing to the cue predicting no change between the two. On the other hand, if lOFC neurons signaled specific reward features or even the particular associative representations (e.g., this odor predicts that larger reward), then neurons might respond relatively specifically to each cue. Our results in the lOFC were consistent with a firing pattern consistent with feature signaling. Only a very small population of neurons exhibited firing associated with general or abstract value signaling and much larger groups of neurons became selective for each individual cue with training. This included, interestingly, a group of neurons that developed selectivity for the odor cue that predicted no change in reward.

In this study, we repeated this experiment, recording in the mOFC. We did this due to the relative dearth of single-unit data from this area, particularly in rats (Burton et al., 2014), combined with the increasing prevalence of reports placing abstract or general value signaling in the medial rather than the lateral part of the OFC (Plassmann et al., 2007, 2010; Hare et al., 2008, 2010; Levy and Glimcher, 2011; McNamee et al., 2013; Strait et al., 2014). While these studies are in primate species, typically examining BOLD response in humans, we wondered whether our task might reveal similar pure value signaling in mOFC in rats. Similarities between mOFC in rats and these medial orbital areas in primates are supported by a recent study of frontostriatal anatomical homology, which compared projections from OFC to the nucleus accumbens shell, hippocampus, and amygdala in each species (Heilbronner et al., 2016). Despite this, we did not observe pure value signaling in rat mOFC in our task. Only a single neuron exhibited firing that was truly consistent with abstract value coding across the three cues. However, while groups of neurons became responsive to each individual cue with training, we also observed several important differences from the neural responses we observed in the lOFC. Overall, these differences suggest less activity related to specific features of the predicted rewards and more related to their relative valuation.

One difference was the overall proportion of odor-responsive cells. Only half as many of the recorded cells in the mOFC responded over baseline during the odor period relative to those in the lOFC. This difference alone suggests that mOFC attends less closely than lOFC to the idiosyncratic or physical properties of the cues and rewards in the task and the specific associations between them.

A second difference was that fewer odor-responsive neurons in the mOFC were responsive to the blocked odor cue, which

predicted no change in reward, than were in lOFC. In mOFC, the number of cells that responded to this cue was significantly lower than the number of neurons responsive to the cues that predicted a change in the amount of reward. In lOFC, there was no significant difference between the number of cells selectively responding to any of the three novel cues, nor was there a difference between the number of cells in the populations responsive to each cue. In lOFC, the blocked cue activated a population approximately similar to that activated by the other two cues. This dichotomy is particularly interesting because this cue predicts the absence of change in reward in a task in which change is otherwise pervasive; thus, in some regards, it might be said to have gained associative significance as a result of our procedure (i.e., it predicts something); however, it should not have gained much, if any, value because value was not formally changed on these trials. The dissociation in neural activity we observed in the two regions would therefore be consistent with representation of particular associative information in lOFC, independent of value, and representation of associative information more closely linked to value in mOFC.

A third difference was in the pattern of activity in the upshift- and downshift- responsive populations. In lOFC, these populations responded more to the cue used to identify them than the other two novel cues; thus, they were selective for each cue. For example, when we screened for upshift-responsive neurons, the neurons identified fired most to that cue. We found similar selectivity in the blocked and downshift-responsive populations. By contrast, in mOFC, only cells responsive to the downshift-predictive odor showed any specificity. This particular result in mOFC echoes the main result from the only other study of which we are aware that has looked at value-related signaling in rat mOFC. In that study, Burton et al. (2014) recorded during an odor choice task in which rats responded for differently valued rewards. The contingencies were reversed across blocks of trials, such that a cue and response that predicted a high value reward in one block subsequently predicted a low value reward in the next (and vice versa). They reported that cells in the mOFC responded most strongly to cues predictive of low value rewards in this setting (Burton et al., 2014). Our results suggest that it is not simply low value that is important in driving this response, but rather it is the change in value because the mOFC appears to be especially interested in cues that predict a value that has changed. This finding is particularly consistent with suggestions that mOFC plays a special role in representing current value to support devaluation/revaluation-sensitive changes in behavior. Overall, these differences are generally consistent with the hypothesis that mOFC encodes more highly processed, more abstract, and less idiosyncratic associative representations of value than lOFC.

References

Burton AC, Kashtelyan V, Bryden DW, Roesch MR (2014) Increased firing to cues that predict low-value reward in the medial orbitofrontal cortex. Cereb Cortex 24:3310–3321. CrossRef Medline

Combrisson E, Jerbi K (2015) Exceeding chance level by chance: the caveat of theoretical chance levels in brain signal classification and statistical assessment of decoding accuracy. J Neurosci Methods 250:126–136. CrossRef Medline

Delamater AR (2007) The role of the orbitofrontal cortex in sensory-specific encoding of associations in Pavlovian and instrumental conditioning. Ann N Y Acad Sci 1121:152–173. CrossRef Medline

Hare TA, O'Doherty J, Camerer CF, Schultz W, Rangel A (2008) Dissociating the role of the orbitofrontal cortex and the striatum in the computa-

tion of goal values and prediction errors. J Neurosci 28:5623–5630. CrossRef Medline

Hare TA, Camerer CF, Knoepfle DT, Rangel A (2010) Value computations in ventral medial prefrontal cortex during charitable decision making incorporate input from regions involved in social cognition. J Neurosci 30:583–590. CrossRef Medline

Heilbronner SR, Rodriguez-Romaguera J, Quirk GJ, Groenewegen HJ, Haber SN (2016) Circuit based cortico-striatal homologies between rat and primate. Biol Psychiatry, in press.

Jones JL, Esber GR, McDannald MA, Gruber AJ, Hernandez A, Mirenzi A, Schoenbaum G (2012) Orbitofrontal cortex supports behavior and learning using inferred but not cached values. Science 338:953–956. CrossRef Medline

Levy DJ, Glimcher PW (2011) Comparing apples and oranges: using reward-specific and reward-general subjective value representation in the brain. J Neurosci 31:14693–14707. CrossRef Medline

Lopatina N, McDannald MA, Styer CV, Sadacca BF, Cheer JF, Schoenbaum G (2015) Lateral orbitofrontal neurons acquire responses to upshifted, downshifted, or blocked cues during unblocking. Elife 4:e11299. CrossRef Medline

Luk CH, Wallis JD (2013) Choice coding in frontal cortex during stimulus-guided or action-guided decision-making. J Neurosci 33:1864–1871. CrossRef Medline

McDannald MA, Esber GR, Wegener MA, Wied HM, Liu TL, Stalnaker TA, Jones JL, Trageser J, Schoenbaum G (2014) Orbitofrontal neurons acquire responses to 'valueless' Pavlovian cues during unblocking. Elife 3:e02653. CrossRef Medline

McNamee D, Rangel A, O'Doherty JP (2013) Category-dependent and category-independent goal-value codes in human ventromedial prefrontal cortex. Nat Neurosci 16:479–485. CrossRef Medline

Noonan MP, Walton ME, Behrens TE, Sallet J, Buckley MJ, Rushworth MF (2010) Separate value comparison and learning mechanisms in macaque medial and lateral orbitofrontal cortex. Proc Natl Acad Sci U S A 107: 20547–20552. CrossRef Medline

Noonan MP, Kolling N, Walton ME, Rushworth MF (2012) Re-evaluating the role of the orbitofrontal cortex in reward and reinforcement. Eur J Neurosci 35:997–1010. CrossRef Medline

Ostlund SB, Balleine BW (2007) Orbitofrontal cortex mediates outcome encoding in Pavlovian but not instrumental conditioning. J Neurosci 27:4819–4825. CrossRef Medline

Padoa-Schioppa C (2009) Range-adapting representation of economic value in the orbitofrontal cortex. J Neurosci 29:14004–14014. CrossRef Medline

Padoa-Schioppa C (2013) Neuronal origins of choice variability in economic decisions. Neuron 80:1322–1336. CrossRef Medline

Padoa-Schioppa C, Assad JA (2006) Neurons in orbitofrontal cortex encode economic value. Nature 441:223–226. CrossRef Medline

Padoa-Schioppa C, Assad JA (2008) The representation of economic value in the orbitofrontal cortex is invariant for changes in menu. Nat Neurosci 11:95–102. CrossRef Medline

Pickens CL, Setlow B, Saddoris MP, 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. Medline

Plassmann H, O'Doherty J, Rangel A (2007) Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. J Neurosci 27: 9984–9988. CrossRef Medline

Plassmann H, O'Doherty JP, Rangel A (2010) Appetitive and aversive goal values are encoded in the medial orbitofrontal cortex at the time of decision making. J Neurosci 30:10799–10808. CrossRef Medline

Schoenbaum G, Chiba AA, Gallagher M (1998) Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. Nat Neurosci 1:155–159. CrossRef Medline

Steiner AP, Redish AD (2012) The road not taken: neural correlates of decision making in orbitofrontal cortex. Front Neurosci 6:131. CrossRef Medline

Strait CE, Blanchard TC, Hayden BY (2014) Reward value comparison via mutual inhibition in ventromedial prefrontal cortex. Neuron 82:1357–1366. CrossRef Medline

West EA, DesJardin JT, Gale K, Malkova L (2011) Transient inactivation of orbitofrontal cortex blocks reinforcer devaluation in macaques. J Neurosci 31:15128–15135. CrossRef Medline